



Complete Software Guide for Junos[®] OS for the QFX Series, Release 13.1

Release

13.1



Published: 2014-06-05

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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

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

Complete Software Guide for Junos® OS for the QFX Series, Release 13.1
Release 13.1
Copyright © 2014, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	cxxi
	Documentation and Release Notes	cxxi
	Supported Platforms	cxxi
	Using the Examples in This Manual	cxxi
	Merging a Full Example	cxxii
	Merging a Snippet	cxxii
	Documentation Conventions	cxxiii
	Documentation Feedback	cxxv
	Requesting Technical Support	cxxv
	Self-Help Online Tools and Resources	cxxv
	Opening a Case with JTAC	cxxvi
Part 1	QFX3500 Switch Overview	
Chapter 1	QFX3500 Switch Overview	3
	QFX3500 Device Overview	3
	Software	3
	Hardware	4
	SFP+ Access Ports	5
	QSFP+ Uplink Ports	6
Part 2	QFX3600 Switch Overview	
Chapter 2	QFX3600 Switch Overview	9
	QFX3600 Device Overview	9
	Software	9
	Hardware	10
Part 3	Software Feature Support	
Chapter 3	Software Feature Support for the QFX Series	15
	QFX Series Software Features Overview	15
Part 4	Junos OS Basics	
Chapter 4	Overview	33
	Software Overview	33
	QFX Series Software Features Overview	33
	Configuration File Terms	48
	Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos	
	OS Configuration Statements	49
	Junos OS Commit Model for Router or Switch Configuration	50

Junos OS Package Names	51
NTP Time Server and Time Services Overview	52
NTP Time Server and Time Services Overview (QFabric System)	53
Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)	53
Understanding Autoinstallation of Configuration Files	55
Typical Uses for Autoinstallation	55
Autoinstallation Configuration Files and IP Addresses	55
Typical Autoinstallation Process on a New Switch	56
Understanding Nonstop Software Upgrade for QFabric Systems	57
Understanding Zero Touch Provisioning	61
Understanding DHCP Services for Switches	63
DHCP Client/Server Model	64
Using DHCP	64
DHCP Relay Servers and DHCP Servers	65
Legacy DHCP and Extended DHCP for Server Versions	65
Configuring DHCP on a Switch	66
How DHCP Works	66
Understanding Software Infrastructure and Processes	67
Routing Engine and Packet Forwarding Engine	68
Junos OS Processes	68
User Interfaces	70
CLI User Interface Overview	70
CLI Overview	70
CLI Key Features	71
CLI Command Modes	71
Configuring CLI Tips	72
Format for Specifying Filenames and URLs in Junos OS CLI Commands	73
Junos OS Operational Mode Commands That Combine Other Commands	74
Overview of Junos OS CLI Operational Mode Commands	75
CLI Command Categories	75
Commonly Used Operational Mode Commands	76
Overview of Navigating the CLI	78
CLI Command Hierarchy	78
CLI Configuration Statements	78
Moving Among Hierarchy Levels	79
Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands	79
Understanding Junos OS CLI Configuration Mode	80
Configuration Mode Commands	81
Configuration Statements and Identifiers	82
Configuration Statement Hierarchy	84
Licenses	86
Junos OS Feature Licenses	86
Software Features That Require Licenses on the QFX Series	87

	Junos OS License Keys	88
	Licensable Ports on MX5, MX10, and MX40 Routers	88
	Generating the License Keys for a Standalone QFX Series Device	90
	Generating the License Keys for a QFabric System	91
	Adding New Licenses (CLI Procedure)	93
	Deleting a License (CLI Procedure)	94
	Saving License Keys	95
	Verifying Junos OS License Installation	96
	Displaying Installed Licenses	96
	Displaying License Usage	97
Chapter 5	Installation	99
	Software Installation	99
	Junos OS Package Names	99
	Configuring Zero Touch Provisioning	101
	Performing a Nonstop Software Upgrade on the QFabric System	105
	Backing Up the Current Configuration Files	106
	Downloading Software Files Using a Browser	107
	Retrieving Software Files for Download	108
	Performing a Nonstop Software Upgrade for Director Devices in a Director Group	108
	Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components	108
	(Optional) Creating Upgrade Groups for Node Groups	109
	Performing a Nonstop Software Upgrade on a Node Group	110
	Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device	111
	Performing a QFabric System Recovery Installation on the Director Group	112
	(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive	113
	Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software	115
	Recovering from a Failed Software Installation	119
	Software Installation Overview	120
	Upgrading Jloader Software on QFX Series Devices	120
	Jloader Software Version 1.1.4 Guidelines	122
	Upgrading Jloader Software on a QFX3500 Switch	123
	Upgrading Jloader Software on a QFabric System	126
	Upgrading Software on QFX3500 and QFX3600 Standalone Switches	132
	Downloading Software Files with a Browser	132
	Accessing Software Downloaded to a Remote Location	133
	Connecting to the Console Port	133
	Backing Up the Current Configuration Files	133
	Installing the Software Package	133
	Upgrading Software on a QFabric System	134
	Backing Up the Current Configuration Files	135
	Downloading Software Files Using a Browser	135
	Retrieving Software Files for Download	136

	Installing the Software Package on the Entire QFabric System	136
Chapter 6	Configuration	141
	Initial Configuration	141
	Configuring a DHCP Client (CLI Procedure)	142
	Configuring a DHCP Server on Switches (CLI Procedure)	143
	Configuring an Extended DHCP Server on a Switch	144
	Configuring a Legacy DHCP Server on a Switch (CLI Procedure)	145
	Configuring a DNS Name Server for Resolving a Hostname into Addresses	146
	Configuring the Domain Name for the Router or Switch	146
	Configuring the Domains to Search When a Router or Switch Is Included in Multiple Domains	147
	Configuring the Hostname of the Router or Switch	147
	Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types	147
	Configuring the Junos OS to Disable Protocol Redirect Messages on the Router or Switch	148
	Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses	148
	Configuring the Junos OS to Display a System Login Announcement	149
	Configuring the Junos OS to Display a System Login Message	149
	Configuring the Junos OS to Extend the Default Port Address Range	151
	Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages	151
	Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets	151
	Configuring NTP Authentication Keys	152
	Configuring NTP Authentication Keys (QFabric System)	153
	Configuring the NTP Time Server and Time Services	153
	Configuring the Router or Switch to Operate in Client Mode	153
	Configuring the Router or Switch to Operate in Symmetric Active Mode	154
	Configuring the Router or Switch to Operate in Broadcast Mode	154
	Configuring the Router or Switch to Operate in Server Mode	155
	Configuring the NTP Time Server and Time Services (QFabric System)	156
	Specifying the Physical Location of the Switch	156
	Configuring the Root Password	158
	Configuring the Router or Switch to Listen for Broadcast Messages Using NTP	160
	Configuring the Router or Switch to Listen for Multicast Messages Using NTP	160
	Configuring System Alarms to Appear Automatically Upon Login	161
	Configuring Time-Based User Access	161
	Configuring the Timeout Value for Idle Login Sessions	162
	Configuring a QFX3500 Device as a Standalone Switch	163
	Creating an Emergency Boot Device for a QFX Series Device	165

Creating a Snapshot and Using It to Boot a QFX Series Switch	166
Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch	166
Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch	167
Including the Year or Millisecond in Timestamps	167
Mapping the Hostname of the Switch to IP Addresses	168
Methods for Configuring Junos OS	169
Junos OS Command-Line Interface (CLI)	170
ASCII File	170
J-Web Package	170
Junos XML Management Protocol Software	171
NETCONF XML Management Protocol Software	171
Configuration Commit Scripts	171
Modifying the Default Time Zone for a Router or Switch Running Junos OS	172
Rebooting and Halting a QFX Series Product	172
Reverting to the Default Factory Configuration	173
Reverting to the Default Factory Configuration by Using the request system zeroize Command	174
Reverting to the Rescue Configuration	175
Saving Core Files Generated by Junos OS Processes	175
Setting a Custom Time Zone on Routers or Switches Running Junos OS	175
Importing and Installing Time Zone Files	176
Configuring a Custom Time Zone	177
Setting the Date and Time	177
Specifying Access Privileges for Junos OS Operational Mode Commands	178
Synchronizing and Coordinating Time Distribution Using NTP	180
Configuring NTP	180
Configuring the NTP Boot Server	180
Specifying a Source Address for an NTP Server	181
Viewing Core Files from Junos OS Processes	181
Configuration Examples	182
Example: Changing the Requirements for Junos OS Plain-Text Passwords	182
Example: Configuring the Domain Name for the Router or Switch	184
Example: Configuring the Name of the Switch, IP Address, and System ID	184
Example: Configuring NTP	185
Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization	187
Example: Configuring a Plain-Text Password for Root Logins	188
Example: Configuring the Root Password	189
Configuration Statements	190
QFX Series CLI Hierarchy	192
[edit access] Hierarchy	193
[edit accounting-options] Hierarchy	193
[edit chassis] Hierarchy	195

[edit class-of-service] Hierarchy	196
[edit ethernet-switching-options] Hierarchy	199
[edit fabric] Hierarchy	200
[edit fc-fabrics] Hierarchy	201
[edit fc-options] Hierarchy	202
[edit firewall] Hierarchy	202
[edit groups] Hierarchy	203
[edit interfaces] Hierarchy	203
[edit policy-options] Hierarchy	209
[edit protocols] Hierarchy	210
[edit security] Hierarchy	222
[edit snmp] Hierarchy	222
[edit system] Hierarchy	226
[edit vlans] Hierarchy	231
access-end	232
access-start	232
accounting	233
accounting-port	234
allow-commands	234
allow-configuration	235
allowed-days	235
allow-transients	236
announcement	236
archival	237
arp (System)	238
authentication (Login)	239
authentication-key	240
authentication-order	241
auxiliary	242
boot-server (NTP)	243
broadcast	244
broadcast-client	245
change-type	245
checksum	246
class (Defining Login Classes)	247
class (Assigning a Class to an Individual User)	248
commit	249
compress-configuration-files (System)	250
console (Physical Port)	251
default-address-selection	252
deny-commands	253
deny-configuration	254
destination (Accounting)	255
destination-override	256
direct-access	256
domain-name	257
domain-search	257
explicit-priority	258
events	259

format	259
host-name	260
icmpv4-rate-limit	260
idle-timeout	261
internet-options	261
load-key-file	262
location	263
login	264
login-alarms	265
login-tip	265
max-configurations-on-flash	266
maximum-length	266
message	267
minimum-changes	267
minimum-length	268
minimum-lower-cases	269
minimum-numeric	270
minimum-punctuations	271
minimum-upper-cases	272
multicast-client	272
name-server	273
no-multicast-echo	274
no-ping-record-route	275
no-ping-time-stamp	275
no-redirects (IPv4 Traffic)	276
ntp	277
ntp (QFabric)	277
optional	278
password (Login)	278
peer	279
permissions	280
port (TACACS+ Server)	280
ports	281
radius (System)	282
refresh (Commit Scripts)	283
refresh-from (Commit Scripts)	283
retry	284
retry-options	285
root-authentication	286
saved-core-context	287
saved-core-files	287
secret	288
server (NTP)	289
server (RADIUS Accounting)	290
server (TACACS+ Accounting)	290
single-connection	291
source (Commit Scripts)	291
source-address (NTP, RADIUS, System Logging, or TACACS+)	292
source-port (Port Addresses)	293

	ssh-dsa	293
	ssh-rsa	294
	static-host-mapping	295
	structured-data	296
	syslog (System)	297
	system	299
	tacplus	304
	tacplus-server	305
	time-format	306
	timeout	307
	time-zone	308
	traceoptions (Commit Scripts)	310
	tracing	312
	trusted-key	313
	uid	313
	use-imported-time-zones	314
	user (Access)	314
Chapter 7	Administration	315
	Routine Monitoring	315
	Monitoring System Process Information	315
	Monitoring System Properties	316
	Monitoring Interface Status and Traffic	317
	Other Tools to Configure and Monitor Devices Running Junos OS	318
	Operational Commands	318
	commit	322
	clear log	327
	clear chassis display message	328
	clear system commit	331
	clear system reboot	332
	file	336
	file archive	337
	file checksum md5	339
	file checksum sha1	340
	file checksum sha-256	341
	file compare	342
	file delete	345
	file list	346
	file rename	348
	file show	350
	load	352
	ping	354
	request chassis beacon	358
	request chassis cb	360
	request chassis fpc	363
	request chassis routing-engine master	367
	request message	372
	request system configuration rescue delete	373
	request system configuration rescue save	374

request system halt	375
request system license add	381
request system license delete	382
request system license save	383
request system logout	384
request system power-off	385
request system reboot	390
request system snapshot	393
request system software add	394
request system software configuration-backup	402
request system software configuration-restore	403
request system software delete	404
request system software download	408
request system software nonstop-upgrade	410
request system software rollback	416
request system software validate	420
request system storage cleanup	423
request system zeroize	433
restart	438
rollback	449
save	450
show chassis alarms	452
show chassis beacon	464
show chassis ethernet-switch interconnect-device fpc	466
show chassis ethernet-switch interconnect-device cb	491
show chassis environment	508
show chassis environment cb	561
show chassis environment fpc	579
show chassis environment pem	604
show chassis environment routing-engine	613
show chassis fan	618
show chassis firmware	631
show chassis fpc	641
show chassis hardware	667
show chassis lcd	785
show chassis led	798
show chassis location	808
show chassis mac-addresses	812
show chassis nonstop-upgrade node-group	817
show chassis pic	818
show chassis routing-engine	831
show chassis temperature-thresholds	852
show chassis zones	868
show cli	874
show cli authorization	876
show cli directory	880
show cli history	881
show host	882
show interfaces diagnostics optics	883

	show log	889
	show ntp associations	892
	show ntp status	894
	show subscribers	897
	show system alarms	914
	show system audit	915
	show system boot-messages	923
	show system buffers	930
	show system certificate	937
	show system commit	939
	show system configuration archival	941
	show system configuration rescue	942
	show system connections	944
	show system core-dumps	963
	show system directory-usage	976
	show system license	980
	show system processes	983
	show system reboot	1010
	show system resource-cleanup processes	1013
	show system rollback	1015
	show system services service-deployment	1017
	show system software	1018
	show system statistics	1026
	show system storage	1061
	show system uptime	1067
	show system users	1072
	show system virtual-memory	1077
	show version	1135
	start shell	1148
	test configuration	1150
	traceroute	1151
	traceroute monitor	1155
Chapter 8	Troubleshooting	1157
	Troubleshooting Procedures	1157
	Rebooting and Halting a QFX Series Product	1157
	Recovering from a Failed Software Installation	1158
	Recovering the Root Password	1159
	Troubleshooting Network Interfaces	1160
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down	1160
	Troubleshooting an Aggregated Ethernet Interface	1161

Part 5	QFabric System Deployment	
Chapter 9	Overview	1165
	Before You Begin	1165
	QFabric System Overview	1165
	Legacy Data Center Architecture	1165
	QFX Series QFabric System Architecture	1167
	Understanding QFabric System Terminology	1169
	Understanding Interfaces on the QFabric System	1173
	Four-Level Interface Naming Convention	1173
	QSFP+ Interfaces	1174
	Link Aggregation	1176
	Hardware Architecture Overview	1177
	Understanding the QFabric System Hardware Architecture	1177
	QFabric System Hardware Architecture Overview	1177
	QFX3000-G QFabric System Features	1180
	QFX3000-M QFabric System Features	1180
	Understanding the Director Group	1180
	Director Group Components	1180
	Director Group Services	1181
	Understanding Routing Engines in the QFabric System	1181
	Hardware-Based Routing Engines	1182
	Software-Based External Routing Engines	1182
	Understanding Interconnect Devices	1183
	Interconnect Device Introduction	1183
	QFX3008-I Interconnect Devices	1184
	QFX3600-I Interconnect Devices	1185
	Understanding Node Devices	1186
	Node Device Introduction	1186
	QFX3500 Node Devices	1187
	QFX3600 Node Devices	1188
	Understanding Node Groups	1190
	Network Node Groups	1190
	Server Node Groups	1190
	Understanding Port Oversubscription on Node Devices	1191
	Software Architecture Overview	1192
	Understanding the QFabric System Software Architecture	1192
	Understanding the Director Software	1193
	Understanding Partitions	1194
	QFabric System Default Partition	1194
	Understanding the QFabric System Control Plane	1196
	Control Plane Elements	1197
	Control Plane Services	1199
	Understanding the QFabric System Data Plane	1199
	Data Plane Components	1199
	QFabric System Fabric	1200

Software Features	1201
QFX Series Software Features Overview	1202
Understanding Software Upgrade on the QFabric System	1217
Operational Software Commands	1217
Operational Reboot Commands	1218
Understanding Nonstop Software Upgrade for QFabric Systems	1218
Understanding Statements and Commands on the QFabric System	1223
Chassis Statements	1223
Chassis Commands	1223
Understanding NTP on the QFabric System	1224
Understanding Network Management Implementation on the QFabric System	1224
Understanding the Implementation of SNMP on the QFabric System	1225
Understanding the Implementation of System Log Messages on the QFabric System	1227
Understanding User and Access Management Features on the QFabric System	1229
Understanding QFabric System Login Classes	1230
Understanding Interfaces on the QFabric System	1231
Four-Level Interface Naming Convention	1231
QSFP+ Interfaces	1231
Link Aggregation	1234
Understanding Layer 3 Features on the QFabric System	1234
Understanding Security Features on the QFabric System	1235
Understanding Port Mirroring on the QFabric System	1236
Understanding Fibre Channel Fabrics on the QFabric System	1236
Understanding CoS Fabric Forwarding Class Sets	1238
Default Fabric Forwarding Class Sets	1239
Fabric Forwarding Class Set Configuration and Implementation	1242
Fabric Forwarding Class Set Scheduling (CoS)	1244
Support for Flow Control and Lossless Transport Across the Fabric	1246
Viewing Fabric Forwarding Class Set Information	1248
Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences	1250
Licenses	1250
Junos OS Feature Licenses	1250
Software Features That Require Licenses on the QFX Series	1251
Junos OS License Keys	1253
Licensable Ports on MX5, MX10, and MX40 Routers	1253

Chapter 10	Configuration	1255
	Initial Setup	1255
	QFabric System Initial and Default Configuration Information	1255
	Converting the Device Mode for a QFabric System Component	1258
	Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane	1263
	Importing a QFX3000-G QFabric System Control Plane Virtual Chassis Configuration with a USB Flash Drive	1308
	Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane	1309
	Importing a QFX3000-M QFabric System Control Plane EX4200 Switch Configuration with a USB Flash Drive	1333
	Generating the MAC Address Range for a QFabric System	1334
	Performing the QFabric System Initial Setup on a QFX3100 Director Group	1335
	Performing an Initial Setup	1336
	Restoring a Backup Configuration	1339
	QFabric System Configuration	1340
	Understanding QFabric System Administration Tasks and Utilities	1340
	Gaining Access to the QFabric System Through the Default Partition	1344
	Example: Configuring QFabric System Login Classes	1345
	Configuring Aliases for the QFabric System	1353
	Configuring Node Groups for the QFabric System	1363
	Configuring the Port Type on QFX3600 Node Devices	1367
	Example: Configuring SNMP	1370
	Example: Configuring System Log Messages	1372
	Configuring Graceful Restart for QFabric Systems	1375
	Enabling Graceful Restart	1375
	Configuring Graceful Restart Options for BGP	1376
	Configuring Graceful Restart Options for OSPF and OSPFv3	1377
	Tracking Graceful Restart Events	1378
	Configuration Statements	1378
	aliases	1380
	archive (QFabric System)	1381
	chassis (QFabric System)	1382
	director-device (Aliases)	1384
	fabric	1385
	file (QFabric System)	1386
	graceful-restart (Enabling Globally)	1387
	graceful-restart (Protocols BGP)	1388
	graceful-restart (Protocols OSPF)	1389
	interconnect-device (Chassis)	1391
	interconnect-device (Aliases)	1392
	multicast (QFabric Routing Options)	1393
	network-domain	1393
	no-make-before-break	1394
	node-device (Aliases)	1395
	node-device (Chassis)	1396
	node-device (Resources)	1397

	node-group (Chassis)	1398
	node-group (Resources)	1399
	pic (Port)	1400
	remote-debug-permission	1401
	resources	1402
	routing-options (QFabric System)	1402
	syslog (QFabric System)	1403
Chapter 11	Administration	1405
	Software Upgrade and Recovery	1405
	Performing a Nonstop Software Upgrade on the QFabric System	1405
	Backing Up the Current Configuration Files	1407
	Downloading Software Files Using a Browser	1407
	Retrieving Software Files for Download	1408
	Performing a Nonstop Software Upgrade for Director Devices in a Director Group	1408
	Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components	1408
	(Optional) Creating Upgrade Groups for Node Groups	1409
	Performing a Nonstop Software Upgrade on a Node Group	1410
	Verifying Nonstop Software Upgrade for QFabric Systems	1410
	Verifying a Director Group Nonstop Software Upgrade	1411
	Verifying a Fabric Nonstop Software Upgrade	1423
	Verifying a Redundant Server Node Group Nonstop Software Upgrade	1425
	Verifying a Network Node Group Nonstop Software Upgrade	1428
	Upgrading Software on a QFabric System	1430
	Backing Up the Current Configuration Files	1430
	Downloading Software Files Using a Browser	1431
	Retrieving Software Files for Download	1432
	Installing the Software Package on the Entire QFabric System	1432
	Performing System Backup and Recovery for a QFabric System	1435
	Operational Mode Commands	1436
	QFabric System Operational Mode Commands	1437
	Filtering Operational Mode Command Output in a QFabric System	1439
	request chassis device-mode	1441
	request chassis fabric fpc	1443
	request component login	1444
	request fabric administration director-group change-master	1446
	request fabric administration remove	1447
	request fabric administration system mac-pool add	1449
	request fabric administration system mac-pool delete	1450
	request system halt	1451
	request system reboot	1457
	request system software format-qfabric-backup	1460
	request system software nonstop-upgrade	1461
	request system software system-backup	1467
	set chassis display message	1469
	show chassis device-mode	1475

	show chassis ethernet-switch interconnect-device cb	1478
	show chassis ethernet-switch interconnect-device fpc	1495
	show chassis fabric connectivity	1520
	show chassis fabric device	1527
	show chassis lcd	1529
	show chassis nonstop-upgrade node-group	1542
	show fabric administration inventory	1543
	show fabric administration inventory director-group status	1548
	show fabric administration inventory infrastructure	1553
	show fabric administration inventory interconnect-devices	1556
	show fabric administration inventory node-devices	1558
	show fabric administration inventory node-groups	1560
	show fabric administration system mac-pool	1562
	show fabric inventory	1563
	show fabric session-host	1566
	show log	1567
	show system software upgrade status	1570
Chapter 12	Troubleshooting	1573
	QFabric System Troubleshooting	1573
	Performing System Backup and Recovery for a QFabric System	1573
	Performing a QFabric System Recovery Installation on the Director Group	1574
	(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive	1575
	Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software	1577
	Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device	1581
	Creating an Emergency Boot Device for a QFX Series Device	1583
Part 6	Configuration and File Management	
Chapter 13	Overview	1587
	Configuration Files Overview	1587
	Configuration File Terms	1587
	Software Overview	1588
	Forms of the configure Command	1588
	Junos OS Commit Model for Router or Switch Configuration	1589
	Understanding Configuration Files	1590
	Understanding How the Junos Configuration Is Stored	1591
Chapter 14	Configuration	1593
	Configuration Tasks	1593
	Comparing Configuration Changes with a Prior Version	1593
	Compressing the Current Configuration File	1595
	Creating and Returning to a Rescue Configuration	1596
	Loading a Configuration from a File	1597
	Loading a Previous Configuration File	1600

	Returning to the Most Recently Committed Junos Configuration	1600
	Returning to a Previously Committed Junos OS Configuration	1601
	Returning to a Configuration Prior to the One Most Recently Committed	1601
	Displaying Previous Configurations	1601
	Comparing Configuration Changes with a Prior Version	1602
	Creating and Returning to a Rescue Configuration	1604
	Saving a Configuration to a File	1605
	Reverting to the Default Factory Configuration	1606
	Reverting to the Rescue Configuration	1606
	Rolling Back Junos OS Configuration Changes	1607
	Saving a Configuration to a File	1608
	Setting or Deleting the Rescue Configuration	1609
	Uploading a Configuration File	1609
	Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site	1611
	Configuring the Router or Switch to Transfer Its Currently Active Configuration to an Archive	1611
	Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site	1611
	Configuring Transfer of the Current Active Configuration When a Configuration Is Committed	1612
	Configuring Archive Sites for Transfer of Active Configuration Files	1612
	Configuration Statements	1613
	archival	1614
	archive-sites (Configuration File)	1615
	configuration	1617
	transfer-interval (Configuration)	1618
	transfer-on-commit	1619
	Default Configurations	1619
	QFX3500 Switch Default Configuration	1619
	Configuration Examples	1625
	Examples: Loading a Configuration from a File	1625
Chapter 15	Administration	1629
	Operational Commands	1629
	clear log	1630
	clear system commit	1631
	file archive	1632
	file checksum md5	1634
	file checksum sha1	1635
	file checksum sha-256	1636
	file compare	1637
	file delete	1640
	file list	1641
	file rename	1643
	file show	1645
	request system configuration rescue delete	1647
	request system configuration rescue save	1648

	show system commit	1649
	show system configuration archival	1651
	show system configuration rescue	1652
	show system rollback	1654
	test configuration	1656
Chapter 16	Troubleshooting	1657
	Troubleshooting Procedures	1657
	Loading a Previous Configuration File	1657
	Reverting to the Default Factory Configuration	1658
	Reverting to the Rescue Configuration	1658
Part 7	User and Access Management	
Chapter 17	Overview	1661
	Software Overview	1661
	Understanding Software Infrastructure and Processes	1661
	Routing Engine and Packet Forwarding Engine	1661
	Junos OS Processes	1662
	Understanding User and Access Management Features on the QFabric System	1663
	Access Control Overview	1664
	Overview of Template Accounts for RADIUS and TACACS+ Authentication	1665
	Understanding Login Authentication	1665
	MAC RADIUS Authentication	1665
	Understanding LLDP	1666
	Understanding RADIUS Accounting	1667
	Understanding VSAs on the QFX Series	1668
	Juniper Networks Vendor-Specific RADIUS Attributes	1668
	Juniper Networks Vendor-Specific TACACS+ Attributes	1670
	Understanding Junos OS Access Privilege Levels	1672
	Junos OS Login Class Permission Flags	1672
	Allowing or Denying Individual Commands for Junos OS Login Classes	1675
	Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication	1676
	Using RADIUS or TACACS+ Authentication	1676
	Using Local Password Authentication	1677
	Order of Authentication Attempts	1677
	Junos OS User Authentication Methods	1681
	Junos OS User Accounts Overview	1681
	Junos OS Login Classes Overview	1683
	Understanding QFabric System Login Classes	1684
	Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies	1685
	Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands	1686
	Special Requirements for Junos OS Plain-Text Passwords	1687

Chapter 18	Configuration	1691
	Configuration Tasks	1691
	Configuring Access Privilege Levels	1692
	Configuring CLI Tips	1692
	Configuring Junos OS User Accounts	1692
	Configuring LLDP	1693
	Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication	1694
	Configuring Local User Template Accounts for User Authentication	1695
	Configuring Management Access	1697
	Configuring RADIUS System Accounting	1697
	Configuring Auditing of User Events on a RADIUS Server	1697
	Specifying RADIUS Server Accounting and Auditing Events	1697
	Configuring RADIUS Server Accounting	1698
	Configuring RADIUS Authentication	1699
	Configuring RADIUS Server Details	1699
	Configuring MS-CHAPv2 for Password-Change Support	1700
	Specifying a Source Address for the Junos OS to Access External RADIUS Servers	1701
	Configuring Remote Template Accounts for User Authentication	1701
	Configuring the Root Password	1702
	Configuring SNMP	1703
	Configuring SSH Host Keys for Secure Copying of Data	1707
	Configuring SSH Known Hosts	1708
	Configuring Support for SCP File Transfer	1708
	Updating SSH Host Key Information	1709
	Configuring SSH Service for Remote Access to the Router or Switch	1709
	Configuring the Root Login Through SSH	1710
	Configuring the SSH Protocol Version	1710
	Configuring the Client Alive Mechanism	1711
	Configuring TACACS+ Authentication	1711
	Configuring TACACS+ Server Details	1711
	Specifying a Source Address for the Junos OS to Access External TACACS+ Servers	1712
	Configuring the Same Authentication Service for Multiple TACACS+ Servers	1713
	Configuring Juniper Networks Vendor-Specific TACACS+ Attributes	1713
	Configuring TACACS+ System Accounting	1714
	Specifying TACACS+ Auditing and Accounting Events	1714
	Configuring TACACS+ Server Accounting	1715
	Defining Junos OS Login Classes	1716
	Limiting the Number of User Login Attempts for SSH and Telnet Sessions	1716
	Recovering the Root Password	1717
	Specifying Access Privileges for Junos OS Configuration Mode Hierarchies	1719
	Specifying Access Privileges for Junos OS Operational Mode Commands	1720

Using Junos OS to Configure Logical System Administrators	1721
Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands	1722
VSA Match Conditions and Actions	1724
Configuration Examples	1726
Example: Changing the Requirements for Junos OS Plain-Text Passwords	1727
Example: Configuring Access Privilege Levels	1729
Example: Configuring Access Privileges for Operational Mode Commands	1729
Example: Configuring a Plain-Text Password for Root Logins	1730
Example: Configuring RADIUS Authentication	1730
Example: Configuring RADIUS Authentication on a QFabric System	1731
Example: Configuring RADIUS System Accounting	1732
Example: Configuring the Root Password	1733
Example: Configuring SSH Authentication for Root Logins	1733
Example: Configuring User Accounts	1733
Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication	1734
Example: Creating Login Classes with Specific Privileges	1736
Example: Configuring QFabric System Login Classes	1737
Example: Configuring User Login Accounts	1744
Example: Configuring RADIUS Template Accounts	1745
Defining Access Privileges Using allow/deny-configuration Statements . .	1745
Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions	1746
Configuration Statements	1747
access	1749
accounting	1750
accounting-options	1751
accounting-server	1753
accounting-stop-on-access-deny	1754
accounting-stop-on-failure	1755
advertisement-interval	1756
agent-address	1757
archival	1758
archive-sites (Configuration File)	1759
authentication-order	1760
authentication-server	1761
authorization	1762
categories	1763
client-list	1763
client-list-name	1764
clients	1764
commit-delay	1765
community (SNMP)	1766
configuration	1767
connection-limit	1768
contact	1769

disable (LLDP)	1769
ethernet-switching-options	1770
falling-threshold (Health Monitor)	1772
filter-duplicates	1772
full-name	1773
health-monitor	1773
hold-multiplier	1774
idle-timeout (Access)	1775
interface (LLDP)	1776
interval (Health Monitor)	1777
lldp	1778
lldp-configuration-notification-interval	1779
location	1779
management-address	1780
name	1780
nas-ip-address	1781
nonvolatile	1781
oid	1782
order	1783
port (RADIUS Server)	1784
profile	1785
protocols	1786
protocol-version	1799
ptopo-configuration-maximum-hold-time	1799
ptopo-configuration-trap-interval	1800
radius	1801
radius-options (edit system)	1802
radius-server	1803
rate-limit	1804
remote-debug-permission	1805
retry	1806
rising-threshold (Health Monitor)	1807
root-login	1808
services (Switches)	1809
snmp	1810
ssh	1814
system	1815
tacplus-options	1821
targets	1822
traceoptions (LLDP)	1823
transfer-interval (Configuration)	1825
transfer-on-commit	1826
trap-group	1827
trap-options	1828
user (Access)	1829
version	1830

Chapter 19	Administration	1831
	Routine Monitoring	1831
	Monitoring SNMP	1831
	Monitoring Commands	1832
	clear lldp neighbors	1834
	clear lldp statistics	1835
	request component login	1836
	show ethernet-switching interfaces	1838
	show lldp	1842
	show lldp local-information	1847
	show lldp neighbors	1849
	show lldp statistics	1853
	show route instance	1855
	show snmp statistics	1859
	ssh	1863
Part 8	Ethernet Features	
Chapter 20	Overview	1867
	Software Features Overview	1867
	Overview of Layer 2 Networking	1867
	Understanding Bridging	1869
	Understanding Unicast	1871
	Introduction to the Media Access Control (MAC) Layer 2 Sublayer	1871
	Understanding Layer 2 Broadcasting	1872
	Bridging and VLANs	1873
	Understanding VLANs	1873
	Introduction to VLANs	1873
	VLAN Tagging	1874
	Assigning Traffic to a VLAN	1874
	Ethernet Switching Tables	1874
	Layer 2 and Layer 3 Forwarding of VLAN Traffic	1874
	Understanding Routed VLAN Interfaces	1875
	Understanding MAC Learning	1876
	Understanding Reflective Relay for Use with VEPA Technology	1877
	VEPA	1877
	Reflective Relay	1877
	Understanding Private VLANs	1878
	Typical Structure and Primary Application of PVLANS	1878
	Using 802.1Q Tags to Identify Packets	1880
	Efficient Use of IP Addresses	1881
	PVLAN Port Types	1881
	Limitations of Private VLANs	1883
	Understanding PVLAN Traffic Flows Across Multiple Switches	1883
	Community VLAN Sending Untagged Traffic	1883
	Isolated VLAN Sending Untagged Traffic	1884

PVLAN Tagged Traffic Sent on a Promiscuous Port	1885
Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS	1887
PVLAN Port Types	1887
Secondary VLAN Trunk Port Details	1888
Use Cases	1889
Understanding Egress Firewall Filters with PVLANS	1896
Understanding Multiple VLAN Registration Protocol (MVRP)	1897
QFabric Requirements	1897
MVRP Operations	1898
MRP Timers Control MVRP Updates	1898
MVRP Uses MRP Messages to Transmit Switch and VLAN States	1899
Spanning Trees Overview	1899
Overview of Spanning-Tree Protocols	1900
Understanding Spanning Tree Protocols on a QFabric System	1900
Understanding MSTP	1901
Understanding RSTP	1901
Understanding VSTP	1902
Understanding BPDU Protection for STP, RSTP, and MSTP	1903
Understanding Loop Protection for STP, RSTP, VSTP, and MSTP	1904
Understanding Root Protection for STP, RSTP, VSTP, and MSTP	1905
Q-in-Q Tunneling	1906
Understanding Q-in-Q Tunneling and VLAN Translation	1906
How Q-in-Q Tunneling Works	1907
How VLAN Translation Works	1908
Mapping C-VLANs to S-VLANs	1908
Routed VLAN Interfaces on Q-in-Q VLANs	1909
Constraints for Q-in-Q Tunneling and VLAN Translation	1909
Layer 2 Protocol Tunneling	1910
Understanding Layer 2 Protocol Tunneling	1910
Layer 2 Protocols Supported by L2PT	1911
How L2PT Works	1911
L2PT Basics	1913
Proxy ARP	1914
Understanding Proxy ARP	1914
What Is ARP?	1914
Proxy ARP Overview	1915
Best Practices for Proxy ARP	1915
Chapter 21 Configuration	1917
Configuration Examples	1917
Example: Configuring Faster Convergence and Improving Network Stability with RSTP	1918
Example: Configuring Network Regions for VLANs with MSTP	1931
Example: Configuring Automatic VLAN Administration Using MVRP	1953
Example: Configuring Reflective Relay for Use with VEPA Technology	1958
Example: Setting Up Basic Bridging and a VLAN on the QFX Series	1962
Example: Setting Up Bridging with Multiple VLANs	1971
Example: Configuring Routing Between VLANs on One Switch	1976

Example: Connecting an Access Switch to a Distribution Switch	1982
Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations	1990
Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees	1994
Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree	1999
Example: Disabling MAC Learning	2003
Example: Disabling MAC Learning in a VLAN	2003
Private VLAN Configuration Examples	2005
Example: Configuring a Private VLAN on a Single Switch	2005
Example: Configuring a Private VLAN Spanning Multiple Switches	2010
Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports	2024
Q-in-Q Tunneling Configuration Example	2036
Example: Setting Up Q-in-Q Tunneling	2036
Configuration Tasks	2039
Configuring the Interface Address	2040
Configuring Interface IPv4 Addresses	2041
Configuring Interface IPv6 Addresses	2041
Configuring MAC Notification	2042
Enabling MAC Notification	2043
Disabling MAC Notification	2043
Setting the MAC Notification Interval	2043
Configuring MAC Table Aging	2044
Configuring Reflective Relay	2044
Configuring Static ARP Entries	2045
Configuring Reflective Relay	2045
Configuring Routed VLAN Interfaces	2046
Configuring the Native VLAN Identifier	2047
Configuring VLANs	2048
Creating a Series of Tagged VLANs	2050
Configuring Multiple VLAN Registration Protocol	2051
Enabling MVRP	2051
Disabling MVRP	2051
Configuring Timer Values	2052
Configuring Passive Mode	2052
Configuring STP	2053
Unblocking an Interface That Receives BPDUs in Error	2054
Configuring VLAN Spanning Tree Protocol	2054
Disabling MAC Learning	2055
Disabling MAC Learning in a VLAN	2056
Creating a Private VLAN on a Single Switch	2057
Creating a Private VLAN Spanning Multiple Switches	2058
Configuring the Forwarding Mode	2059
Private VLAN Configuration Tasks	2059
Creating a Private VLAN on a Single Switch	2060
Creating a Private VLAN Spanning Multiple Switches	2061

Q-in-Q Tunneling Configuration Tasks	2062
Configuring Q-in-Q Tunneling	2062
Configuring Layer 2 Protocol Tunneling	2063
Proxy ARP Configuration Task	2065
Configuring Proxy ARP	2065
Configuration Statements	2066
address	2069
alarm (STP)	2071
arp (Interfaces)	2072
block	2073
bpdu-block	2074
bpdu-block-on-edge	2075
bpdu-timeout-action	2076
bridge-priority	2077
broadcast	2078
cost (STP)	2079
configuration-name (MSTP)	2080
cut-through	2080
description (VLAN)	2081
disable (STP)	2082
disable-timeout (BPDU)	2083
drop-threshold	2084
edge (STP)	2085
ethernet-switching-options	2086
eui-64	2088
filter (VLANs)	2088
force-version	2089
forward-delay	2090
group-limit	2091
hello-time	2092
interface (BPDU)	2093
interface (Spanning Trees)	2094
interface (Storm Control)	2095
interface (STP)	2096
interface (VLANs)	2097
interfaces	2097
l3-interface (VLAN)	2098
mac-limit	2099
mac-notification	2100
mac-table-aging-time	2101
max-age	2102
max-hops	2103
members	2104
mode (STP)	2105
msti	2106
mstp	2107
native-vlan-id	2108
no-mac-learning (Global)	2109
no-mac-learning (Per VLAN)	2109

no-root-port	2110
notification-interval	2111
port-mode	2112
preferred	2113
primary (Address on Interface)	2114
priority (STP)	2115
protocols	2116
proxy-arp	2129
reflective-relay	2130
revision-level	2130
rstp	2131
storm-control	2132
stp	2133
traceoptions (Ethernet Switching Options)	2134
traceoptions (STP)	2137
unknown-unicast-forwarding	2140
vlan (Ethernet)	2141
vlan (STP)	2142
vlan (Unknown Unicast)	2143
vlan-id (VLANs)	2144
vlan-range	2144
vlangs	2145
vlan-tagging	2146
vstp	2147
MVRP Configuration Statements	2148
disable (MVRP)	2148
interface (MVRP)	2149
join-timer (MVRP)	2150
leave-timer (MVRP)	2151
leaveall-timer (MVRP)	2152
passive (MVRP)	2153
Private VLAN Configuration Statements	2153
extend-secondary-vlan-id	2154
isolated	2154
isolation-vlan-id	2155
primary-vlan	2155
pvlan	2156
promiscuous	2156
pvlan-trunk	2157
vlangs	2158
Q-in-Q Tunneling Configuration Statements	2159
customer-vlangs	2160
dot1q-tunneling (Ethernet Switching)	2161
dot1q-tunneling (VLANs)	2162
ether-type	2163
layer2-protocol-tunneling	2164
mapping	2166
mapping-range	2167
no-local-switching	2167

	shutdown-threshold	2168
	vlan-id-start	2169
	vlangs	2170
Chapter 22	Administration	2173
	Routine Monitoring	2173
	Verifying That MAC Notification Is Working Properly	2173
	Verifying That a Series of Tagged VLANs Has Been Created	2173
	Verifying That Q-in-Q Tunneling Is Working	2175
	Verifying That a Private VLAN Is Working	2176
	Verifying That Proxy ARP Is Working Correctly	2181
	Verifying That MVRP Is Working Correctly	2182
	Monitoring Commands	2183
	clear ethernet-switching bpd-error	2185
	clear ethernet-switching layer2-protocol-tunneling error	2186
	clear ethernet-switching layer2-protocol-tunneling statistics	2187
	clear ethernet-switching table	2188
	clear spanning-tree statistics	2190
	show ethernet-switching interfaces	2191
	show ethernet-switching layer2-protocol-tunneling interface	2195
	show ethernet-switching layer2-protocol-tunneling statistics	2197
	show ethernet-switching layer2-protocol-tunneling vlan	2200
	show ethernet-switching mac-learning-log	2202
	show ethernet-switching mac-notification	2204
	show ethernet-switching statistics aging	2205
	show ethernet-switching statistics mac-learning	2207
	show ethernet-switching table	2211
	show interfaces xe	2217
	show spanning-tree bridge	2235
	show spanning-tree interface	2240
	show spanning-tree mstp configuration	2246
	show spanning-tree statistics	2248
	show system statistics arp	2250
	show vlans	2251
Chapter 23	Troubleshooting	2261
	Troubleshooting Procedures	2261
	Troubleshooting Ethernet Switching	2261
	Troubleshooting Layer 2 Protocol Tunneling	2262
	Drop Threshold Statistics Might Be Incorrect	2262
	Egress Filtering of L2PT Traffic Not Supported	2262
	Troubleshooting Private VLANs	2263
	Limitations of Private VLANs	2263
	Forwarding with Private VLANs	2263
	Egress Firewall Filters with Private VLANs	2264
	Egress Port Mirroring with Private VLANs	2265
	Troubleshooting Q-in-Q and VLAN Translation Configuration	2266
	Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	2266
	Egress Port Mirroring with VLAN Translation	2266

Part 9	High Availability	
Chapter 24	Overview	2269
	Software Features Overview	2269
	Graceful Restart Concepts	2269
	Understanding Nonstop Software Upgrade for QFabric Systems	2270
	Understanding VRRP	2274
	Overview of VRRP	2274
	Sample VRRP Topology	2275
	Understanding VRRP Between QFabric Systems	2277
	VRRP Differences on QFabric Systems	2277
	Configuration Details	2277
Chapter 25	Configuration	2281
	Configuration Tasks for Graceful Restart	2281
	Configuring Graceful Restart for QFabric Systems	2281
	Enabling Graceful Restart	2282
	Configuring Graceful Restart Options for BGP	2282
	Configuring Graceful Restart Options for OSPF and OSPFv3	2283
	Tracking Graceful Restart Events	2284
	Configuring Routing Protocols Graceful Restart	2285
	Enabling Graceful Restart	2285
	Configuring Graceful Restart Options for BGP	2286
	Configuring Graceful Restart Options for ES-IS	2287
	Configuring Graceful Restart Options for IS-IS	2287
	Configuring Graceful Restart Options for OSPF and OSPFv3	2288
	Configuring Graceful Restart Options for RIP and RIPng	2289
	Configuring Graceful Restart Options for PIM Sparse Mode	2289
	Tracking Graceful Restart Events	2291
	Configuration Example for VRRP	2291
	Example: Configuring VRRP for Load Sharing	2291
	Configuration Tasks for VRRP	2296
	Configuring Basic VRRP Support	2296
	Configuring VRRP Authentication (IPv4 Only)	2297
	Configuring the Startup Period for VRRP Operations	2298
	Configuring the Advertisement Interval for the VRRP Master	2299
	Modifying the Advertisement Interval in Seconds	2299
	Modifying the Advertisement Interval in Milliseconds	2299
	Configuring VRRP Preemption and Hold Time	2300
	Configuring VRRP Preemption	2300
	Configuring the Preemption Hold Time	2300
	Overriding the Hold Time	2301
	Configuring a Route to Be Tracked	2301
	Configuring a Logical Interface to Be Tracked	2302
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	2304
	Address	2304
	Configuring Passive ARP Learning for VRRP Backups	2304
	Configuring the Silent Period	2305
	Configuring Inheritance for a VRRP Group	2305

Configuration Statements for Graceful Restart	2306
disable	2307
disable (BGP Graceful Restart)	2308
graceful-restart (Enabling Globally)	2309
graceful-restart (Protocols BGP)	2310
graceful-restart (Protocols OSPF)	2311
helper-disable (OSPF)	2313
no-strict-lsa-checking	2314
notify-duration	2315
restart-duration	2316
restart-time (BGP Graceful Restart)	2317
stale-routes-time	2318
Configuration Statements for VRRP	2318
accept-data	2320
advertise-interval	2321
asymmetric-hold-time	2322
authentication-key	2323
authentication-type	2324
bandwidth-threshold	2325
failover-delay	2326
fast-interval	2327
hold-time (VRRP)	2328
interface (VRRP Group)	2329
preempt (VRRP)	2330
priority (Protocols VRRP)	2331
priority-cost (VRRP)	2332
priority-hold-time	2333
route (Interfaces)	2334
startup-silent-period	2335
traceoptions	2336
track (VRRP)	2338
virtual-address	2339
vrrp-group	2340
Chapter 26 Administration	2343
Operational Mode Commands for Graceful Restart	2343
Verifying Graceful Restart Operation	2343
Graceful Restart Operational Mode Commands	2343
Verifying BGP Graceful Restart	2344
Verifying IS-IS and OSPF Graceful Restart	2344
Verifying CCC and TCC Graceful Restart	2345
show bgp neighbor	2346
show log	2360
show (ospf ospf3) overview	2363
Operational Mode Commands for Nonstop Software Upgrade	2367
Performing a Nonstop Software Upgrade on the QFabric System	2367
Backing Up the Current Configuration Files	2368
Downloading Software Files Using a Browser	2369
Retrieving Software Files for Download	2370

	Performing a Nonstop Software Upgrade for Director Devices in a Director Group	2370
	Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components	2370
	(Optional) Creating Upgrade Groups for Node Groups	2371
	Performing a Nonstop Software Upgrade on a Node Group	2372
	Verifying Nonstop Software Upgrade for QFabric Systems	2372
	Verifying a Director Group Nonstop Software Upgrade	2373
	Verifying a Fabric Nonstop Software Upgrade	2385
	Verifying a Redundant Server Node Group Nonstop Software Upgrade	2387
	Verifying a Network Node Group Nonstop Software Upgrade	2390
	request system software nonstop-upgrade	2393
	show chassis nonstop-upgrade node-group	2399
	Operational Mode Commands for VRRP	2399
	show vrrp	2400
Chapter 27	Troubleshooting	2411
	Troubleshooting Procedures	2411
	Troubleshooting VRRP	2411
Part 10	Interfaces	
Chapter 28	Overview	2415
	Interfaces Overview	2415
	Interfaces Overview	2415
	Network Interfaces	2415
	Special Interfaces	2416
	Understanding Aggregated Ethernet Interfaces and LACP	2417
	Link Aggregation Group	2417
	Link Aggregation Control Protocol (LACP)	2418
	Understanding Interface Naming Conventions	2419
	Physical Part of an Interface Name	2419
	Logical Part of an Interface Name	2431
	Wildcard Characters in Interface Names	2432
	Understanding Interface Ranges on the QFX Series	2432
	Understanding Layer 3 Logical Interfaces	2434
	Understanding Management Interfaces	2434
	Understanding Multichassis Link Aggregation	2435
	Active-Active Mode	2437
	ICCP and ICL-PL	2437
	Failure Handling	2437
	Multichassis Link Protection	2438
	MC-LAG Packet Forwarding	2438
	Layer 3 Routing	2438
	Spanning Tree Protocol (STP) Guidelines	2439
	MC-LAG Upgrade Guidelines	2439
	Layer 2 Unicast Features Supported	2440
	Layer 2 Multicast Features Supported	2440
	IGMP Snooping on an Active-Active MC-LAG	2440

	Layer 3 Unicast Features Supported	2441
	VRRP Active-Standby Support	2441
	Routed VLAN Interface (RVI) MAC Address Synchronization	2442
	Address Resolution Protocol (ARP)	2442
	DHCP Relay with Option 82	2442
	Private VLAN (PVLAN)	2443
	Layer 3 Multicast	2443
	Overview of Uplink Failure Detection	2445
	Uplink Failure Detection Configuration	2445
	Failure Detection Pair	2446
Chapter 29	Configuration	2449
	Configuration Examples	2449
	Example: Configuring Interfaces for Uplink Failure Detection	2449
	Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch	2454
	Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch	2458
	Example: Configuring Multichassis Link Aggregation	2462
	Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP	2477
	Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization	2500
	Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP)	2514
	Configuration Tasks	2532
	Configuring Gigabit and 10-Gigabit Ethernet Interfaces	2533
	Configuring Port Mode	2533
	Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces	2533
	Configuring the IP Options	2534
	Configuring Aggregated Ethernet LACP	2534
	Configuring Ethernet Loopback Capability	2535
	Configuring Interfaces for Uplink Failure Detection	2536
	Configuring a Layer 3 Logical Interface	2537
	Configuring Link Aggregation	2537
	Creating an Aggregated Ethernet Interface	2538
	Configuring the VLAN Name and VLAN ID Number	2538
	Configuring Aggregated Ethernet LACP	2538
	Configuring Multichassis Link Aggregation	2540
	Configuring the QSFP+ Port Type on QFX3500 Standalone Switches	2543
	Configuring the Port Type on QFX3600 Standalone Switches	2545
	Configuration Statements	2546
	802.3ad	2549
	address	2550
	aggregated-devices	2552
	aggregated-ether-options	2553
	alarm (chassis)	2554
	authentication-key (ICCP)	2555

backup-liveness-detection	2555
backup-peer-ip	2556
chassis	2557
chassis (QFabric System)	2558
chassis-id	2559
container-devices	2560
craft-lockout	2561
description (Interfaces)	2562
detection-time (Liveness Detection)	2563
device-count	2563
disk-failure-action	2564
ethernet	2564
ethernet (Alarm)	2565
ethernet-switching	2566
ether-options	2567
eui-64	2568
family	2569
fibre-channel (Alarm)	2571
flow-control	2572
force-up	2573
fpc	2574
fpc (Interconnect Device)	2575
gratuitous-arp-reply	2575
group	2576
hold-time (Physical Interface)	2577
iccp	2578
inet (interfaces)	2579
inet6 (interfaces)	2580
interconnect-device (Chassis)	2581
interface-range	2582
interface (Multichassis Protection)	2583
interfaces	2584
lacp (802.3ad)	2590
lacp (Aggregated Ethernet)	2591
link-to-disable	2591
link-to-monitor	2592
link-down	2593
link-mode	2594
link-speed	2595
liveness-detection	2596
local-ip-addr (ICCP)	2596
loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)	2597
management-ethernet (Alarm)	2597
mc-ae	2598
mc-ae-id	2599
member	2599
member-range	2600
minimum-interval (Liveness Detection)	2600

minimum-receive-interval (Liveness Detection)	2601
minimum-links	2601
mode (QFX Series)	2602
mtu	2603
multi-chassis	2604
multi-chassis-protection	2604
multiplier (Liveness Detection)	2605
no-adaptation (Liveness Detection)	2605
no-gratuitous-arp-request	2606
node-device (Chassis)	2607
node-group (Chassis)	2608
on-disk-failure	2609
peer (ICCP)	2610
peer (Multichassis)	2611
periodic	2611
pic	2612
port-mode	2613
reflective-relay	2614
routing-engine	2614
session-establishment-hold-time	2615
speed	2616
status-control	2616
targeted-broadcast	2617
threshold (Detection Time)	2617
traceoptions (Individual Interfaces)	2618
traceoptions (ICCP)	2619
transmit-interval (Liveness Detection)	2620
traps	2621
unit	2622
uplink-failure-detection	2623
version (Liveness Detection)	2623
vlan-id	2624
vlan-tagging	2624
xe (Port)	2625
xle (Port)	2626
Chapter 30 Administration	2627
Routine Monitoring	2627
Monitoring System Process Information	2627
Monitoring System Properties	2628
Monitoring Interface Status and Traffic	2629
Verifying That Layer 3 Logical Interfaces Are Working	2630
Verifying the Status of a LAG Interface	2630
Verifying That LACP Is Configured Correctly and Bundle Members Are	
Exchanging LACP Protocol Packets	2631
Verifying the LACP Setup	2631
Verifying That LACP Packets Are Being Exchanged	2632

Verifying That Uplink Failure Detection Is Working Correctly	2632
Monitoring Commands	2633
monitor interface	2635
show iccp	2643
show interfaces diagnostics optics	2645
show interfaces fabric	2651
show interfaces ge	2672
show interfaces mc-ae	2687
show interfaces statistics fabric	2689
show interfaces queue	2707
show interfaces queue fabric	2743
show interfaces xe	2765
show interfaces xle	2783
show lacp interfaces	2801
show lacp statistics interfaces (View)	2806
show uplink-failure-detection	2808
Chapter 31 Troubleshooting	2811
Troubleshooting Procedures	2811
Troubleshooting an Aggregated Ethernet Interface	2811
Troubleshooting Multichassis Link Aggregation	2811
MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table	2812
MC-LAG Peer Does Not Go into Standby Mode	2813
Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive	2813
Redirect Filters Take Priority over User-Defined Filters	2813
Operational Command Output Is Wrong	2813
ICCP Connection Might Take Up to 60 Seconds to Become Active	2814
MAC Address Age Learned on an MC-AE Interface Is Reset to Zero	2814
MAC Address Is Not Learned Remotely in a Default VLAN	2814
Snooping Entries Learned on MC-AE Interfaces Are Not Removed	2814
ICCP Does Not Come Up After You Add or Delete an Authentication Key	2815
Local Status Is Standby When It Should Be Active	2815
Packets Loop on the Server When ICCP Fails	2815
Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change	2815
No Commit Checks Are Done for ICL-PL Interfaces	2816
Double Failover Scenario	2816
Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up	2816
Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer	2816
AE Interfaces Go Down	2816
Flooding of Upstream Traffic	2817
Troubleshooting Network Interfaces	2817
The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down	2817

Part 11	Routing Options	
Chapter 32	Overview	2821
	Routing Options Overview	2821
	Overview of Routing Options	2821
	Understanding Virtual Router Routing Instances	2822
	Understanding Distributed Periodic Packet Management	2822
Chapter 33	Configuration	2825
	Configuration Tasks	2825
	Configuring Static Routing	2826
	Configuring Per-Packet Load Balancing	2826
	Configuring Distributed Periodic Packet Management	2828
	Disabling or Enabling Distributed Periodic Packet Management	
	Globally	2828
	Disabling or Enabling Distributed Periodic Packet Management for LACP	
	Packets	2828
	Configuring Virtual Router Routing Instances	2829
	Configuration Examples	2830
	Examples: Configuring Per-Packet Load Balancing	2830
	Examples: Configuring BFD for Static Routes	2831
	Understanding BFD for Static Routes	2831
	Example: Configuring BFD for Static Routes	2835
	Example: Enabling BFD on Qualified Next Hops in Static Routes	2840
	Example: Configuring BFD Authentication for Static Routes	2846
	Understanding BFD Authentication for Static Routes	2846
	Example: Configuring BFD Authentication for Static Routes	2848
	Configuration Statements	2854
	active	2857
	aggregate (Routing)	2858
	as-path (Routing Options)	2860
	autonomous-system	2862
	backup-pe-group	2864
	backups	2865
	bandwidth (Multicast Flow Map)	2866
	bfd-liveness-detection (Routing Options Static Route)	2867
	bgp-orf-cisco-mode	2871
	bmp	2872
	brief	2873
	centralized	2874
	community (Routing Options)	2875
	confederation	2877
	description (Routing Instances)	2878
	discard	2879
	export (Routing Options)	2880
	export-rib	2881
	fate-sharing	2882
	flow	2883
	flow-map	2884
	forwarding-cache (Flow Maps)	2885

forwarding-cache (Multicast)	2886
forwarding-table	2887
generate	2888
import (Routing Options)	2889
import-policy	2890
import-rib	2891
indirect-next-hop	2892
install (Routing Options)	2893
instance-type	2894
interface (Multicast Static Routes)	2895
interface (Routing Instances)	2896
interface (Routing Options)	2897
interface-routes	2898
local-address (Routing Options)	2899
martians	2900
maximum-bandwidth (Routing Options)	2901
maximum-paths	2902
maximum-prefixes	2904
med-igp-update-interval	2905
metric (Aggregate, Generated, or Static Route)	2906
multicast (Routing Options)	2907
no-qos-adjust	2908
options (Routing Options)	2909
pim-to-igmp-proxy	2910
policy (Aggregate and Generated Routes)	2911
policy (Flow Maps)	2912
policy-options	2913
policy-statement	2914
ppm	2916
ppm (Ethernet Switching)	2917
preference (Routing Options)	2918
prefix	2919
protocols	2920
qualified-next-hop (Static Routes)	2922
readvertise	2924
redundant-sources	2925
resolution	2926
resolution-ribs	2927
resolve	2928
retain	2929
reverse-oif-mapping	2930
rpf-check-policy (Routing Options RPF)	2931
rib (General)	2932
rib (Route Resolution)	2934
rib-group (Routing Options)	2935
rib-groups	2936
route-record	2937
router-id	2938
routing-instances	2939

	routing-options	2939
	scope	2940
	scope-policy	2941
	source-routing	2942
	static (Routes)	2943
	subscriber-leave-timer	2945
	tag (Routing Options)	2946
	threshold (Multicast Forwarding Cache)	2947
	timeout (Flow Maps)	2948
	timeout (Multicast)	2949
	traceoptions (Routing Options)	2950
	upstream-interface	2953
Chapter 34	Administration	2955
	Routine Monitoring	2955
	Monitoring Routing Information	2955
	Verifying That Virtual Router Routing Instances Are Working	2956
	Operational Commands	2957
	clear ipv6 neighbors	2959
	show as-path	2960
	show as-path domain	2964
	show as-path summary	2967
	show ipv6 neighbors	2969
	show ipv6 router-advertisement	2971
	show route	2974
	show route active-path	2979
	show route all	2984
	show route aspath-regex	2986
	show route best	2988
	show route brief	2991
	show route community	2993
	show route community-name	2995
	show route damping	2997
	show route detail	3002
	show route exact	3017
	show route export	3019
	show route extensive	3022
	show route flow validation	3037
	show route forwarding-table	3039
	show route inactive-path	3052
	show route inactive-prefix	3055
	show route instance	3057
	show route label	3064
	show route label-switched-path	3066
	show route martians	3068
	show route next-hop	3070
	show route no-community	3076
	show route protocol	3079
	show route range	3091

	show route receive-protocol	3095
	show route resolution	3103
	show route snooping	3106
	show route source-gateway	3114
	show route summary	3120
	show route table	3124
	show route terse	3135
Chapter 35	Troubleshooting	3139
	Troubleshooting Procedures	3139
	Troubleshooting Virtual Routing Instances	3139
	Direct Routes Not Leaked Between Routing Instances	3139
Part 12	Border Gateway Protocol	
Chapter 36	Overview	3143
	BGP Overview	3143
	Understanding BGP	3144
	Autonomous Systems	3144
	AS Paths and Attributes	3144
	External and Internal BGP	3145
	Multiple Instances of BGP	3145
	BGP Routes Overview	3146
	BGP Messages Overview	3147
	Open Messages	3147
	Update Messages	3148
	Keepalive Messages	3148
	Notification Messages	3148
Chapter 37	Configuration	3149
	Basic BGP Configuration	3149
	Examples: Configuring External BGP Peering	3149
	Understanding External BGP Peering Sessions	3149
	Example: Configuring External BGP Point-to-Point Peer Sessions	3150
	Example: Configuring External BGP on Logical Systems with IPv6 Interfaces	3157
	Examples: Configuring Internal BGP Peering	3172
	Understanding Internal BGP Peering Sessions	3172
	Example: Configuring Internal BGP Peer Sessions	3173
	Example: Configuring Internal BGP Peering Sessions on Logical Systems	3184
	BGP Path Attribute Configuration	3194
	Example: Configuring BGP Local Preference	3194
	Understanding the BGP Local Preference	3194
	Example: Configuring the Local Preference Value for BGP Routes	3195
	Examples: Configuring BGP MED	3208
	Understanding the MED Attribute	3208
	Example: Configuring the MED Attribute Directly	3210
	Example: Configuring the MED Using Route Filters	3223
	Example: Configuring the MED Using Communities	3236

Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates	3236
Examples: Configuring BGP Local AS	3246
Understanding the BGP Local AS Attribute	3246
Example: Configuring a Local AS for EBGp Sessions	3251
Example: Configuring a Private Local AS for EBGp Sessions	3261
Example: Configuring the Accumulated IGP Attribute for BGP	3266
Understanding the Accumulated IGP Attribute for BGP	3267
Example: Configuring the Accumulated IGP Attribute for BGP	3267
BGP Policy Configuration	3305
Example: Configuring BGP Interactions with IGP's	3305
Understanding Routing Policies	3305
Example: Injecting OSPF Routes into the BGP Routing Table	3306
Example: Configuring BGP Route Advertisement	3309
Understanding Route Advertisement	3309
Example: Configuring BGP Prefix-Based Outbound Route Filtering	3313
Example: Configuring EBGp Multihop	3317
Understanding BGP Multihop	3317
Example: Configuring EBGp Multihop Sessions	3317
Example: Configuring BGP Route Preference (Administrative Distance)	3326
Understanding Route Preference Values	3326
Example: Configuring the Preference Value for BGP Routes	3327
Example: Configuring BGP Path Selection	3332
Understanding BGP Path Selection	3332
Example: Ignoring the AS Path Attribute When Selecting the Best Path	3336
Example: Removing Private AS Numbers	3342
Understanding Private AS Number Removal from AS Paths	3342
Example: Removing Private AS Numbers from AS Paths	3343
BGP BFD Configuration	3348
Example: Configuring BFD for BGP	3348
Understanding BFD for BGP	3348
Example: Configuring BFD on Internal BGP Peer Sessions	3349
Example: Configuring BFD Authentication for BGP	3357
Understanding BFD Authentication for BGP	3357
Example: Configuring BFD Authentication for BGP	3359
BGP Load Balancing Configuration	3362
Examples: Configuring BGP Multipath	3362
Understanding BGP Multipath	3362
Example: Load Balancing BGP Traffic	3363
Example: Configuring Single-Hop EBGp Peers to Accept Remote Next Hops	3367
Example: Advertising Multiple BGP Paths to a Destination	3379
Understanding the Advertisement of Multiple Paths to a Single Destination in BGP	3379
Example: Advertising Multiple Paths in BGP	3380

IBGP Scaling Configuration	3405
Example: Configuring BGP Route Reflectors	3405
Understanding BGP Route Reflectors	3405
Example: Configuring a Route Reflector	3407
Example: Configuring BGP Confederations	3422
Understanding BGP Confederations	3422
Example: Configuring BGP Confederations	3423
BGP Security Configuration	3428
Example: Configuring BGP Route Authentication	3429
Understanding Route Authentication	3429
Example: Configuring Route Authentication for BGP	3430
Examples: Configuring TCP and BGP Security	3435
Understanding Security Options for BGP with TCP	3435
Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers	3436
Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List	3441
Example: Limiting TCP Segment Size for BGP	3444
BGP Flap Configuration	3449
Example: Preventing BGP Session Resets	3449
Understanding BGP Session Resets	3449
Example: Preventing BGP Session Flaps When VPN Families Are Configured	3449
Examples: Configuring BGP Flap Damping	3456
Understanding Damping Parameters	3456
Example: Configuring Damping Parameters	3457
Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family	3466
BGP Monitoring Configuration	3475
Example: Configuring BGP Trace Operations	3476
Understanding Trace Operations for BGP Protocol Traffic	3476
Example: Viewing BGP Trace Files on Logical Systems	3477
Configuration Statements	3482
accept-remote-nexthop	3485
advertise-external	3486
advertise-inactive	3487
advertise-peer-as	3488
algorithm (BGP BFD Authentication)	3489
apply-groups	3491
apply-groups-except	3491
authentication (BGP BFD Liveness Detection)	3492
authentication-algorithm	3494
authentication-key (Protocols BGP)	3495
authentication-key-chain (Protocols BGP)	3496
bfd-liveness-detection (Protocols BGP)	3497
bgp	3501
bgp-orf-cisco-mode	3502
cluster	3504
damping (Protocols BGP)	3506

description (Protocols BGP)	3508
detection-time (BFD for BGP)	3509
disable (Protocols BGP)	3510
disable (BGP Graceful Restart)	3511
export (Protocols BGP)	3512
family (Protocols BGP)	3513
graceful-restart (Protocols BGP)	3517
group (Protocols BGP)	3518
hold-down-interval (BGP BFD Liveness Detection)	3521
hold-time (Protocols BGP)	3523
import (Protocols BGP)	3525
include-mp-next-hop	3526
keep	3527
key-chain (BGP BFD Authentication)	3529
local-address (Protocols BGP)	3531
local-as	3533
local-preference	3536
log-updown (Protocols BGP)	3537
loops	3538
loose-check (BGP BFD Authentication)	3540
metric-out (Protocols BGP)	3542
minimum-interval (BGP BFD Liveness Detection)	3544
minimum-interval (BGP BFD Transmit Interval)	3545
minimum-receive-interval (BGP BFD Liveness Detection)	3547
mtu-discovery	3548
multihop	3550
multiplier (BGP BFD Liveness Detection)	3552
neighbor (Protocols BGP)	3553
no-adaptation (BGP BFD Liveness Detection)	3556
no-advertise-peer-as	3557
no-aggregator-id	3558
no-client-reflect	3559
out-delay	3560
outbound-route-filter	3561
passive (Protocols BGP)	3562
path-selection	3563
peer-as (Protocols BGP)	3565
preference (Protocols BGP)	3567
remove-private	3568
restart-time (BGP Graceful Restart)	3570
session-mode	3571
stale-routes-time	3572
tcp-mss (Protocols BGP)	3573
threshold (BGP BFD Detection Time)	3574
threshold (BGP BFD Transmit Interval)	3576
traceoptions (Protocols BGP)	3577
transmit-interval (BGP BFD Liveness Detection)	3580
version (BGP BFD Liveness Detection)	3581

Chapter 38	Administration	3583
	Routine Monitoring	3583
	Monitoring BGP Routing Information	3583
	Operational Commands	3583
	clear bgp damping	3584
	clear bgp neighbor	3585
	clear bgp table	3587
	show bgp group	3589
	show bgp neighbor	3596
	show bgp summary	3610
	show policy damping	3615
	show route damping	3617
Part 13	Intermediate System to Intermediate System	
Chapter 39	Overview	3625
	IS-IS Overview	3625
	IS-IS Overview	3625
	IS-IS Terminology	3626
	ISO Network Addresses	3626
	IS-IS Packets	3628
	Persistent Route Reachability	3629
	IS-IS Support for Multipoint Network Clouds	3629
	Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels	3629
	Understanding BFD Authentication for IS-IS	3630
	BFD Authentication Algorithms	3630
	Security Authentication Keychains	3631
	Strict Versus Loose Authentication	3631
	Understanding Hitless Authentication Key Rollover for IS-IS	3631
Chapter 40	Configuration	3633
	Configuration Guidelines	3633
	Example: Configuring IS-IS	3633
	Configuration Examples	3638
	Example: Configuring Multi-Level IS-IS	3639
	Example: Configuring Hitless Authentication Key Rollover for IS-IS	3647
	Example: Redistributing OSPF Routes into IS-IS	3651
	Example: Configuring BFD for IS-IS	3659
	Example: Configuring BFD Authentication for IS-IS	3665
	Example: Configuring IS-IS Multicast Topology	3668
	IS-IS Multicast Topologies Overview	3669
	Example: Configuring IS-IS Multicast Topology	3670
	Example: Configuring IS-IS for CLNS	3684
	Understanding IS-IS for CLNS	3684
	Example: Configuring IS-IS for CLNS	3684
	Example: Configuring IS-IS Designated Routers	3686
	Understanding IS-IS Designated Routers	3687
	Example: Configuring Designated Router Election Priority for IS-IS	3687

Example: Enabling Packet Checksums on IS-IS Interfaces	3687
Configuration Tasks	3690
Configuring IS-IS Authentication	3690
Configuring Authentication Without Network-Wide Deployment	3691
Configuration Statements	3692
authentication-key (Protocols IS-IS)	3694
authentication-key-chain (Protocols IS-IS)	3695
authentication-type (Protocols IS-IS)	3696
bfd-liveness-detection (Protocols IS-IS)	3697
checksum (Protocols IS-IS)	3699
csnp-interval	3700
disable (Protocols IS-IS)	3701
export (Protocols IS-IS)	3702
external-preference (Protocols IS-IS)	3703
family (Protocols IS-IS)	3704
hello-authentication-key	3705
hello-authentication-key-chain	3706
hello-authentication-type	3707
hello-interval (Protocols IS-IS)	3708
hello-padding	3709
hold-time (Protocols IS-IS)	3711
ignore-attached-bit	3712
interface (Protocols IS-IS)	3713
ipv4-multicast	3715
ipv4-multicast-metric	3716
ipv6-multicast	3716
ipv6-multicast-metric	3717
ipv6-unicast	3718
ipv6-unicast-metric	3719
isis	3720
level (Global IS-IS)	3721
loose-authentication-check	3722
lsp-interval	3723
lsp-lifetime	3724
max-areas	3725
mesh-group (Protocols IS-IS)	3726
metric (Protocols IS-IS)	3727
no-adjacency-holddown	3728
no-authentication-check	3729
no-csnp-authentication	3729
no-hello-authentication	3730
no-ipv4-multicast	3730
no-ipv4-routing	3731
no-ipv6-multicast	3732
no-ipv6-routing	3733
no-ipv6-unicast	3734
no-psnp-authentication	3734
no-unicast-topology	3735
overload (Protocols IS-IS)	3736

	passive (Protocols IS-IS)	3739
	point-to-point	3740
	preference (Protocols IS-IS)	3741
	prefix-export-limit (Protocols IS-IS)	3742
	priority (Protocols IS-IS)	3743
	reference-bandwidth (Protocols IS-IS)	3744
	rib-group (Protocols IS-IS)	3745
	topologies (Protocols IS-IS)	3746
	traceoptions (Protocols IS-IS)	3747
	traffic-engineering (Protocols IS-IS)	3750
	wide-metrics-only	3753
Chapter 41	Administration	3755
	Operational Commands	3755
	clear isis adjacency	3756
	clear isis database	3758
	clear isis overload	3760
	clear isis statistics	3762
	show isis adjacency	3764
	show isis authentication	3768
	show isis database	3770
	show isis hostname	3777
	show isis interface	3779
	show isis overview	3784
	show isis route	3787
	show isis statistics	3791
Part 14	Open Shortest Path First	
Chapter 42	Overview	3797
	OSPF Overview	3797
	OSPF Overview	3798
	OSPF Default Route Preference Values	3800
	OSPF Routing Algorithm	3800
	OSPF Three-Way Handshake	3801
	OSPF Version 3	3802
	OSPF Areas and Router Functionality Overview	3803
	Areas	3803
	Area Border Routers	3803
	Backbone Areas	3803
	AS Boundary Routers	3804
	Backbone Router	3804
	Internal Router	3804
	Stub Areas	3804
	Not-So-Stubby Areas	3804
	Transit Areas	3805
	Packets Overview	3805
	OSPF Packet Header	3805
	Hello Packets	3806
	Database Description Packets	3806

	Link-State Request Packets	3806
	Link-State Update Packets	3806
	Link-State Acknowledgment Packets	3807
	Link-State Advertisement Packet Types	3807
	OSPF External Metrics Overview	3808
Chapter 43	Configuration	3809
	Basic OSPF Area Configuration	3809
	Examples: Configuring OSPF Designated Routers	3809
	OSPF Designated Router Overview	3809
	Example: Configuring an OSPF Router Identifier	3810
	Example: Controlling OSPF Designated Router Election	3812
	Examples: Configuring OSPF Areas	3814
	Understanding OSPF Areas and Backbone Areas	3814
	Example: Configuring a Single-Area OSPF Network	3815
	Example: Configuring a Multiarea OSPF Network	3817
	Advanced OSPF Area Configuration	3820
	Examples: Configuring OSPF Stub and Not-So-Stubby Areas	3821
	Understanding OSPF Stub Areas, Totally Stubby Areas, and	
	Not-So-Stubby Areas	3821
	Example: Configuring OSPF Stub and Totally Stubby Areas	3822
	Example: Configuring OSPF Not-So-Stubby Areas	3826
	Example: Configuring OSPF Multiarea Adjacency	3831
	Multiarea Adjacency for OSPF	3831
	Example: Configuring Multiarea Adjacency for OSPF	3832
	Example: Disabling OSPFv2 Compatibility with RFC 1583	3835
	OSPFv2 Compatibility with RFC 1583 Overview	3836
	Example: Disabling OSPFv2 Compatibility with RFC 1583	3836
	OSPF Interface Configuration	3837
	Examples: Configuring OSPF Interfaces	3837
	About OSPF Interfaces	3838
	Example: Configuring an Interface on a Broadcast or Point-to-Point	
	Network	3839
	Example: Configuring an OSPFv2 Interface on a Nonbroadcast	
	Multiaccess Network	3841
	Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint	
	Network	3844
	Example: Configuring OSPF Demand Circuits	3846
	Example: Configuring a Passive OSPF Interface	3848
	Example: Configuring OSPFv2 Peer interfaces	3850

OSPF Route Control Configuration	3852
Examples: Configuring OSPF Route Summarization	3852
Understanding OSPF Route Summarization	3852
Example: Summarizing Ranges of Routes in OSPF Link-State	
Advertisements	3852
Example: Limiting the Number of Prefixes Exported to OSPF	3858
Configuring OSPF Refresh and Flooding Reduction in Stable	
Topologies	3860
Examples: Configuring OSPF Traffic Control	3861
Understanding OSPF Traffic Control	3861
Example: Controlling the Cost of Individual OSPF Network	
Segments	3863
Example: Dynamically Adjusting OSPF Interface Metrics Based on	
Bandwidth	3867
Example: Controlling OSPF Route Preferences	3869
Example: Configuring OSPF Overload Mode	3871
OSPF Overload Function Overview	3871
Example: Configuring OSPF to Make Routing Devices Appear	
Overloaded	3872
OSPF Fault Detection Configuration	3875
Example: Configuring OSPF Timers	3875
OSPF Timers Overview	3875
Example: Configuring OSPF Timers	3876
Example: Configuring BFD for OSPF	3881
BFD for OSPF Overview	3881
Example: Configuring BFD for OSPF	3883
Example: Configuring BFD Authentication for OSPF	3887
BFD Authentication for OSPF Overview	3887
Configuring BFD Authentication for OSPF	3889
OSPF Redundancy Features Configuration	3892
Examples: Configuring Graceful Restart for OSPF	3892
Graceful Restart for OSPF Overview	3892
Example: Configuring Graceful Restart for OSPF	3894
Example: Configuring the Helper Capability Mode for OSPFv2 Graceful	
Restart	3898
Example: Configuring the Helper Capability Mode for OSPFv3 Graceful	
Restart	3901
Example: Disabling Strict LSA Checking for OSPF Graceful Restart . .	3904
OSPF Traffic Engineering Configuration	3907
Examples: Configuring OSPF Traffic Engineering	3907
OSPF Support for Traffic Engineering	3908
Example: Enabling OSPF Traffic Engineering Support	3910
Example: Configuring the Traffic Engineering Metric for a Specific OSPF	
Interface	3914
Example: Configuring OSPF Passive Traffic Engineering Mode	3916
OSPF Passive Traffic Engineering Mode	3916
Example: Configuring OSPF Passive Traffic Engineering Mode	3917

OSPF Database Protection Configuration	3919
Example: Configuring OSPF Database Protection	3919
OSPF Database Protection Overview	3919
Configuring OSPF Database Protection	3920
OSPF Policy Configuration	3921
Examples: Configuring OSPF Routing Policy	3921
Understanding OSPF Routing Policy	3921
Example: Injecting OSPF Routes into the BGP Routing Table	3923
Example: Redistributing Static Routes into OSPF	3927
Example: Configuring an OSPF Import Policy	3929
Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF	3933
Examples: Configuring Routing Policy for Network Summaries	3937
Import and Export Policies for Network Summaries Overview	3937
Example: Configuring an OSPF Export Policy for Network Summaries	3938
Example: Configuring an OSPF Import Policy for Network Summaries	3946
OSPF Monitoring Configuration	3954
Example: Configuring OSPF Trace Options	3954
Tracing OSPF Protocol Traffic	3954
Example: Tracing OSPF Protocol Traffic	3956
Configuration Statements	3961
area	3963
area-range	3965
authentication (Protocols OSPF)	3967
backup-spf-options	3968
bandwidth-based-metrics	3969
bfd-liveness-detection (Protocols OSPF)	3971
context-identifier (Protocols OSPF)	3974
database-protection	3975
dead-interval	3977
default-lsa	3978
disable (OSPF)	3979
export (Protocols OSPF)	3981
external-preference (Protocols OSPF)	3982
flood-reduction	3983
graceful-restart (Protocols OSPF)	3984
hello-interval (Protocols OSPF)	3986
helper-disable (Multiple Protocols)	3987
ignore-lsp-metrics	3988
import (Protocols OSPF)	3989
inter-area-prefix-export	3990
inter-area-prefix-import	3991
interface (Protocols OSPF)	3992
interface-type (Protocols OSPF)	3995
lsa-refresh-interval	3996
metric (Protocols OSPF Interface)	3997
no-eligible-backup (Protocols OSPF)	3999

no-nssa-abr	4000
no-rfc-1583	4001
no-strict-lsa-checking	4002
node-link-protection (Protocols OSPF)	4003
notify-duration	4004
nssa	4005
ospf	4006
ospf3	4006
overload (Protocols OSPF)	4007
passive (Protocols OSPF)	4009
preference (Protocols OSPF)	4010
prefix-export-limit (Protocols OSPF)	4011
priority (Protocols OSPF)	4012
realm	4013
reference-bandwidth (Protocols OSPF)	4014
retransmit-interval (OSPF)	4015
rib-group (Protocols OSPF)	4016
shortcuts (Protocols OSPF)	4017
spf-options (Protocols OSPF)	4018
stub	4020
summaries	4021
topology (OSPF)	4022
traceoptions (Protocols OSPF)	4023
traffic-engineering (OSPF)	4026
transmit-interval (Protocols OSPF)	4028
transit-delay (OSPF)	4029
Chapter 44 Administration	4031
Routine Monitoring	4031
Monitoring OSPF Routing Information	4031
Operational Commands	4031
clear (ospf ospf3) database	4033
clear (ospf ospf3) database-protection	4036
clear (ospf ospf3) io-statistics	4037
clear (ospf ospf3) neighbor	4038
clear (ospf ospf3) overload	4040
clear (ospf ospf3) statistics	4041
show (ospf ospf3) backup coverage	4043
show (ospf ospf3) backup neighbor	4046
show ospf context-identifier	4048
show ospf database	4050
show ospf3 database	4058
show (ospf ospf3) interface	4069
show (ospf ospf3) io-statistics	4075
show (ospf ospf3) log	4077
show (ospf ospf3) neighbor	4080
show (ospf ospf3) overview	4086
show (ospf ospf3) route	4091
show (ospf ospf3) statistics	4097

Part 15	Routing Information Protocol	
Chapter 45	Overview	4103
	RIP Overview	4103
	RIP Overview	4103
	Distance-Vector Routing Protocols	4103
	RIP Protocol Overview	4104
	RIP Packets	4105
	Maximizing Hop Count	4106
	Split Horizon and Poison Reverse Efficiency Techniques	4106
	Limitations of Unidirectional Connectivity	4107
Chapter 46	Configuration	4109
	RIP Configuration Tasks	4109
	Example: Configuring RIP	4109
	Understanding Basic RIP Routing	4109
	Example: Configuring a Basic RIP Network	4110
	Example: Configuring Authentication for RIP Routes	4116
	Understanding RIP Authentication	4116
	Example: Configuring Route Authentication for RIP	4116
	Enabling Authentication with Plain-Text Passwords (CLI Procedure)	4121
	Enabling Authentication with MD5 Authentication (CLI Procedure)	4121
	Example: Configuring BFD for RIP	4122
	Understanding BFD for RIP	4122
	Example: Configuring BFD for RIP	4123
	Example: Configuring BFD Authentication for RIP	4128
	Understanding BFD Authentication for RIP	4128
	Example: Configuring BFD Authentication for RIP	4130
	Example: Applying Policies to RIP Routes Imported from Neighbors	4136
	Understanding RIP Import Policy	4136
	Example: Applying Policies to RIP Routes Imported from Neighbors	4136
	Examples: Controlling Traffic with Metrics in a RIP Network	4142
	Understanding Traffic Control with Metrics in a RIP Network	4142
	Example: Controlling Traffic in a RIP Network with an Incoming Metric	4143
	Example: Controlling Traffic in a RIP Network with an Outgoing Metric	4144
	Example: Configuring the Metric Value Added to Imported RIP Routes	4146
	Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets	4150
	Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets	4150
	Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets	4150
	Example: Redistributing Routes Among RIP Instances	4154
	Understanding Route Redistribution Among RIP instances	4154
	Example: Redistributing Routes Between Two RIP Instances	4155

Example: Configuring RIP Timers	4159
Understanding RIP Timers	4160
Example: Configuring RIP Timers	4161
Example: Tracing RIP Protocol Traffic	4166
Understanding RIP Trace Operations	4166
Example: Tracing RIP Protocol Traffic	4167
RIP Configuration Statements	4171
any-sender	4172
authentication-key (Protocols RIP)	4173
authentication-type (Protocols RIP)	4174
bfd-liveness-detection (Protocols RIP)	4175
check-zero	4178
export (Protocols RIP)	4179
group (Protocols RIP)	4180
holddown (Protocols RIP)	4182
import (Protocols RIP)	4183
message-size	4184
metric-in (Protocols RIP)	4185
metric-out (Protocols RIP)	4186
neighbor (Protocols RIP)	4187
preference (Protocols RIP)	4188
receive (Protocols RIP)	4189
rib-group (Protocols RIP)	4190
rip	4190
route-timeout (Protocols RIP)	4191
send (Protocols RIP)	4192
traceoptions (Protocols RIP)	4193
update-interval (Protocols RIP)	4196
Chapter 47 Administration	4197
Routine Monitoring	4197
Monitoring RIP Routing Information	4197
RIP Operational Commands	4197
clear rip general-statistics	4198
clear rip statistics	4199
show rip general-statistics	4200
show rip neighbor	4202
show rip statistics	4204

Part 16

Chapter 48

Multicast

Overview	4209
Introduction to PIM Basics	4209
PIM Overview	4209
Basic PIM Network Components	4211
PIM on Aggregated Interfaces	4212
Introduction to PIM Sparse Mode	4212
Understanding PIM Sparse Mode	4213
Rendezvous Point	4214
RP Mapping Options	4215
Designated Router	4215
Introduction to Static RP	4216
Understanding Static RP	4216
Introduction to Anycast RP	4216
Understanding RP Mapping with Anycast RP	4217
Introduction to PIM Bootstrap Router	4217
Understanding the PIM Bootstrap Router	4217
Introduction to PIM Filtering	4218
Understanding Multicast Message Filters	4218
Filtering MAC Addresses	4219
Filtering RP and DR Register Messages	4219
Introduction to PIM RPT and SPT Cutover	4220
Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees	4220
Building an RPT Between the RP and Receivers	4221
PIM Sparse Mode Source Registration	4222
Multicast Shortest-Path Tree	4225
SPT Cutover	4226
SPT Cutover Control	4229
Introduction to IGMP	4229
Understanding Group Membership Protocols	4229
Understanding IGMP	4230
Introduction to IGMP Snooping	4232
IGMP Snooping Overview	4232
How IGMP Snooping Works	4232
How IGMP Snooping Works with Routed VLAN Interfaces	4233
How Hosts Join and Leave Multicast Groups	4233
IGMP Snooping Support for IGMPv3	4233
IGMP Snooping and Forwarding Interfaces	4234
General Forwarding Rules	4234
Introduction to MSDP	4235
Understanding MSDP	4235
Filtering MSDP SA Messages	4236
Introduction to Source-Specific Multicast	4237
Source-Specific Multicast Groups Overview	4237
Understanding PIM Source-Specific Mode	4238
PIM SSM	4239

	Introduction to Multicast VLAN Registration	4241
	Understanding Multicast VLAN Registration	4241
	How MVR Works	4241
Chapter 49	Configuration	4245
	Optimizing Multicast Flows on QFabric Systems	4245
	Optimizing the Number of Multicast Flows on QFabric Systems	4246
	PIM Basics	4246
	Changing the PIM Version	4246
	Modifying the PIM Hello Interval	4246
	Preserving Multicast Performance by Disabling Response to the ping Utility	4247
	Configuring PIM Trace Options	4248
	Disabling PIM	4250
	Disabling the PIM Protocol	4251
	Disabling PIM On an Interface	4251
	Disabling PIM for a Family	4252
	Disabling PIM for a Rendezvous Point	4252
	PIM Designated Router	4253
	Configuring Interface Priority for PIM Designated Router Selection	4253
	Configuring PIM Designated Router Election on Point-to-Point Links	4254
	PIM Sparse Mode	4254
	Enabling PIM Sparse Mode	4255
	Configuring PIM Join Load Balancing	4256
	Modifying the Join State Timeout	4259
	Example: Enabling Join Suppression	4259
	Static RP	4264
	Configuring Local PIM RPs	4264
	Configuring the Static PIM RP Address on the Non-RP Routing Device	4266
	Anycast RP	4267
	Example: Configuring PIM Anycast With or Without MSDP	4267
	Configuring a PIM Anycast RP Router with MSDP	4271
	Configuring a PIM Anycast RP Router Using Only PIM	4272
	Configuring All PIM Anycast Non-RP Routers	4273
	Example: Configuring Multiple RPs in a Domain with Anycast RP	4273
	PIM Bootstrap Router	4276
	Configuring PIM Bootstrap Properties for IPv4 or IPv6	4276
	Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain	4277
	Example: Configuring PIM BSR Filters	4278
	PIM Filtering	4278
	Configuring Interface-Level PIM Neighbor Policies	4279
	Filtering Outgoing PIM Join Messages	4280
	Filtering Incoming PIM Join Messages	4281
	Configuring Register Message Filters on a PIM RP and DR	4282
	PIM RPT and SPT Cutover	4284
	Example: Configuring the PIM Assert Timeout	4284
	Example: Configuring the PIM SPT Threshold Policy	4287

PIM and the BFD Protocol	4290
Configuring BFD for PIM	4290
Configuring BFD Authentication for PIM	4292
Configuring BFD Authentication Parameters	4292
Viewing Authentication Information for BFD Sessions	4293
IGMP	4295
Configuring IGMP	4295
Enabling IGMP	4297
Changing the IGMP Version	4298
Modifying the IGMP Host-Query Message Interval	4299
Modifying the IGMP Last-Member Query Interval	4299
Specifying Immediate-Leave Host Removal for IGMP	4300
Filtering Unwanted IGMP Reports at the IGMP Interface Level	4301
Accepting IGMP Messages from Remote Subnetworks	4302
Modifying the IGMP Query Response Interval	4303
Modifying the IGMP Robustness Variable	4304
Limiting the Maximum IGMP Message Rate	4305
Enabling IGMP Static Group Membership	4305
Recording IGMP Join and Leave Events	4312
Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces	4313
Tracing IGMP Protocol Traffic	4314
Disabling IGMP	4316
IGMP Snooping	4316
Configuring IGMP Snooping	4317
Changing the IGMP Snooping Group Timeout Value	4318
Example: Configuring IGMP Snooping	4318
Configuring Multicast VLAN Registration (CLI Procedure)	4321
Example: Configuring Multicast VLAN Registration	4322
MSDP	4327
Configuring MSDP	4327
Tracing MSDP Protocol Traffic	4328
Configuring the Interface to Accept Traffic from a Remote Source	4330
Example: Configuring MSDP	4331
Example: Configuring MSDP with Active Source Limits and Mesh Groups	4332
Example: Configuring PIM Anycast With or Without MSDP	4338
Configuring a PIM Anycast RP Router with MSDP	4341
Source-Specific Multicast	4342
Example: Configuring PIM SSM on a Network	4342
Example: Configuring an SSM-Only Domain	4344
Example: Configuring SSM Mapping	4344
Example: Configuring Source-Specific Multicast Groups with Any-Source Override	4347
Example: Configuring SSM Maps for Different Groups to Different Sources	4350
Multiple SSM Maps and Groups for Interfaces	4350
Example: Configuring Multiple SSM Maps Per Interface	4351

PIM Configuration Statements	4354
address (Anycast RPs)	4356
address (Local RPs)	4357
address (Static RPs)	4358
algorithm	4359
anycast-pim	4360
assert-timeout	4361
authentication (Protocols PIM)	4362
bfd-liveness-detection (Protocols PIM)	4363
bootstrap	4364
bootstrap-export	4365
bootstrap-import	4366
bootstrap-priority	4367
detection-time (BFD for PIM)	4368
disable (PIM)	4369
dr-election-on-p2p	4370
dr-register-policy	4370
embedded-rp	4371
export (Protocols PIM Bootstrap)	4372
export (Protocols PIM)	4372
family (Bootstrap)	4373
family (Protocols PIM)	4374
family (Local RP)	4375
group (RPF Selection)	4376
group-ranges	4377
hello-interval (Protocols PIM)	4378
hold-time (Protocols PIM)	4379
import (Protocols PIM Bootstrap)	4380
import (Protocols PIM)	4381
infinity	4382
interface	4383
join-load-balance	4384
join-prune-timeout	4385
key-chain (Protocols PIM)	4386
local	4387
local-address (Protocols PIM)	4388
loose-check	4389
maximum-rps	4390
minimum-interval (PIM BFD Liveness Detection)	4391
minimum-interval (PIM BFD Transmit Interval)	4392
minimum-receive-interval	4393
mode (Protocols PIM)	4394
multiplier	4394
neighbor-policy	4395
next-hop (PIM RPF Selection)	4395
no-adaptation (PIM BFD Liveness Detection)	4396
override-interval	4397
pim	4398
prefix-list (PIM RPF Selection)	4401

priority (Bootstrap)	4402
priority (PIM Interfaces)	4403
priority (PIM RPs)	4404
propagation-delay	4405
reset-tracking-bit	4406
rib-group (Protocols PIM)	4407
rp	4408
rp-register-policy	4410
rp-set	4411
rpf-selection	4412
source (PIM RPF Selection)	4413
spt-threshold	4414
static (Protocols PIM)	4415
threshold (PIM BFD Detection Time)	4416
threshold (PIM BFD Transmit Interval)	4417
transmit-interval (PIM BFD Liveness Detection)	4418
traceoptions (Protocols PIM)	4419
version (BFD)	4422
version (PIM)	4423
wildcard-source (PIM RPF Selection)	4424
IGMP Configuration Statements	4424
accounting (Protocols IGMP)	4425
accounting (Protocols IGMP Interface)	4426
asm-override-ssm	4426
disable (Protocols IGMP)	4427
exclude (Protocols IGMP)	4427
group (Protocols IGMP)	4428
group-count (Protocols IGMP)	4429
group-increment (Protocols IGMP)	4429
group-limit	4430
group-policy (Protocols IGMP)	4431
igmp	4432
immediate-leave (Protocols IGMP)	4434
interface (Protocols IGMP)	4435
maximum-transmit-rate (Protocols IGMP)	4436
oif-map	4436
passive (IGMP)	4437
promiscuous-mode (Protocols IGMP)	4438
query-interval (Protocols IGMP)	4438
query-last-member-interval (Protocols IGMP)	4439
query-response-interval (Protocols IGMP)	4440
robust-count (Protocols IGMP)	4441
source (Protocols IGMP)	4442
source-count (Protocols IGMP)	4443
source-increment (Protocols IGMP)	4443
ssm-map (Protocols IGMP)	4444
ssm-map-policy (IGMP)	4444
static (Protocols IGMP)	4445
traceoptions (Protocols IGMP)	4446

version (Protocols IGMP)	4448
IGMP Snooping Configuration Statements	4448
data-forwarding	4449
disable (IGMP Snooping)	4450
group (IGMP Snooping)	4450
groups (Multicast VLAN Registration)	4451
igmp-snooping	4452
install (Multicast VLAN Registration)	4453
interface (IGMP Snooping)	4453
multicast-router-interface (IGMP Snooping)	4454
proxy (Multicast VLAN Registration)	4454
receiver	4455
robust-count (IGMP Snooping)	4455
source (Multicast VLAN Registration)	4456
source-vlans	4456
static (IGMP Snooping)	4457
traceoptions (IGMP Snooping)	4458
vlan (IGMP Snooping)	4460
version (IGMP Snooping)	4461
MSDP Configuration Statements	4461
active-source-limit	4462
authentication-key	4463
data-encapsulation	4464
default-peer	4465
disable (Protocols MSDP)	4466
export (Protocols MSDP)	4467
group	4468
import (Protocols MSDP)	4469
local-address	4470
maximum	4471
mode (Protocols MSDP)	4472
msdp	4473
peer (Protocols MSDP)	4475
rib-group (Protocols MSDP)	4476
source	4477
threshold	4478
traceoptions (Protocols MSDP)	4479
Source-Specific Multicast Configuration Statements	4481
asm-override-ssm	4482
policy (SSM Maps)	4483
ssm-groups	4484
ssm-map (Protocols IGMP)	4485
ssm-map (Routing Options Multicast)	4485
ssm-map-policy (IGMP)	4486

Chapter 50	Administration	4487
	Routine Monitoring	4487
	Monitoring IGMP Snooping	4487
	Verifying the IGMP Snooping Group Timeout Value	4488
	Monitoring Commands for Multicast Protocols	4488
	clear igmp membership	4491
	clear igmp-snooping membership	4494
	clear igmp statistics	4495
	clear igmp-snooping statistics	4497
	clear msdp cache	4498
	clear msdp statistics	4499
	clear multicast bandwidth-admission	4500
	clear multicast scope	4502
	clear multicast sessions	4503
	clear multicast statistics	4504
	clear pim join	4505
	clear pim register	4507
	clear pim statistics	4509
	mtrace	4512
	mtrace from-source	4515
	mtrace monitor	4518
	mtrace to-gateway	4520
	show configuration protocols igmp	4523
	show igmp group	4525
	show igmp interface	4529
	show igmp statistics	4533
	show igmp-snooping membership	4536
	show igmp-snooping route	4539
	show igmp-snooping statistics	4541
	show igmp-snooping vlans	4543
	show msdp	4545
	show msdp source	4547
	show msdp source-active	4549
	show msdp statistics	4552
	show multicast flow-map	4556
	show multicast interface	4558
	show multicast minfo	4560
	show multicast next-hops	4562
	show multicast pim-to-igmp-proxy	4565
	show multicast pim-to-mld-proxy	4567
	show multicast route	4569
	show multicast rpf	4575
	show multicast scope	4579
	show multicast sessions	4581
	show multicast usage	4584
	show pim bootstrap	4587
	show pim interfaces	4589
	show pim join	4592
	show pim neighbors	4601

	show pim rps	4605
	show pim source	4612
	show pim statistics	4615
	show system statistics igmp	4628
	test msdp	4632
Part 17	Security	
Chapter 51	Overview	4635
	Firewall Filters	4635
	Overview of Firewall Filters	4635
	Firewall Filter Types	4636
	Firewall Filter Components	4636
	Firewall Filter Processing	4637
	Understanding Filter-Based Forwarding	4637
	Understanding How Firewall Filters Are Evaluated	4638
	Understanding How Firewall Filters Control Packet Flows	4640
	Understanding Firewall Filter Match Conditions	4641
	Filter Match Conditions	4641
	Numeric Filter Match Conditions	4641
	Interface Filter Match Conditions	4642
	IP Address Filter Match Conditions	4642
	MAC Address Filter Match Conditions	4643
	Bit-Field Filter Match Conditions	4643
	Firewall Filter Match Conditions and Actions	4644
	Understanding How a Firewall Filter Tests a Protocol	4656
	Understanding Firewall Filter Planning	4657
	Planning the Number of Firewall Filters to Create	4658
	Understanding How Many Firewall Filters Are Supported	4658
	Egress Filters	4660
	Avoid Configuring too Many Filters	4660
	Policers can Limit Egress Filters	4660
	Planning for Filter-Specific Policers	4661
	Planning for Filter-Based Forwarding	4662
	Understanding Firewall Filter Processing Points for Bridged and Routed Packets	4662
	Applying Firewall Filters to Interfaces	4663
	Policers	4663
	Overview of Policers	4664
	Policer Overview	4664
	Policer Types	4664
	Policer Actions	4665
	Policer Colors	4666
	Filter-Specific Policers	4666
	Suggested Naming Convention for Policers	4667
	Policer Counters	4667
	Policer Algorithms	4667
	How Many Policers are Supported?	4668

Policers can Limit Egress Firewall Filters	4668
Understanding Policers with Link Aggregation Groups	4669
Understanding Color-Blind Mode for Single-Rate Tricolor Marking	4669
Understanding Color-Aware Mode for Single-Rate Tricolor Marking	4670
Summary of PLP Changes	4670
Understanding Color-Blind Mode for Two-Rate Tricolor Marking	4671
Understanding Color-Aware Mode for Two-Rate Tricolor Marking	4672
Summary of PLP Changes	4672
Effect on Green Packets (Low PLP)	4672
Effect on Yellow Packets (Medium PLP)	4673
Effect on Red Packets (High PLP)	4673
Port Security	4673
Overview of Access Port Protection	4674
Mitigation of Ethernet Switching Table Overflow Attacks	4674
Mitigation of Rogue DHCP Server Attacks	4674
Protection Against ARP Spoofing Attacks	4675
Protection Against DHCP Snooping Database Alteration Attacks	4675
Protection Against DHCP Starvation Attacks	4676
Port Security Overview	4676
Understanding DHCP Snooping for Port Security	4678
DHCP Snooping Basics	4679
DHCP Snooping Process	4680
DHCP Server Access	4681
DHCP Snooping Table	4684
Static IP Address Additions to the DHCP Snooping Database	4684
Snooping DHCP Packets That Have Invalid IP Addresses	4684
Prioritizing Snooped Packets	4685
Understanding DAI for Port Security	4686
Address Resolution Protocol	4686
ARP Spoofing	4686
Dynamic ARP Inspection (DAI)	4687
Prioritizing Inspected Packets	4687
Understanding MAC Limiting and MAC Move Limiting for Port Security	4688
MAC Limiting	4688
MAC Move Limiting	4689
Actions for MAC Limiting	4689
MAC Addresses That Exceed the MAC Limit or MAC Move Limit	4689
Understanding Trusted and Untrusted Ports	4690
Understanding Trusted DHCP Servers for Port Security	4690
Understanding DHCP Option 82 for Port Security	4691
DHCP Option 82 Processing	4691
Suboption Components of Option 82	4692
Configurations That Support Option 82	4692
Understanding Static ARP Entries	4693
Device Security	4694
Understanding Storm Control	4694

Chapter 52	Configuration	4697
	Firewall and Policer Configuration Examples	4697
	Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device	4697
	Example: Using Two-Color Policers and Prefix Lists	4701
	Example: Using Policers to Manage Oversubscription	4704
	Port Security Configuration Examples	4706
	Example: Configuring Basic Port Security Features	4706
	Example: Configuring Storm Control	4713
	Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks	4715
	Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks	4718
	Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks	4722
	Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch	4725
	Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks	4732
	Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks	4737
	Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server	4740
	Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server	4744
	Firewall and Policer Configuration Tasks	4746
	Configuring Firewall Filters	4747
	Configuring a Firewall Filter	4747
	Applying a Firewall Filter to a Port	4749
	Applying a Firewall Filter to a VLAN	4749
	Applying a Firewall Filter to a Layer 3 (Routed) Interface	4749
	Applying Firewall Filters to Interfaces	4750
	Assigning Forwarding Classes and Loss Priority	4751
	Configuring Color-Blind Egress Policers for Medium-Low PLP	4752
	Configuring Two-Color and Three-Color Policers to Control Traffic Rates	4753
	Configuring Two-Color Policers	4753
	Configuring Three-Color Policers	4754
	Specifying Policers in a Firewall Filter Configuration	4754
	Applying a Firewall Filter That Includes a Policer	4755
	Port Security Configuration Tasks	4755
	Configuring Port Security (CLI Procedure)	4756
	Enabling DHCP Snooping	4757
	Enabling Dynamic ARP Inspection (DAI)	4757
	Limiting Dynamic MAC Addresses on an Interface	4757
	Enabling Persistent MAC Learning on an Interface	4757
	Limiting MAC Address Movement	4758

Configuring Trusted DHCP Servers on an Interface	4758
Configuring MAC Limiting	4758
Configuring MAC Move Limiting (CLI Procedure)	4760
Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)	4762
Configuring the none Action to Override a MAC Limit Applied to All Interfaces (CLI Procedure)	4762
Configuring Static ARP Entries	4763
Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)	4763
Enabling DHCP Snooping (CLI Procedure)	4764
Enabling DHCP Snooping	4765
Applying CoS Forwarding Classes to Prioritize Snooped Packets	4765
Enabling Dynamic ARP Inspection (CLI Procedure)	4766
Enabling DAI	4767
Applying CoS Forwarding Classes to Prioritize Inspected Packets . . .	4767
Enabling a Trusted DHCP Server (CLI Procedure)	4768
Enabling a Trusted Port for DHCP	4769
Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)	4770
Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)	4773
Configuration Statements for Firewall Filters	4775
family	4776
filter	4777
filter (Layer 2 and Layer 3 Interfaces)	4778
filter (VLANs)	4779
firewall	4780
from	4781
interface-specific	4782
term	4782
then (Filters)	4783
Configuration Statements for Policers	4783
action	4784
bandwidth-limit	4784
burst-size-limit	4785
color-aware	4786
color-blind	4787
committed-burst-size	4788
committed-information-rate	4789
excess-burst-size	4790
filter-specific	4791
firewall	4792
if-exceeding	4793
loss-priority high then discard (Three-Color Policer)	4794
peak-burst-size	4795
peak-information-rate	4796
policer	4797
single-rate	4798

then (Policers)	4799
three-color-policer	4800
two-rate	4801
Configuration Statements for Port Security	4801
allowed-mac	4803
arp-inspection	4804
circuit-id	4805
dhcp-trusted	4806
dhcp-option82	4807
dhcp-snooping-file	4808
dhcp-trusted	4808
disable-timeout (Port Error Disable)	4809
ethernet-switching-options	4810
examine-dhcp	4812
examine-fip	4813
fc-map	4814
fcoe-trusted	4815
forwarding-class (for DHCP Snooping or DAI Packets)	4816
interface (Secure Access Port)	4817
location	4818
mac	4818
mac-limit	4819
mac-move-limit	4820
no-allowed-mac-log	4821
no-dhcp-trusted	4821
no-gratuitous-arp-request	4822
persistent-learning	4822
port-error-disable	4823
prefix (Remote ID for Option 82)	4824
remote-id	4825
secure-access-port	4826
static-ip	4827
timeout (DHCP Snooping)	4828
use-interface-description	4829
use-string	4830
use-vlan-id	4831
vendor-id	4832
vlan (Secure Access Port)	4833
vlan (Static IP)	4834
write-interval	4835
Configuration Statements for Device Security	4835
action-shutdown	4836
bandwidth	4837
ethernet-switching-options	4838
interface (Storm Control)	4840
no-broadcast	4841
no-multicast	4841
no-unknown-unicast	4842
storm-control	4843

Chapter 53	Administration	4845
	Routine Monitoring	4845
	Monitoring Firewall Filter Traffic	4845
	Monitoring Traffic for All Firewall Filters and Policers That Are Configured	4845
	Monitoring Traffic for a Specific Firewall Filter	4846
	Monitoring Traffic for a Specific Policer	4846
	Monitoring Port Security	4847
	Verifying That Firewall Filters Are Operational	4848
	Verifying That DAI Is Working Correctly	4848
	Verifying That DHCP Snooping Is Working Correctly	4849
	Verifying That MAC Limiting Is Working Correctly	4850
	Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly	4851
	Verifying That Allowed MAC Addresses Are Working Correctly	4851
	Verifying That Interfaces Are Shut Down	4852
	Customizing the Ethernet Switching Table Display to View Information for a Specific Interface	4852
	Verifying That MAC Move Limiting Is Working Correctly	4853
	Verifying That the Port Error Disable Setting Is Working Correctly	4854
	Verifying That a Trusted DHCP Server Is Working Correctly	4854
	Verifying That Three-Color Policers Are Operational	4855
	Verifying That Two-Color Policers Are Operational	4856
	Monitoring Commands	4856
	clear arp inspection statistics	4857
	clear dhcp snooping binding	4858
	clear ethernet-switching port-error	4859
	clear firewall	4860
	show arp inspection statistics	4861
	show dhcp snooping binding	4862
	show firewall	4864
	show firewall policer	4868
	show interfaces filters	4870
Chapter 54	Troubleshooting	4873
	Troubleshooting Procedures	4873
	Troubleshooting Firewall Filter Configuration	4873
	Firewall Filter Configuration Returns a No Space Available in TCAM Message	4873
	Filter Counts Previously Dropped Packet	4875
	Matching Packets Not Counted	4876
	Counter Reset When Editing Filter	4876
	Cannot Include loss-priority and policer Actions in Same Term	4876
	Cannot Egress Filter Certain Traffic Originating on QFX Switch	4877
	Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	4877
	Egress Firewall Filters with Private VLANs	4877
	Egress Filtering of L2PT Traffic Not Supported	4878
	Cannot Drop BGP Packets in Certain Circumstances	4878
	Invalid Statistics for Policer	4878

	Policers can Limit Egress Filters	4878
	Troubleshooting Policer Configuration	4880
	Incomplete Count of Packet Drops	4880
	Counter Reset When Editing Filter	4880
	Invalid Statistics for Policer	4880
	Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	4881
	Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	4882
	Policers Can Limit Egress Filters	4882
Part 18	Services	
Chapter 55	Overview	4887
	Port Mirroring	4887
	Understanding Port Mirroring	4887
	Port Mirroring Overview	4887
	Port-Mirroring Terminology	4888
	Port Mirroring Constraints and Limitations	4889
	Understanding Layer 3 Logical Interfaces	4892
	DHCP Relay	4892
	DHCP and BOOTP Relay Overview	4892
Chapter 56	Configuration	4895
	Configuration Examples	4895
	Example: Configuring Port Mirroring for Local Analysis	4895
	Example: Configuring Port Mirroring for Remote Analysis	4900
	Configuration Tasks	4904
	Configuring Port Mirroring	4904
	Configuring Port Mirroring for Local Analysis	4905
	Configuring Port Mirroring for Remote Analysis	4905
	Filtering the Traffic Entering an Analyzer	4906
	Configuring DHCP and BOOTP Relay	4907
	Configuring a DHCP and BOOTP Relay Agent	4907
	Configuring DHCP Smart Relay	4909
	Configuration Statements for Port Mirroring	4910
	analyzer	4911
	egress	4912
	ethernet-switching-options	4913
	ingress (ethernet-switching-options)	4915
	input	4916
	interface (Port Mirroring)	4917
	ip-address (Port Mirroring)	4918
	output	4918
	vlan (Port Mirroring)	4919
	Configuration Statements for Encryption	4919
	authentication-key-chains	4920
	cache-size	4921
	cache-timeout-negative	4922
	ca-name	4922

	certificates	4923
	certification-authority	4924
	crl (Encryption Interface)	4924
	encoding	4925
	enrollment-retry	4925
	enrollment-url	4926
	file	4926
	key (Authentication Keychain)	4927
	key-chain (Security)	4928
	ldap-url	4929
	local	4930
	maximum-certificates	4931
	path-length	4931
	secret	4932
	security	4933
	ssh-known-hosts	4934
	start-time (Authentication Key Transmission)	4935
	traceoptions	4937
	Configuration Statements for DHCP Relay	4938
	apply-secondary-as-giaddr	4939
	bootp	4940
	broadcast	4941
	client-response-ttl	4941
	description (Forwarding Options)	4942
	interface (BOOTP)	4943
	maximum-hop-count	4944
	minimum-wait-time	4944
	no-listen	4945
	server (DHCP and BOOTP Relay Agent)	4945
Chapter 57	Administration	4947
	Monitoring Commands for Port Mirroring	4947
	show analyzer	4948
Chapter 58	Troubleshooting	4949
	Troubleshooting Procedures	4949
	Troubleshooting Port Mirroring	4949
	Port Mirroring Constraints and Limitations	4949
	Egress Port Mirroring with VLAN Translation	4951
	Egress Port Mirroring with Private VLANs	4951
Part 19	Storage	
Chapter 59	Overview	4955
	Software Features Overview	4955
	Overview of Fibre Channel on the QFX Series	4956
	Fibre Channel Transport Protocol	4956
	How FC Works on the QFX Series	4957
	Supported FC Features and Functions	4959

Lossless Transport Support	4959
Overview of FIP	4960
Fibre Channel, FCoE, and FIP	4961
Understanding Fibre Channel	4961
FC Fabrics	4962
FC Port Types	4962
FC Switches	4962
Adapters	4963
N_Port ID Virtualization (NPIV)	4963
FC Services	4963
Understanding DCB Features and Requirements	4965
Lossless Transport	4965
ETS	4966
DCBX	4967
Understanding FCoE	4968
FCoE Devices	4968
FCoE Frames	4970
Virtual Links	4970
FCoE VLANs	4971
Understanding FCoE Transit Switch Functionality	4972
Understanding an FCoE-FC Gateway	4975
Gateway FC Fabric	4976
Fabric Services	4977
FCoE-FC Gateway Traffic Switching	4977
Understanding FCoE-FC Gateway Functions	4979
Login and Logout	4979
FCoE and FC Frame Handling	4979
Data Center Bridging	4979
Disabling the Fabric WWN Verification Check	4980
Load Balancing	4981
Understanding FCoE and FIP Session High Availability	4981
High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode)	4982
High Availability for FIP Snooping	4982
Nonstop Software Upgrade (QFabric Systems)	4983
Understanding FIP Functions	4984
FIP VLAN Discovery	4984
FIP Discovery	4985
FIP FLOGI	4986
FIP FDISC	4986
FIP Maintenance (Keepalive Messages)	4987
FIP LOGO	4987
Understanding FIP Implementation	4988
FIP Basics	4988
Fabric Login and FIP Login Overview	4989
Proxy FIP Discovery	4990
Proxy FIP Initialization	4990
Proxy FIP Maintenance	4991
Proxy FIP Logout	4991

Understanding FIP Parameters on an FCoE-FC Gateway	4992
FIP Keepalive Advertisement Period	4992
Addressing Mode	4992
FC-MAP	4993
FCoE Trusted Fabric	4994
Maximum Number of FCoE Sessions Per ENode	4994
Priority	4994
Understanding Fibre Channel Virtual Links	4995
Understanding Interfaces on an FCoE-FC Gateway	4996
Native FC Interfaces to the FC Switch	4996
FIP Login Session Limits	4998
Trusted and Untrusted Interfaces	5001
Buffer-to-Buffer Credit Recovery	5002
FCoE VLAN Interface to FCoE Devices	5003
Assigning Interfaces to a Fibre Channel Fabric	5006
Deleting a Fibre Channel Interface	5006
Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric . . .	5008
Load-Balancing Algorithms	5009
Load-Rebalancing Methods	5013
NP_Port Interface FIP Session Limit Effect on Load Balancing	5014
Load-Balancing Triggers and Timing	5014
Load Rebalancing Behavior When a Link Goes Down	5016
Interface Load Calculation Algorithm	5017
Load-Balancing Scenarios	5018
Load Balancing on the FCoE Interfaces (Ethernet Links)	5023
Understanding OxID Hash Control for FCoE Traffic Load Balancing	5024
Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit	
Switch	5025
FC Network Security	5026
VN2VF_Port FIP Snooping Functions	5026
FIP Snooping Firewall Filters	5027
Session Scalability	5027
VN2VF_Port FIP Snooping Implementation	5027
T11 VN2VF_Port FIP Snooping Specification	5030
Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit	
Switch	5031
VN2VN_Port FIP Snooping and FIP Snooping Virtual Links	5031
VN2VN_Port Communication Modes	5032
Network Security	5033
VN2VN_Port FIP Snooping Functions	5033
Scalability	5033
VN2VN_Port FIP Snooping Implementation	5033
ENode-Facing Interfaces	5034
Network-Facing Interfaces (Connecting to Another Transit Switch) . .	5035
Beacon Period (VN2VN_Port FIP Snooping Link Maintenance)	5035
QFabric System Differences in VN2VN_Port FIP Snooping Traffic	
Handling	5036

Understanding FIP Snooping, FBF, and MVR Filter Scalability	5038
VFP TCAM Architecture and Allocation	5038
VFP TCAM Entry Consumption	5039
Rejected Filter Configurations (No Available VFP TCAM Space)	5042
VFP TCAM Allocation and Consumption (Scaling) Examples	5043
Filter Configuration Recommendations	5045
Understanding MC-LAGs on an FCoE Transit Switch	5047
Supported Topology	5047
FIP Snooping and FCoE Trusted Ports	5049
CoS and Data Center Bridging (DCB)	5050
Understanding DCBX	5051
DCBX Basics	5051
DCBX Modes and Support	5052
DCBX Attribute Types	5055
DCBX Application Protocol TLV Exchange	5056
DCBX and PFC	5057
DCBX and ETS	5057
Understanding DCBX Application Protocol TLV Exchange	5060
Applications	5060
Application Maps	5061
Classifying and Prioritizing Application Traffic	5062
Enabling Interfaces to Exchange Application Protocol Information . .	5063
Disabling DCBX Application Protocol Exchange	5063
Understanding CoS Flow Control (Ethernet PAUSE and PFC)	5064
Ethernet PAUSE	5064
PFC	5068
Lossless Transport Support Summary	5072
Understanding Fibre Channel Terminology	5074
QFabric Specific	5084
Understanding Fibre Channel Fabrics on the QFabric System	5084
Understanding CoS Fabric Forwarding Class Sets	5086
Default Fabric Forwarding Class Sets	5087
Fabric Forwarding Class Set Configuration and Implementation . . .	5090
Fabric Forwarding Class Set Scheduling (CoS)	5092
Support for Flow Control and Lossless Transport Across the Fabric . .	5094
Viewing Fabric Forwarding Class Set Information	5096
Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences	5098
Chapter 60 Configuration	5099
Configuration Examples	5099
Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric	5099
Example: Configuring CoS PFC for FCoE Traffic	5113
Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG	5121
Example: Configuring DCBX Application Protocol TLV Exchange	5144
Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)	5155

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)	5159
Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)	5167
Example: Configuring Automated Fibre Channel Interface Load Rebalancing	5175
Configuration Tasks	5178
Configuring an FCoE-FC Gateway Fibre Channel Fabric	5179
Disabling the Fabric WWN Verification Check	5180
Configuring a Physical Fibre Channel Interface	5181
Configuring a Fibre Channel Interface	5182
Configuring an FCoE VLAN Interface	5185
Assigning Interfaces to a Fibre Channel Fabric	5187
Deleting a Fibre Channel Interface	5188
Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway	5189
Defining the Proxy Load-Balancing Algorithm	5190
Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test)	5191
Enabling and Disabling CoS OxID Hash Control	5192
Configuring FIP on an FCoE-FC Gateway	5192
Setting the Maximum Number of FIP Login Sessions per ENode	5195
Setting the Maximum Number of FIP Login Sessions per FC Interface	5196
Setting the Maximum Number of FIP Login Sessions per FC Fabric	5197
Setting the Maximum Number of FIP Login Sessions per Node Device	5198
Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch	5199
Disabling VN2VF_Port FIP Snooping on an FCoE-FC Gateway Switch Interface	5201
Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch	5202
Configuring the DCBX Mode	5203
Configuring DCBX Autonegotiation	5204
Disabling the ETS Recommendation TLV	5207
Defining an Application for DCBX Application Protocol TLV Exchange	5207
Configuring an Application Map for DCBX Application Protocol TLV Exchange	5209
Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	5210
Configuration Statements	5210
application (Application Maps)	5212
application (Applications)	5213
application-map	5214
application-maps	5215
applications (Applications)	5216
applications (DCBX)	5217
auto-load-rebalance	5217
bb-sc-n	5218
beacon-period	5219
code-points (Application Maps)	5220
dcbx	5221

dcbx-version	5222
description (Fibre Channel Fabrics)	5223
destination-port (Applications)	5224
disable (DCBX)	5225
enhanced-transmission-selection	5226
ether-type	5227
examine-fip	5228
examine-vn2vn	5229
fabric-id	5230
fabric-type	5230
family fcoe	5231
fc2	5231
fc-fabrics	5232
fc-map	5234
fc-options	5235
fcoe-trusted	5236
fibre-channel (Family Interfaces)	5237
fibre-channel (Port)	5238
fibrechannel-options	5238
fip	5239
fka-adv-period	5240
interface (DCBX)	5241
interface (Fibre Channel Fabric)	5242
interface (FIP)	5243
load-balance-algorithm	5244
loopback (Fibre Channel Interface)	5245
max-login-sessions	5246
max-login-sessions-per-node	5247
max-sessions-per-enode	5248
no-fabric-wwn-verify	5249
oxid	5250
policy-options	5251
port-mode (Fibre Channel Interfaces)	5252
port-range	5253
priority (FIP)	5254
priority-flow-control	5255
protocol (Applications)	5256
protocols (FIP)	5257
proxy (Fibre Channel)	5258
recommendation-tlv	5258
speed (Fibre Channel Interfaces)	5259
traceoptions (FC-2 Fibre Channel)	5260
traceoptions (Fibre Channel)	5262
traceoptions (FIP Protocol Fibre Channel)	5265
traceoptions (Proxy Fibre Channel)	5267

Chapter 61	Administration	5269
	Routine Monitoring	5269
	Monitoring Fibre Channel Interface Load Balancing	5269
	Monitoring the Interface Load-Balancing State	5269
	Monitoring the Fabric Load-Balancing Algorithm	5270
	Operational Commands	5274
	clear fibre-channel fc2 statistics	5276
	clear fibre-channel fip enode	5277
	clear fibre-channel fip statistics	5278
	clear fibre-channel fip vn-port	5279
	clear fibre-channel flogi statistics	5280
	clear fibre-channel proxy statistics	5281
	clear fip snooping enode	5282
	clear fip snooping statistics	5283
	clear fip snooping vlan	5284
	clear fip vlan-discovery statistics	5285
	request fibre-channel proxy load-rebalance	5286
	restart	5288
	show dcbx	5299
	show dcbx neighbors	5300
	show fibre-channel fabric	5322
	show fibre-channel fc2 sessions	5324
	show fibre-channel fc2 statistics	5326
	show fibre-channel fip	5328
	show fibre-channel fip enode	5333
	show fibre-channel fip fabric	5337
	show fibre-channel fip fcf	5340
	show fibre-channel fip interface	5343
	show fibre-channel fip statistics	5346
	show fibre-channel flogi fport	5350
	show fibre-channel flogi nport	5352
	show fibre-channel flogi statistics	5354
	show fibre-channel interfaces	5357
	show fibre-channel next-hops	5360
	show fibre-channel routes	5362
	show fibre-channel proxy fabric-state	5364
	show fibre-channel proxy login-table	5368
	show fibre-channel proxy np-port	5371
	show fibre-channel proxy statistics	5374
	show fip snooping	5377
	show fip snooping enode	5381
	show fip snooping fcf	5385
	show fip snooping interface	5388
	show fip snooping statistics	5391
	show fip snooping vlan	5394
	show fip vlan-discovery	5398
	show route forwarding-table family fibre-channel	5400

Chapter 62	Troubleshooting	5403
	Troubleshooting Procedures	5403
	Troubleshooting Dropped FCoE Traffic	5403
	Troubleshooting Fibre Channel Interface Deletion	5406
	Troubleshooting Dropped FIP Traffic	5407
Part 20	Traffic Management	
Chapter 63	Overview	5411
	Software Features Overview	5411
	Overview of Junos OS CoS for the QFX Series	5412
	CoS Standards	5412
	How Junos CoS Works	5413
	Default CoS Behavior	5414
	Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)	5414
	Overview of CoS Upgrade Requirements to Junos OS Release 12.2	5416
	Overview of CoS Upgrade Requirements to Junos OS Release 12.3 (QFX3500 and QFX3600 Switches) or to Junos OS Release 13.1 (QFabric Systems)	5418
	Support for Six Lossless Forwarding Classes	5418
	Scheduling on QFabric System Node Device Fabric (fte) Ports	5419
	Strict-High Priority Scheduling on QFabric System Node Device Fabric (fte) Ports	5420
	Overview of CoS Changes Introduced in Junos OS Release 11.3	5420
	CoS Default Value Changes	5420
	Queue Priority Configuration Changes	5425
	Minimum Guaranteed Bandwidth (Transmit Rate and Guaranteed Rate) Changes	5426
	Excess Rate Statement Disabled	5426
	Queue Scheduling (Low and Strict-High Priority Queues)	5427
	Multidestination Traffic Changes	5427
	Overview of CoS Changes Introduced in Junos OS Release 12.2	5428
	Lossless Forwarding Classes (fcoe and no-loss)	5428
	Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes	5429
	CoS for Layer 3 Physical Interfaces	5429
	DSCP IPv6 Classifiers and Rewrite Rules	5429
	Overview of Policers	5430
	Policer Overview	5430
	Policer Types	5431
	Policer Actions	5432
	Policer Colors	5432
	Filter-Specific Policers	5433
	Suggested Naming Convention for Policers	5433
	Policer Counters	5434
	Policer Algorithms	5434
	How Many Policers are Supported?	5434

Policers can Limit Egress Firewall Filters	5434
CoS Overview	5435
Understanding Junos CoS Components	5436
Code-Point Aliases	5436
Policers	5436
Classifiers	5437
Forwarding Classes	5437
Forwarding Class Sets	5437
Flow Control (Ethernet PAUSE and Priority-Based Flow Control)	5438
Tail-Drop Profiles	5438
Schedulers	5439
Rewrite Rules	5439
Understanding CoS Packet Flow	5440
CoS Inputs and Outputs Overview	5442
Understanding Default CoS Settings	5443
Default Forwarding Classes and Queue Mapping	5443
Default Forwarding Class Sets (Priority Groups)	5444
Default Code-Point Aliases	5444
Default Classifiers	5446
Default Rewrite Rules	5449
Default Drop Profile	5449
Default Schedulers	5449
Default Scheduler Maps	5449
Default Shared Buffer Configuration	5450
Understanding Host Inbound Traffic Classification	5451
Understanding Host Routing Engine Outbound Traffic Queues and Defaults	5451
Understanding CoS Code-Point Aliases	5453
Default Code-Point Aliases	5453
Understanding CoS Classifiers	5455
Interfaces and Output Queues	5455
Behavior Aggregate Classifiers	5456
Fixed Classifiers on Ethernet Interfaces	5459
Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports)	5459
Multifield Classifiers	5459
Packet Classification for Routed VLAN Interfaces (RVIs)	5460
Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces	5460
Supported Classifier and Rewrite Rule Types	5461
Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration	5462
Default Classifiers	5463
Default Rewrite Rules	5464
Classifier Precedence	5464
Classifier Behavior and Limitations	5465
Rewrite Rule Precedence and Behavior	5466
Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration	5466

Understanding CoS Forwarding Classes	5469
Default Forwarding Classes	5469
Forwarding Class Configuration Rules	5471
Lossless Transport Support	5472
Understanding CoS Forwarding Class Sets (Priority Groups)	5474
Understanding Default CoS Scheduling and Classification	5475
Default Classification	5475
Default Scheduling	5478
Default DCBX Advertisement	5479
Default Scheduling and Classification Summary	5479
Understanding CoS Hierarchical Port Scheduling (ETS)	5481
Hierarchical Scheduling and ETS	5482
ETS Advertisement in DCBX	5483
Hierarchical Scheduling Process	5483
Strict-High Priority Queues and Hierarchical Scheduling	5484
Default Hierarchical Scheduling	5485
Understanding CoS Output Queue Schedulers	5486
Output Queue Scheduling Components	5486
Default Schedulers	5487
Transmit Rate (Minimum Guaranteed Bandwidth)	5488
Sharing Extra Bandwidth	5489
Shaping Rate (Maximum Bandwidth)	5489
Scheduling Priority	5490
Scheduler Drop-Profile Maps	5491
Buffer Size	5491
Scheduler Maps	5492
Understanding CoS Priority Group Scheduling	5493
Priority Group Scheduling Components	5493
Default Traffic Control Profile	5494
Guaranteed Rate (Minimum Guaranteed Bandwidth)	5494
Sharing Extra Bandwidth	5494
Shaping Rate (Maximum Bandwidth)	5495
Scheduler Maps	5495
Understanding CoS Traffic Control Profiles	5496
Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum Bandwidth)	5497
Guaranteeing Bandwidth Using Hierarchical Scheduling	5497
Priority Group Guaranteed Rate (Minimum Bandwidth)	5499
Queue Transmit Rate (Minimum Bandwidth)	5499
Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)	5500
Priority Group Shaping	5500
Queue Shaping	5500
Shaping Maximum Bandwidth Using Hierarchical Scheduling	5501
Understanding CoS Scheduling Behavior and Configuration Considerations	5502
Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows	5506
Lossless Transport Features Introduced in Junos OS Release 12.3	5506
Default Lossless Priority Configuration	5507

Configuring Lossless Priorities	5509
Backward Compatibility with Junos OS Releases Earlier Than Release 12.3	5522
Configuration Rules and Recommendations	5523
Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway	5524
Priority Remapping Configuration	5524
Configuration Rules	5525
Fate Sharing	5526
Understanding CoS Buffer Configuration	5527
Buffer Pools	5528
Default Buffer Pool Values	5536
Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios	5538
Optimizing the Buffer Configuration	5542
Understanding CoS Tail-Drop Profiles	5545
Drop Profile Parameters	5545
Default Drop Profile	5546
Packet Drop Method	5547
Drop Profile Maps	5548
Congestion Prevention	5548
Understanding CoS Rewrite Rules	5549
Understanding DCB Features and Requirements	5550
Lossless Transport	5551
ETS	5552
DCBX	5552
Understanding CoS Flow Control (Ethernet PAUSE and PFC)	5554
Ethernet PAUSE	5554
PFC	5558
Lossless Transport Support Summary	5562
Understanding DCBX	5564
DCBX Basics	5564
DCBX Modes and Support	5565
DCBX Attribute Types	5568
DCBX Application Protocol TLV Exchange	5569
DCBX and PFC	5570
DCBX and ETS	5571
Understanding DCBX Application Protocol TLV Exchange	5573
Applications	5573
Application Maps	5574
Classifying and Prioritizing Application Traffic	5575
Enabling Interfaces to Exchange Application Protocol Information	5576
Disabling DCBX Application Protocol Exchange	5576
QFabric-Specific CoS Overview	5577
Understanding CoS Fabric Forwarding Class Sets	5578
Default Fabric Forwarding Class Sets	5579
Fabric Forwarding Class Set Configuration and Implementation	5582
Fabric Forwarding Class Set Scheduling (CoS)	5584
Support for Flow Control and Lossless Transport Across the Fabric	5586

	Viewing Fabric Forwarding Class Set Information	5588
	Summary of Fabric Forwarding Class Set and Node Device Forwarding	
	Class Set Differences	5590
	Understanding CoS Scheduling on QFabric System Node Device Fabric (fte)	
	Ports	5591
	Hierarchical Scheduling Architecture on QFabric System Node	
	Devices	5591
	Default Scheduling on Node Device Fabric Interfaces	5592
	Configuring Scheduling on Node Device Fabric Interfaces	5593
	Understanding Default CoS Scheduling on QFabric System Interconnect	
	Devices (Junos OS Release 13.1 and Later Releases)	5595
	Hierarchical CoS Architecture Across a QFabric System Interconnect	
	Device	5595
	Default CoS on Interconnect Device Fabric Interfaces	5597
Chapter 64	Configuration	5605
	Configuration Examples	5605
	Example: Configuring CoS Hierarchical Port Scheduling (ETS)	5606
	Example: Configuring CoS PFC for FCoE Traffic	5627
	Example: Configuring CoS for FCoE Transit Switch Traffic Across an	
	MC-LAG	5635
	Example: Configuring Unicast Classifiers	5658
	Example: Configuring Multidestination (Multicast, Broadcast, DLF)	
	Classifiers	5661
	Example: Configuring Tail-Drop Profiles	5663
	Example: Configuring Drop Profile Maps	5666
	Example: Configuring Forwarding Classes	5668
	Example: Configuring Forwarding Class Sets	5671
	Example: Configuring Queue Schedulers	5674
	Example: Configuring Queue Scheduling Priority	5679
	Example: Configuring Traffic Control Profiles (Priority Group	
	Scheduling)	5682
	Example: Configuring Minimum Guaranteed Output Bandwidth	5684
	Example: Configuring Maximum Output Bandwidth	5689
	Example: Configuring Lossless FCoE Traffic When the Converged Ethernet	
	Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE	
	Transit Switch)	5693
	Example: Configuring Two or More Lossless FCoE Priorities on the Same	
	FCoE Transit Switch Interface	5701
	Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on	
	Different FCoE Transit Switch Interfaces	5710
	Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces	
	for Multiple Applications (FCoE and iSCSI)	5723
	Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC	
	Gateway	5739
	Example: Recommended Configuration of the Shared Buffer Pool for	
	Networks with Mostly Best-Effort Unicast Traffic	5748

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled	5754
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic	5759
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic	5764
Example: Configuring DCBX Application Protocol TLV Exchange	5770
Configuration Tasks	5781
Configuring CoS	5781
Defining CoS Code-Point Aliases	5783
Defining CoS Unicast BA Classifiers	5784
Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers	5785
Configuring CoS Tail-Drop Profiles	5786
Configuring CoS Drop Profile Maps	5787
Defining CoS Forwarding Classes	5787
Defining CoS Forwarding Class Sets	5789
Defining CoS Queue Schedulers	5789
Defining CoS Queue Scheduling Priority	5792
Changing the Host Outbound Traffic Default Queue Mapping	5793
Defining CoS Traffic Control Profiles (Priority Group Scheduling)	5794
Configuring CoS PFC (Congestion Notification Profiles)	5795
Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control ...	5798
Configuring CoS Asymmetric Ethernet PAUSE Flow Control	5799
Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP_Ports)	5801
Configuring Global Ingress and Egress Shared Buffers	5803
Defining CoS Rewrite Rules	5805
Assigning CoS Components to Interfaces	5807
Configuring the DCBX Mode	5808
Configuring DCBX Autonegotiation	5809
Disabling the ETS Recommendation TLV	5812
Defining an Application for DCBX Application Protocol TLV Exchange ...	5812
Configuring an Application Map for DCBX Application Protocol TLV Exchange	5814
Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	5815
Configuration Statements	5815
application (Application Maps)	5818
application (Applications)	5819
application-map	5820
application-maps	5821
applications (Applications)	5822
applications (DCBX)	5823
buffer-partition (Egress)	5824
buffer-partition (Ingress)	5826
buffer-size	5828
cable-length (Congestion Notification)	5830

class-of-service	5831
class (Forwarding Classes)	5835
class (Forwarding Class Sets)	5836
classifiers	5837
code-point (Fibre Channel Interfaces)	5838
code-point (Input Congestion Notification)	5839
code-point (Output Congestion Notification)	5840
code-point (Rewrite Rules)	5841
code-point-aliases	5841
code-points (Application Maps)	5842
code-points (CoS)	5842
configured-flow-control	5843
congestion-notification-profile	5844
dcbx	5846
dcbx-version	5847
destination-port (Applications)	5848
disable (DCBX)	5849
drop-probability	5850
drop-profile	5851
drop-profile-map	5851
drop-profiles	5852
dscp	5853
dscp-ipv6	5855
dscp-code-point	5856
egress (Buffer Configuration)	5857
enhanced-transmission-selection	5858
ether-type	5859
fill-level	5860
flow-control	5861
flow-control-queue (Output Congestion Notification)	5862
forwarding-class	5863
forwarding-class (Fibre Channel Interfaces)	5864
forwarding-class (Host Outbound Traffic)	5865
forwarding-class-set (Node Device)	5866
forwarding-class-sets	5866
forwarding-classes	5867
guaranteed-rate	5869
host-outbound-traffic	5870
ieee-802.1	5871
ieee-802.1 (Fibre Channel Interfaces)	5872
ieee-802.1 (Input Congestion Notification)	5873
ieee-802.1 (Output Congestion Notification)	5874
import	5875
ingress (Buffer Configuration)	5876
input (Congestion Notification)	5877
input (Fibre Channel Interfaces)	5878
interface (DCBX)	5879
interfaces (Class of Service)	5880
interpolate	5881

loss-priority (Classifiers)	5882
loss-priority (Drop Profiles)	5883
loss-priority (Rewrite Rules)	5884
multi-destination	5885
mru	5886
output (Congestion Notification)	5887
output-traffic-control-profile	5888
pfc (Input Congestion Notification)	5889
policy-options	5890
priority (Schedulers)	5891
priority-flow-control	5892
protocol (Applications)	5893
protocol (Drop Profile Map)	5894
queue-num	5895
recommendation-tlv	5896
rewrite-rules	5897
rewrite-value (Fibre Channel Interfaces)	5898
rx-buffers	5900
scheduler (Node Device)	5901
scheduler-map	5901
scheduler-maps	5902
schedulers	5903
shaping-rate	5904
shared-buffer	5906
traceoptions (Class of Service)	5908
traffic-control-profiles	5910
transmit-rate	5911
tx-buffers	5913
unit	5914
Chapter 65 Administration	5915
Routine Monitoring	5915
Monitoring CoS Classifiers	5915
Monitoring CoS Forwarding Classes	5916
Monitoring Interfaces That Have CoS Components	5917
Monitoring CoS Rewrite Rules	5918
Monitoring CoS Scheduler Maps	5919
Monitoring CoS Value Aliases	5920
Operational Commands	5921
show class-of-service	5923
show class-of-service classifier	5927
show class-of-service code-point-aliases	5929
show class-of-service congestion-notification	5931
show class-of-service drop-profile	5933
show class-of-service forwarding-class	5936
show class-of-service forwarding-class-set	5938
show class-of-service forwarding-table	5940
show class-of-service forwarding-table classifier	5944
show class-of-service forwarding-table classifier mapping	5946

	show class-of-service forwarding-table drop-profile	5948
	show class-of-service forwarding-table rewrite-rule	5950
	show class-of-service forwarding-table rewrite-rule mapping	5952
	show class-of-service forwarding-table scheduler-map	5953
	show class-of-service interface	5955
	show class-of-service multi-destination	5981
	show class-of-service rewrite-rule	5982
	show class-of-service scheduler-map	5984
	show class-of-service shared-buffer	5986
	show class-of-service traffic-control-profile	5988
	show dcbx	5992
	show dcbx neighbors	5993
	show interfaces queue	6015
	show pfe next-hop	6051
	show pfe route	6056
	show pfe terse	6062
	show pfe version	6064
Chapter 66	Troubleshooting	6065
	Troubleshooting Procedures	6065
	Troubleshooting Dropped FCoE Traffic	6065
	Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth	6068
	Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth	6069
	Troubleshooting Egress Queue Bandwidth Impacted by Congestion	6070
	Troubleshooting an Unexpected Rewrite Value	6071
	Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic	6072
Part 21	Network Management and Monitoring	
Chapter 67	Overview	6077
	Network Management	6077
	Understanding Device and Network Management Features	6077
	Understanding Network Management Implementation on the QFabric System	6080
	Understanding Telnet on the QFabric System	6081
	Understanding Tracing and Logging Operations	6081
	Automation Scripts	6083
	Understanding Automation Scripts Support	6083
	How Commit Scripts Work	6084
	Commit Script Input	6085
	Commit Script Output	6086
	Commit Scripts and the Junos OS Commit Model	6087
	Avoiding Potential Conflicts When Using Multiple Commit Scripts	6089
	Overview of Generating Persistent or Transient Configuration Changes . .	6090
	Differences Between Persistent and Transient Changes	6091
	Interaction of Configuration Changes and Configuration Groups	6094

Tag Elements and Templates for Generating Changes	6094
Required Boilerplate for Commit Scripts	6095
How Op Scripts Work	6096
Required Boilerplate for Op Scripts	6097
Fabric OAM	6099
Overview of Internal Fabric Monitoring	6099
Junos Space	6101
Understanding Junos Space Support	6101
Preparing the Device for Junos Space Management	6104
sFlow Technology	6104
Overview of sFlow Technology	6105
SNMP	6107
Understanding the Implementation of SNMP	6107
Understanding the Implementation of SNMP on the QFabric System	6110
Fabric Chassis MIB	6112
Utility MIB	6116
SNMPv3 Overview	6117
Minimum SNMPv3 Configuration on a Device Running Junos OS	6118
Understanding RMON	6119
RMON Overview	6119
Alarm Thresholds and Events	6120
RMON MIB Event, Alarm, Log, and History Control Tables	6121
Understanding Health Monitoring	6123
SNMP MIBs Support	6124
MIBs Supported on QFX3500 and QFX3600 Switches	6124
MIBs Supported on QFabric Systems	6133
SNMP Traps Support	6139
SNMP Traps Supported on QFX3500 and QFX3600 Switches	6140
SNMP Traps Supported on QFabric Systems	6148
MIB Objects for the QFX Series	6152
QFX3500 and QFX3600 Switches	6152
QFabric Systems	6152
QFabric System QFX3100 Director Device	6153
QFabric System QFX3008-I Interconnect Device	6153
QFabric System QFX3600-I Interconnect Device	6153
QFabric System Node Devices	6154
System Logging	6154
Overview of Junos OS System Log Messages	6154
Overview of Single-Chassis System Logging Configuration	6155
Understanding the Implementation of System Log Messages on the QFabric System	6156
Chapter 68 Configuration	6159
Configuration Examples	6159
Examples: Configuring System Logging	6159
Examples: Assigning an Alternative Facility	6161
Example: Configuring System Log Messages	6162
Example: Monitoring Network Traffic Using sFlow Technology	6165
Example: Configuring SNMP	6169

Example: Configuring Internal Fabric OAM Monitoring	6171
Configuration Tasks for Network Management	6177
Configuring Console and Auxiliary Port Properties	6177
Configuring SSH Service for Remote Access to the Router or Switch	6178
Configuring the Root Login Through SSH	6179
Configuring the SSH Protocol Version	6179
Configuring the Client Alive Mechanism	6180
Configuring Telnet Service for Remote Access to a Switch	6180
Configuration Tasks for Automation Scripts	6181
Controlling the Execution of Commit Scripts	6181
Enabling Commit Scripts to Execute	6181
Removing Commit Scripts from the Configuration	6182
Deactivating Commit Scripts	6183
Activating Inactive Commit Scripts	6183
Configuration Tasks for Fabric OAM	6183
Configuring a Fabric Maintenance Association	6184
Configuring Flow Specifications	6185
Configuring a Unicast Ethernet Flow Specification	6185
Configuring a Unicast Ethernet IPv4 Flow Specification	6186
Configuring a Multicast IPv4 Flow Specification	6187
Configuring a Multicast VLAN Flood Flow Specification	6188
Configuration Tasks for sFlow Technology	6188
Configuring sFlow Technology	6189
Configuration Tasks for SNMP	6190
Configuring SNMP	6190
Configuring the SNMP Community String	6194
Configuring SNMP Trap Groups	6195
Adding a Group of Clients to an SNMP Community	6196
Configuring the Interfaces on Which SNMP Requests Can Be Accepted	6197
Configuring MIB Views	6197
Configuring RMON Alarms and Events	6198
Configuring SNMP	6199
Configuring an Event	6199
Configuring an Alarm	6200
Configuring Health Monitoring	6200
Creating SNMPv3 Users	6201
Configuring Access Privileges for a Group	6202
Assigning a Security Name to a Group	6204
Configuring SNMPv3 Traps on a Device Running Junos OS	6204
Configuring SNMP Informs	6206
Configuration Tasks for System Log Messages	6207
Junos OS Minimum System Logging Configuration	6207
Junos OS System Log Configuration Statements	6208
Adding a Text String to System Log Messages	6209
Directing System Log Messages to a Log File	6209
Directing System Log Messages to a Remote Machine	6210
Directing System Log Messages to a User Terminal	6211
Directing System Log Messages to the Console	6211
Disabling the System Logging of a Facility	6212

Displaying a Log File from a Single-Chassis System	6212
Including Priority Information in System Log Messages	6213
Including the Year or Millisecond in Timestamps	6215
Logging Messages in Structured-Data Format	6215
Interpreting Messages Generated in Structured-Data Format	6216
Interpreting Messages Generated in Standard Format	6219
Specifying Log File Size, Number, and Archiving Properties	6220
Specifying the Facility and Severity of Messages to Include in the Log	6222
Junos OS System Logging Facilities and Message Severity Levels	6222
System Log Default Facilities for Messages Directed to a Remote Destination	6224
Junos OS System Log Alternate Facilities for Remote Logging	6224
Changing the Alternative Facility Name for Remote System Log Messages	6225
Using Regular Expressions to Refine the Set of Logged Messages	6227
Configuration Statements for Network Management	6229
connection-limit	6230
destination-override	6231
no-remote-trace	6231
protocol-version	6232
rate-limit	6233
ssh	6234
telnet	6235
tracing	6236
Configuration Statements for Automation Scripts	6236
allow-transients	6237
apply-macro	6238
checksum	6239
command	6240
commit	6241
description	6242
direct-access	6242
file (Commit Scripts)	6243
file (Op Scripts)	6244
no-allow-url	6245
op	6246
optional	6247
refresh (Commit Scripts)	6248
refresh (Op Scripts)	6249
refresh-from (Commit Scripts)	6250
refresh-from (Op Scripts)	6251
scripts	6252
source (Commit Scripts)	6253
source (Op Scripts)	6254
Configuration Statements for Fabric OAM	6254
ethernet-frame-size	6255
fabric (OAM)	6256
fabric-maintenance-associations	6258
fabric-maintenance-end-points	6259

flow-specs	6260
unicast-ethernet-ipv4	6262
multicast-ipv4	6264
multicast-vlan-flood	6265
unicast-ethernet	6266
Configuration Statements for sFlow Technology	6267
agent-id	6267
collector	6268
interfaces (sFlow)	6268
polling-interval	6269
sample-rate	6270
sflow	6271
source-ip	6272
traceoptions (sFlow Technology)	6273
udp-port	6274
Configuration Statements for SNMP	6274
access (SNMP)	6278
address (SNMP)	6278
address-mask	6279
agent-address	6279
alarm (SNMP RMON)	6280
authentication-md5	6281
authentication-none	6282
authentication-password	6283
authentication-sha	6284
authorization	6285
bucket-size	6286
categories	6286
client-list	6287
client-list-name	6287
clients	6288
commit-delay	6288
community (SNMP)	6289
community (RMON)	6290
community-name (SNMP)	6291
contact	6292
description (SNMP)	6292
description (RMON)	6293
destination-port (SNMP)	6293
engine-id	6294
event	6295
falling-event-index (RMON)	6296
falling-threshold (Health Monitor)	6297
falling-threshold (RMON)	6298
falling-threshold-interval	6299
filter-duplicates	6299
filter-interfaces	6300
group (Associating a Security Name)	6300
group (Configuring Access Privileges)	6301

health-monitor	6302
history	6303
interface (SNMP)	6304
interface (RMON)	6305
interval (Health Monitor)	6305
interval (RMON)	6306
local-engine	6307
location	6308
logical-system (SNMP)	6309
message-processing-model	6310
name	6310
nonvolatile	6311
notify	6311
notify-filter (Applying to the Management Target)	6312
notify-filter (Configuring the Profile Name)	6312
notify-view	6313
oid	6313
oid (SNMPv3)	6314
owner	6314
parameters	6315
port (SNMP)	6315
privacy-3des	6316
privacy-aes128	6317
privacy-des	6318
privacy-none	6318
privacy-password	6319
read-view	6320
remote-engine	6321
request-type	6322
retry-count (SNMPv3)	6323
rising-event-index	6324
rising-threshold (Health Monitor)	6325
rising-threshold (RMON)	6326
rmon	6327
routing-instance (SNMP)	6328
sample-type	6329
security-level (Defining Access Privileges)	6330
security-level (Generating SNMP Notifications)	6331
security-model (Access Privileges)	6332
security-model (Group)	6333
security-model (SNMP Notifications)	6334
security-name (Community String)	6335
security-name (Security Group)	6336
security-name (SNMP Notifications)	6337
security-to-group	6338
snmp	6339
snmp-community	6343
source-address (SNMP)	6343
startup-alarm	6344

syslog-subtag	6345
tag (Configuring Notification Targets)	6345
tag (Configuring the SNMP Community)	6346
tag-list	6346
target-address	6347
target-parameters	6348
targets	6349
timeout	6349
traceoptions (SNMP)	6350
trap-group	6352
trap-options	6353
type (RMON Notification)	6354
type (SNMPv3)	6355
user	6355
usm	6356
v3	6358
vacm	6360
variable	6361
version	6362
view (Configuring a MIB View)	6363
view (Associating MIB View with a Community)	6364
write-view	6364
Configuration Statements for System Log Messages	6364
archive (All System Log Files)	6366
archive (Individual System Log File)	6368
archive (QFabric System)	6369
console (System Logging)	6370
explicit-priority	6371
facility-override	6371
file (QFabric System)	6372
file (System Logging)	6373
files	6374
host (System)	6375
log-prefix (System)	6377
match	6377
size (System)	6378
structured-data	6379
syslog (System)	6380
syslog (QFabric System)	6382
time-format	6383
user (System Logging)	6384

Chapter 69	Administration	6385
	Routine Monitoring Using the CLI	6385
	Displaying a Log File from a Single-Chassis System	6385
	Monitoring Traffic Through the Router or Switch	6386
	Displaying Real-Time Statistics About All Interfaces on the Router or Switch	6386
	Displaying Real-Time Statistics About an Interface on the Router or Switch	6387
	Monitoring RMON MIB Tables	6389
	Monitoring SNMP	6389
	Monitoring System Log Messages	6391
	Pinging Hosts	6392
	Tracing SNMP Activity on a Device Running Junos OS	6393
	Configuring the Number and Size of SNMP Log Files	6394
	Configuring Access to the Log File	6394
	Configuring a Regular Expression for Lines to Be Logged	6394
	Configuring the Trace Operations	6395
	Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage	6396
	Displaying Commit Script Output	6398
	Monitoring Commands	6400
	clear sflow collector statistics	6401
	clear snmp history	6402
	clear snmp statistics	6403
	monitor traffic	6405
	ping	6415
	ping fabric multicast-flow	6419
	ping fabric unicast-flow	6421
	request snmp spoof-trap	6424
	request snmp utility-mib clear instance	6430
	request snmp utility-mib set instance	6431
	show oam fabric flow-specification	6432
	show oam fabric interfaces	6435
	show log	6437
	show sflow	6440
	show sflow collector	6442
	show sflow interface	6443
	show snmp health-monitor	6445
	show snmp inform-statistics	6450
	show snmp mib	6452
	show snmp rmon	6455
	show snmp rmon history	6459
	show snmp statistics	6460
	show snmp v3	6464
	traceroute fabric unicast-flow	6467

Chapter 70	Troubleshooting	6471
	Troubleshooting Overview	6471
	Understanding Troubleshooting Resources	6471
	Troubleshooting Overview	6473
	Troubleshooting Procedures	6475
	Recovering from a Failed Software Installation	6475
	Loading a Previous Configuration File	6476
	Reverting to the Default Factory Configuration	6477
	Reverting to the Rescue Configuration	6478
	Recovering the Root Password	6478
Part 22	Troubleshooting	
Chapter 71	Overview	6483
	General Troubleshooting	6483
	Understanding Troubleshooting Resources	6483
	Troubleshooting Overview	6485
	Alarms	6487
	Understanding Alarms	6487
	Chassis Alarm Messages on a QFX3500 Device	6488
	Interface Alarm Messages	6491
Chapter 72	Administration	6493
	Routine Monitoring Using the CLI	6493
	Monitoring SNMP	6493
	Tracing SNMP Activity on a Device Running Junos OS	6495
	Configuring the Number and Size of SNMP Log Files	6496
	Configuring Access to the Log File	6496
	Configuring a Regular Expression for Lines to Be Logged	6496
	Configuring the Trace Operations	6496
	Monitoring RMON MIB Tables	6498
	Displaying a Log File from a Single-Chassis System	6499
	Monitoring System Log Messages	6500
	Monitoring Traffic Through the Router or Switch	6501
	Displaying Real-Time Statistics About All Interfaces on the Router or Switch	6501
	Displaying Real-Time Statistics About an Interface on the Router or Switch	6501
	Pinging Hosts	6503
Chapter 73	Troubleshooting	6505
	Configuration and File Management	6505
	Loading a Previous Configuration File	6505
	Reverting to the Default Factory Configuration	6506

Reverting to the Rescue Configuration	6507
Cleaning Up the System File Storage Space	6507
Ethernet Switching	6508
Troubleshooting Ethernet Switching	6508
Troubleshooting Layer 2 Protocol Tunneling	6509
Drop Threshold Statistics Might Be Incorrect	6509
Egress Filtering of L2PT Traffic Not Supported	6509
Troubleshooting Private VLANs	6510
Limitations of Private VLANs	6510
Forwarding with Private VLANs	6510
Egress Firewall Filters with Private VLANs	6511
Egress Port Mirroring with Private VLANs	6512
Troubleshooting Q-in-Q and VLAN Translation Configuration	6513
Firewall Filter Match Condition Not Working with Q-in-Q Tunneling ..	6513
Egress Port Mirroring with VLAN Translation	6513
Hardware	6513
Troubleshooting QFX3100 Director Device Isolation	6514
High Availability	6515
Troubleshooting VRRP	6515
Interfaces	6516
Troubleshooting an Aggregated Ethernet Interface	6516
Troubleshooting Network Interfaces	6516
The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down	6516
Troubleshooting Multichassis Link Aggregation	6517
MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table	6517
MC-LAG Peer Does Not Go into Standby Mode	6518
Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive	6518
Redirect Filters Take Priority over User-Defined Filters	6518
Operational Command Output Is Wrong	6519
ICCP Connection Might Take Up to 60 Seconds to Become Active ..	6519
MAC Address Age Learned on an MC-AE Interface Is Reset to Zero ..	6519
MAC Address Is Not Learned Remotely in a Default VLAN	6520
Snooping Entries Learned on MC-AE Interfaces Are Not Removed ..	6520
ICCP Does Not Come Up After You Add or Delete an Authentication Key	6520
Local Status Is Standby When It Should Be Active	6520
Packets Loop on the Server When ICCP Fails	6520
Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change	6521
No Commit Checks Are Done for ICL-PL Interfaces	6521
Double Failover Scenario	6521
Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up	6521
Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer	6522
AE Interfaces Go Down	6522

Flooding of Upstream Traffic	6522
Junos OS Basics	6522
Rebooting and Halting a QFX Series Product	6523
Recovering from a Failed Software Installation	6524
Recovering the Root Password	6524
Creating an Emergency Boot Device for a QFX Series Device	6526
Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device	6528
Performing a QFabric System Recovery Installation on the Director Group	6529
(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive	6530
Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software	6532
Layer 3 Protocols	6536
Troubleshooting Virtual Routing Instances	6536
Direct Routes Not Leaked Between Routing Instances	6536
Security	6537
Troubleshooting Firewall Filter Configuration	6537
Firewall Filter Configuration Returns a No Space Available in TCAM Message	6537
Filter Counts Previously Dropped Packet	6539
Matching Packets Not Counted	6539
Counter Reset When Editing Filter	6540
Cannot Include loss-priority and policer Actions in Same Term	6540
Cannot Egress Filter Certain Traffic Originating on QFX Switch	6540
Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	6541
Egress Firewall Filters with Private VLANs	6541
Egress Filtering of L2PT Traffic Not Supported	6542
Cannot Drop BGP Packets in Certain Circumstances	6542
Invalid Statistics for Policer	6542
Policers can Limit Egress Filters	6542
Troubleshooting Policer Configuration	6543
Incomplete Count of Packet Drops	6544
Counter Reset When Editing Filter	6544
Invalid Statistics for Policer	6544
Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	6544
Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	6545
Policers Can Limit Egress Filters	6546
Services	6547
Troubleshooting Port Mirroring	6547
Port Mirroring Constraints and Limitations	6547
Egress Port Mirroring with VLAN Translation	6549
Egress Port Mirroring with Private VLANs	6549
Storage	6550
Troubleshooting Dropped FCoE Traffic	6550
Troubleshooting Fibre Channel Interface Deletion	6553

Troubleshooting Dropped FIP Traffic	6554
Traffic Management	6555
Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth	6555
Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth	6556
Troubleshooting Egress Queue Bandwidth Impacted by Congestion	6557
Troubleshooting an Unexpected Rewrite Value	6558
Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic	6559

Part 23

Index

Index	6565
-----------------	------

List of Figures

Part 1	QFX3500 Switch Overview	
Chapter 1	QFX3500 Switch Overview	3
	Figure 1: QFX3500 Device Front	4
	Figure 2: QFX3500 Device Rear	4
	Figure 3: SFP+ Access Port Locations	5
	Figure 4: QSFP+ Uplink Port Locations	6
Part 2	QFX3600 Switch Overview	
Chapter 2	QFX3600 Switch Overview	9
	Figure 5: QFX3600 Chassis Front	10
	Figure 6: QFX3600 Chassis Rear	10
Part 4	Junos OS Basics	
Chapter 4	Overview	33
	Figure 7: DHCP Client/Server Model	64
	Figure 8: DHCP Four-Step Transfer	67
	Figure 9: Commands That Combine Other Commands	75
	Figure 10: CLI Command Hierarchy	78
	Figure 11: Command Output Options	80
	Figure 12: Configuration Mode Hierarchy of Statements	84
Part 5	QFabric System Deployment	
Chapter 9	Overview	1165
	Figure 13: Legacy Data Center Architecture	1166
	Figure 14: QFX Series QFabric System Architecture	1167
	Figure 15: QFabric System Hardware Architecture	1178
	Figure 16: External Routing Engine Types	1182
	Figure 17: Clos Switching for QFX3008-I Interconnect Devices	1184
	Figure 18: QFX3008-I Data Plane and Control Plane Connections	1185
	Figure 19: QFX3600-I Data Plane and Control Plane Connections	1186
	Figure 20: QFX3500 Data Plane and Control Plane Connections	1188
	Figure 21: QFX3600 Data Plane and Control Plane Connections	1189
	Figure 22: QFabric System Topology - Default Partition	1195
	Figure 23: QFabric System Control Plane Network	1197
	Figure 24: QFabric System Data Plane Network	1200
	Figure 25: QFX3008-I Interconnect Device Cross-Connect System	1201
Chapter 10	Configuration	1255

	Figure 26: QFX3000-G QFabric System Control Plane—Virtual Chassis Port Ranges	1265
	Figure 27: QFX3000-G QFabric System Control Plane—Director Group to Virtual Chassis Connections	1270
	Figure 28: QFX3000-G QFabric System Control Plane—Interconnect Device to Virtual Chassis Connections	1272
	Figure 29: QFX3000-G QFabric System Control Plane—Node Device to Virtual Chassis Connections	1273
	Figure 30: QFX3000-G QFabric System Control Plane—Inter-Virtual Chassis LAG Connections	1274
	Figure 31: QFX3000-M QFabric System Control Plane—EX4200 Switch Port Ranges	1311
	Figure 32: QFX3000-M QFabric System Control Plane—Director Group to EX4200 Switch Connections	1314
	Figure 33: QFX3000-M QFabric System Control Plane—Interconnect Device to EX4200 Switch Connections	1315
	Figure 34: QFX3000-M QFabric System Control Plane—Node Device to EX4200 Switch Connections	1316
	Figure 35: QFX3000-M QFabric System Control Plane—Inter-EX4200 Switch LAG Connections	1317
Part 6	Configuration and File Management	
Chapter 14	Configuration	1593
	Figure 36: Overriding the Current Configuration	1625
	Figure 37: Using the replace Option	1626
	Figure 38: Using the merge Option	1626
	Figure 39: Using a Patch File	1627
	Figure 40: Using the set Option	1627
Part 8	Ethernet Features	
Chapter 20	Overview	1867
	Figure 41: Subdomains in a PVLAN	1879
	Figure 42: PVLAN Spanning Multiple Switches	1880
	Figure 43: Community VLAN Sends Untagged Traffic	1884
	Figure 44: Isolated VLAN Sends Untagged Traffic	1885
	Figure 45: PVLAN Tagged Traffic Sent on a Promiscuous Port	1886
	Figure 46: Two Secondary VLAN Trunk Ports on One Interface	1890
	Figure 47: Secondary VLAN Trunk and Promiscuous Trunk on One Interface	1892
	Figure 48: Secondary VLAN Trunk and PVLAN Trunk on One Interface	1893
	Figure 49: Secondary VLAN Trunk and Non-Private VLAN Port on One Interface	1894
	Figure 50: Traffic Ingressing on Promiscuous Access Port	1895
	Figure 51: L2PT Example	1912
Chapter 21	Configuration	1917
	Figure 52: Network Topology for RSTP	1919
	Figure 53: Network Topology for MSTP	1932
	Figure 54: Reflective Relay Topology	1959

	Figure 55: RVI with One Switch	1977
	Figure 56: BPDU Protection Topology	1991
	Figure 57: Network Topology for Root Protection	1996
	Figure 58: Network Topology for Loop Protection	2000
	Figure 59: PVLAN Topology Spanning Multiple Switches	2012
	Figure 60: PVLAN Topology with Secondary VLAN Trunk Ports and Promiscuous Access Port	2026
Part 9	High Availability	
Chapter 24	Overview	2269
	Figure 61: Basic VRRP Topology	2276
Chapter 25	Configuration	2281
	Figure 62: VRRP Load-Sharing Configuration	2292
Part 10	Interfaces	
Chapter 28	Overview	2415
	Figure 63: Uplink Failure Detection Configuration on Switches	2446
Chapter 29	Configuration	2449
	Figure 64: Uplink Failure Detection Configuration on Switches	2450
	Figure 65: Configuring a Multichassis LAG Between Switch A and Switch B . .	2463
	Figure 66: Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP . .	2478
	Figure 67: Configuring a Multichassis LAG Between Switch A and Switch B . .	2501
	Figure 68: Configuring a Multichassis LAG Between Switch A and Switch B . .	2516
Part 11	Routing Options	
Chapter 33	Configuration	2825
	Figure 69: Customer Routes Connected to a Service Provider	2836
	Figure 70: BFD Enabled on Qualified Next Hops	2841
	Figure 71: Customer Routes Connected to a Service Provider	2849
Part 12	Border Gateway Protocol	
Chapter 36	Overview	3143
	Figure 72: ASs, EBGp, and IBGP	3145
Chapter 37	Configuration	3149
	Figure 73: BGP Peering Session	3150
	Figure 74: Typical Network with BGP Peer Sessions	3151
	Figure 75: Typical Network with BGP Peer Sessions	3158
	Figure 76: Internal and External BGP	3172
	Figure 77: Typical Network with IBGP Sessions	3175
	Figure 78: Typical Network with IBGP Sessions	3185
	Figure 79: Typical Network with IBGP Sessions and Multiple Exit Points	3196
	Figure 80: Default MED Example	3209
	Figure 81: Typical Network with IBGP Sessions and Multiple Exit Points	3211
	Figure 82: Typical Network with IBGP Sessions and Multiple Exit Points	3224

	Figure 83: Topology for Delaying the MED Update	3238
	Figure 84: Local AS Configuration	3249
	Figure 85: Topology for Configuring the Local AS	3252
	Figure 86: Topology for Configuring a Private Local AS	3262
	Figure 87: Advertisement of Multiple Paths in BGP	3269
	Figure 88: BGP Prefix-Based Outbound Route Filtering	3314
	Figure 89: Typical Network with EBGP Multihop Sessions	3318
	Figure 90: BGP Preference Value Topology	3329
	Figure 91: Topology for Ignoring the AS-Path Length	3337
	Figure 92: Topology for Removing a Private AS from the Advertised AS Path	3344
	Figure 93: Typical Network with IBGP Sessions	3350
	Figure 94: BGP Load Balancing	3364
	Figure 95: Topology for Accepting a Remote Next Hop	3369
	Figure 96: Advertisement of Multiple Paths in BGP	3381
	Figure 97: Simple Route Reflector Topology (One Cluster)	3406
	Figure 98: Basic Route Reflection (Multiple Clusters)	3406
	Figure 99: Hierarchical Route Reflection (Clusters of Clusters)	3407
	Figure 100: IBGP Network Using a Route Reflector	3409
	Figure 101: BGP Confederations	3423
	Figure 102: Typical Network Using BGP Confederations	3424
	Figure 103: Authentication for BGP	3431
	Figure 104: Typical Network with BGP Peer Sessions	3436
	Figure 105: TCP Maximum Segment Size for BGP	3445
	Figure 106: Topology for the EBGP Case	3452
	Figure 107: Topology for the RR Case	3452
	Figure 108: BGP Flap Damping Topology	3458
	Figure 109: MBGP MVPN with BGP Route Flap Damping	3466
Part 13	Intermediate System to Intermediate System	
Chapter 39	Overview	3625
	Figure 110: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2	3629
Chapter 40	Configuration	3633
	Figure 111: Simple IS-IS Topology	3634
	Figure 112: IS-IS Multi-Level Topology	3639
	Figure 113: Hitless Authentication Key Rollover for IS-IS	3648
	Figure 114: IS-IS Route Redistribution Topology	3652
	Figure 115: Configuring BFD for IS-IS	3660
	Figure 116: IS-IS BFD Authentication Topology	3665
	Figure 117: Configuring IS-IS Multicast Topology	3671
	Figure 118: IS-IS Checksum Topology	3688
Part 14	Open Shortest Path First	
Chapter 42	Overview	3797
	Figure 119: OSPF Three-Way Handshake	3801
Chapter 43	Configuration	3809

	Figure 120: Multiarea OSPF Topology	3814
	Figure 121: Typical Single-Area OSPF Network Topology	3816
	Figure 122: Typical Multiarea OSPF Network Topology	3818
	Figure 123: OSPF AS Network with Stub Areas and NSSAs	3821
	Figure 124: OSPF Network Topology with Stub Areas and NSSAs	3824
	Figure 125: OSPF Network Topology with Stub Areas and NSSAs	3828
	Figure 126: Summarizing Ranges of Routes in OSPF	3853
	Figure 127: OSPF Metric Configuration	3864
	Figure 128: Sample Topology Used for an OSPF Export Network Summary Policy	3939
	Figure 129: Sample Topology Used for an OSPF Import Network Summary Policy	3947
Part 15	Routing Information Protocol	
Chapter 45	Overview	4103
	Figure 130: Distance-Vector Protocol	4104
	Figure 131: Split Horizon Example	4106
	Figure 132: Poison Reverse Example	4107
	Figure 133: Limitations of Unidirectional Connectivity	4107
Chapter 46	Configuration	4109
	Figure 134: Sample RIP Network Topology	4110
	Figure 135: RIP Authentication Network Topology	4117
	Figure 136: RIP BFD Network Topology	4125
	Figure 137: RIP BFD Authentication Network Topology	4131
	Figure 138: RIP Import Policy Network Topology	4137
	Figure 139: Controlling Traffic in a RIP Network with the Incoming Metric	4143
	Figure 140: Controlling Traffic in a RIP Network with the Outgoing Metric	4145
	Figure 141: RIP Incoming Metrics Network Topology	4146
	Figure 142: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology	4151
	Figure 143: Redistributing Routes Between RIP Instances Network Topology . .	4155
	Figure 144: RIP Timers Network Topology	4161
	Figure 145: RIP Trace Operations Network Topology	4168
Part 16	Multicast	
Chapter 48	Overview	4209
	Figure 146: Rendezvous Point as Part of the RPT and SPT	4215
	Figure 147: Building an RPT Between the RP and the Receiver	4222
	Figure 148: PIM Register Message and PIM Join Message Exchanged	4223
	Figure 149: Traffic Sent from the Source to the RP Router	4224
	Figure 150: Traffic Sent from the RP Router Toward the Receiver	4224
	Figure 151: Receiver DR Sends a PIM Join Message to the Source	4226
	Figure 152: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router	4227
	Figure 153: RP Router Receives PIM Prune Message	4227
	Figure 154: RP Router Sends a PIM Prune Message to the Source DR	4228

	Figure 155: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router	4228
	Figure 156: Receiver Announces Desire to Join Group G and Source S	4240
	Figure 157: Router 3 (Last-Hop Router) Joins the Source Tree	4240
	Figure 158: (S,G) State Is Built Between the Source and the Receiver	4241
Chapter 49	Configuration	4245
	Figure 159: Join Suppression	4261
	Figure 160: PIM Assert Topology	4286
	Figure 161: MVR Topology in Transparent Mode	4324
	Figure 162: MVR Topology in Proxy Mode	4325
	Figure 163: Source-Active Message Flooding	4335
	Figure 164: Network on Which to Configure PIM SSM	4343
	Figure 165: Receiver Sends Messages to Join Group G and Source S	4348
	Figure 166: Router 3 (Last-Hop Router) Joins the Source Tree	4348
	Figure 167: (S,G) State Is Built Between the Source and the Receiver	4348
	Figure 168: Simple RPF Topology	4348
Part 17	Security	
Chapter 51	Overview	4635
	Figure 169: Evaluation of Terms Within a Firewall Filter	4639
	Figure 170: Application of Firewall Filters to Control Packet Flow	4640
	Figure 171: Flow of Tricolor Marking Policer Operation	4664
	Figure 172: DHCP Server Connected Directly to Switch	4681
	Figure 173: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port	4682
	Figure 174: Switch Is the DHCP Server	4683
	Figure 175: Switch Acting as Relay Agent Through Router to DHCP Server	4684
	Figure 176: Switch Relays DHCP Requests to Server	4693
Chapter 52	Configuration	4697
	Figure 177: Network Topology for Basic Port Security	4708
	Figure 178: Network Topology for Basic Port Security	4716
	Figure 179: Network Topology for Basic Port Security	4719
	Figure 180: Network Topology for Basic Port Security	4723
	Figure 181: Network Topology for Port Security Setup with Two Switches on the Same VLAN	4726
	Figure 182: Network Topology for Basic Port Security	4734
	Figure 183: Network Topology for Basic Port Security	4738
	Figure 184: Network Topology for Configuring DHCP Option 82 on a Switch That Is on the Same VLAN as the DHCP Clients and the DHCP Server	4742
Part 18	Services	
Chapter 56	Configuration	4895
	Figure 185: Network Topology for Local Port Mirroring Example	4896
Part 19	Storage	
Chapter 59	Overview	4955

	Figure 186: ENode Components	4969
	Figure 187: FCoE Transit Switch Connecting FCoE Devices to an FC Switch . . .	4973
	Figure 188: FCoE-FC Gateway Topology	4975
	Figure 189: Traffic Switching Between FCoE Hosts Connected to the FC Network by an FCoE-FC Gateway	4978
	Figure 190: FCoE-FC Gateway Fabric Login and FIP Login	4989
	Figure 191: Sample Load-Balancing Topology	5019
	Figure 192: FCoE Transit Switch Performs VN2VF_Port FIP Snooping	5026
	Figure 193: VN2VN_Port Traffic Across a QFabric Interconnect Device	5036
	Figure 194: Supported Topology for an MC-LAG on an FCoE Transit Switch . .	5048
Chapter 60	Configuration	5099
	Figure 195: Fibre Channel Interface Configuration Topology	5102
	Figure 196: PFC for FCoE Traffic Configuration Components Block Diagram . . .	5115
	Figure 197: Supported Topology for an MC-LAG on an FCoE Transit Switch . . .	5123
	Figure 198: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Same Transit Switch) Topology	5157
	Figure 199: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Different Transit Switches) Topology	5162
	Figure 200: VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected) Topology	5169
Part 20	Traffic Management	
Chapter 63	Overview	5411
	Figure 201: Packet Flow Across the Network	5413
	Figure 202: Flow of Tricolor Marking Policer Operation	5431
	Figure 203: CoS Classifier, Queues, and Scheduler	5441
	Figure 204: Packet Flow Through Configurable CoS Components	5441
	Figure 205: Hierarchical Scheduling Tiers	5482
	Figure 206: Hierarchical Scheduling Packet Flow	5484
	Figure 207: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling . .	5498
	Figure 208: Setting Maximum Bandwidth Using Hierarchical Scheduling . . .	5501
	Figure 209: Tail-Drop Profile Packet Drop	5546
Chapter 64	Configuration	5605
	Figure 210: Hierarchical Port Scheduling Components Block Diagram	5610
	Figure 211: Hierarchical Port Scheduling Packet Flow Block Diagram	5610
	Figure 212: PFC for FCoE Traffic Configuration Components Block Diagram . .	5630
	Figure 213: Supported Topology for an MC-LAG on an FCoE Transit Switch . . .	5637
	Figure 214: Tail-Drop Profile Packet Drop Example	5665
	Figure 215: Topology of the Two Lossless FCoE Priorities Example	5712
	Figure 216: Topology of the Lossless FCoE and iSCSI Priorities Example	5725
	Figure 217: Topology of the IEEE 802.1p Priority Remapping Example	5741
Part 21	Network Management and Monitoring	
Chapter 67	Overview	6077
	Figure 218: Commit Script Input and Output	6085
	Figure 219: Standard Commit Model	6087

	Figure 220: Commit Model with Commit Scripts Added	6088
	Figure 221: Configuration Evaluation by Multiple Commit Scripts	6090
	Figure 222: Op Script Input and Output	6097
	Figure 223: SNMP Communication Flow	6109
	Figure 224: Setting Thresholds	6120
Chapter 68	Configuration	6159
	Figure 225: sFlow Technology Monitoring System	6166
	Figure 226: Inform Request and Response	6206

List of Tables

	About the Documentation	cxxi
	Table 1: Notice Icons	cxiii
	Table 2: Text and Syntax Conventions	cxiii
Part 3	Software Feature Support	
Chapter 3	Software Feature Support for the QFX Series	15
	Table 3: Access Control Features	16
	Table 4: Administration Features	16
	Table 5: CoS Features	16
	Table 6: High Availability and Resiliency Features	17
	Table 7: Interface Features	19
	Table 8: IP Address Management Features	19
	Table 9: Layer 2 Network Protocol Features	20
	Table 10: Layer 3 Protocol Features	22
	Table 11: Multicast Protocol Features	23
	Table 12: Network Management and Monitoring Features	24
	Table 13: Port Security Features	26
	Table 14: QFabric System-Specific Features	26
	Table 15: Security	27
	Table 16: Storage and Fibre Channel Features	28
	Table 17: System Management Features	29
Part 4	Junos OS Basics	
Chapter 4	Overview	33
	Table 18: Access Control Features	34
	Table 19: Administration Features	35
	Table 20: CoS Features	35
	Table 21: High Availability and Resiliency Features	36
	Table 22: Interface Features	37
	Table 23: IP Address Management Features	37
	Table 24: Layer 2 Network Protocol Features	38
	Table 25: Layer 3 Protocol Features	40
	Table 26: Multicast Protocol Features	42
	Table 27: Network Management and Monitoring Features	43
	Table 28: Port Security Features	44
	Table 29: QFabric System-Specific Features	45
	Table 30: Security	45
	Table 31: Storage and Fibre Channel Features	46
	Table 32: System Management Features	48

	Table 33: Configuration File Terms	49
	Table 34: Legacy DHCP and Extended DHCP Server Hierarchy Levels	65
	Table 35: Junos OS Processes	68
	Table 36: Commonly Used Operational Mode Commands	77
	Table 37: Summary of Configuration Mode Commands	81
	Table 38: Configuration Mode Top-Level Statements	83
	Table 39: Junos OS Feature Licenses and Model Numbers for QFX Series Devices	87
	Table 40: Upgrade Licenses for Enhancing Port Capacity	89
Chapter 5	Installation	99
	Table 41: Junos OS and Jloader Software Compatibility Matrix for the QFX3500 Switch and QFX3500 Node Device	121
	Table 42: Junos OS and Jloader Software Compatibility Matrix for the QFX3008-I Interconnect Device	121
	Table 43: Junos OS and Jloader Software Compatibility Matrix for the QFX3600-I Interconnect Device and QFX3600 Node Device	122
	Table 44: Uboot Software Release and Jloader Software Compatibility Matrix	122
Chapter 6	Configuration	141
	Table 45: DHCP Client Settings	143
	Table 46: Methods for Configuring Junos OS	169
Chapter 7	Administration	315
	Table 47: Summary of System Process Information Output Fields	315
	Table 48: Summary of Key System Properties Output Fields	316
	Table 49: request system storage cleanup Output Fields	425
	Table 50: show chassis alarms Output Fields	458
	Table 51: show chassis led Output Fields	464
	Table 52: show chassis ethernet-switch interconnect-device fpc Output Fields	466
	Table 53: show chassis ethernet-switch interconnect-device fpc Output Fields	491
	Table 54: show chassis environment Output Fields	514
	Table 55: show chassis environment cb Output Fields	563
	Table 56: show chassis environment fpc Output Fields	582
	Table 57: show chassis environment pem Output Fields	606
	Table 58: show chassis environment routing-engine Output Fields	615
	Table 59: show chassis fan Output Fields	620
	Table 60: show chassis firmware Output Fields	634
	Table 61: show chassis fpc Output Fields	648
	Table 62: Routing Engines Displaying DIMM Information	670
	Table 63: show chassis hardware Output Fields	673
	Table 64: show chassis lcd Output Fields	787
	Table 65: show chassis led Output Fields	799
	Table 66: show chassis location Output Fields	810
	Table 67: show chassis mac-addresses Output Fields	814
	Table 68: show chassis nonstop-upgrade node-group Output Fields	817
	Table 69: show chassis pic Output Fields	821

Table 70: show chassis routing-engine Output Fields	834
Table 71: show chassis temperature-thresholds Output Fields	854
Table 72: show chassis zones Output Fields	869
Table 73: show cli Output Fields	874
Table 74: show cli authorization Output Fields	876
Table 75: show cli directory Output Fields	880
Table 76: show cli history Output Fields	881
Table 77: show interfaces diagnostics optics Output Fields	883
Table 78: show ntp associations Output Fields	892
Table 79: show ntp status Output Fields	894
Table 80: show subscribers Output Fields	899
Table 81: show system buffers Output Fields	933
Table 82: show system certificate Output Fields	937
Table 83: show system commit Output Fields	939
Table 84: show system connections Output Fields	946
Table 85: show system core-dumps Output Fields	965
Table 86: show system directory-usage Output Fields	978
Table 87: show system license Output Fields	980
Table 88: show system processes Output Fields	990
Table 89: show system resource-cleanup processes Output Fields	1013
Table 90: show system services service-deployment Output Fields	1017
Table 91: show system storage Output Fields	1063
Table 92: show system uptime Output Fields	1069
Table 93: show system users Output Fields	1074
Table 94: show system virtual-memory Output Fields	1080
Table 95: traceroute Output Fields	1153
Table 96: traceroute monitor Output Fields	1156

Part 5

Chapter 9

QFabric System Deployment

Overview	1165
Table 97: QFabric System Terms	1169
Table 98: QFX3600 Node Device Port Mappings	1174
Table 99: Supported QFabric System Hardware Configurations	1179
Table 100: Oversubscription Ratio on Node Devices	1191
Table 101: Access Control Features	1203
Table 102: Administration Features	1203
Table 103: CoS Features	1203
Table 104: High Availability and Resiliency Features	1204
Table 105: Interface Features	1205
Table 106: IP Address Management Features	1206
Table 107: Layer 2 Network Protocol Features	1207
Table 108: Layer 3 Protocol Features	1209
Table 109: Multicast Protocol Features	1210
Table 110: Network Management and Monitoring Features	1211
Table 111: Port Security Features	1213
Table 112: QFabric System-Specific Features	1213
Table 113: Security	1214
Table 114: Storage and Fibre Channel Features	1215

	Table 115: System Management Features	1216
	Table 116: QFX3600 Node Device Port Mappings	1232
	Table 117: Default Fabric Forwarding Class Sets	1239
	Table 118: Default Forwarding Class to Fabric Forwarding Class Set Mapping	1240
	Table 119: Class Group Scheduling Properties and Membership	1245
	Table 120: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	1248
	Table 121: show class-of-service forwarding-class-set Command Output Fields	1249
	Table 122: Summary of Differences Between Fabric fc-sets and Node Device fc-sets	1250
	Table 123: Junos OS Feature Licenses and Model Numbers for QFX Series Devices	1252
	Table 124: Upgrade Licenses for Enhancing Port Capacity	1253
Chapter 10	Configuration	1255
	Table 125: Support for device mode options	1258
	Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments	1265
	Table 127: Director Group Port Mappings	1270
	Table 128: Hardware to Software Port Mappings for Director Device Network Modules	1271
	Table 129: Interconnect Device Port Mappings	1272
	Table 130: Interconnect Device Port Mappings for Two Additional Devices	1273
	Table 131: Node Device Port Mappings	1274
	Table 132: Virtual Chassis LAG Port Mappings	1275
	Table 133: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX4200 Switch Port Mappings	1311
	Table 134: Director Group Port Mappings	1314
	Table 135: Hardware to Software Port Mappings for Director Device Network Modules	1315
	Table 136: Interconnect Device Port Mappings	1315
	Table 137: Node Device to EX4200 Switch Port Mappings	1316
	Table 138: EX4200 Switch LAG Port Mappings	1317
Chapter 11	Administration	1405
	Table 139: QFabric System Operational Mode Commands	1437
	Table 140: show chassis device-mode Output Fields	1476
	Table 141: show chassis ethernet-switch interconnect-device fpc Output Fields	1478
	Table 142: show chassis ethernet-switch interconnect-device fpc Output Fields	1495
	Table 143: show chassis fabric connectivity Output Fields	1520
	Table 144: show chassis fabric device Output Fields	1527
	Table 145: show chassis lcd Output Fields	1531
	Table 146: show chassis nonstop-upgrade node-group Output Fields	1542
	Table 147: show fabric administration inventory Output Fields	1544
	Table 148: show fabric administration inventory director-group status Output Fields	1548

	Table 149: show fabric administration inventory infrastructure Output Fields . .	1553
	Table 150: show fabric administration inventory interconnect-devices Output Fields	1556
	Table 151: show fabric administration inventory node-devices Output Fields . .	1558
	Table 152: show fabric administration inventory node-groups Output Fields . .	1560
	Table 153: show fabric administration system mac-pool Output Fields	1562
	Table 154: show fabric inventory Output Fields	1564
	Table 155: show fabric session-host Output Fields	1566
	Table 156: show system software upgrade status Output Fields	1570
Part 6	Configuration and File Management	
Chapter 13	Overview	1587
	Table 157: Configuration File Terms	1587
	Table 158: Forms of the configure Command	1588
Chapter 14	Configuration	1593
	Table 159: Options for the load Command	1610
Chapter 15	Administration	1629
	Table 160: show system commit Output Fields	1649
Part 7	User and Access Management	
Chapter 17	Overview	1661
	Table 161: Junos OS Processes	1662
	Table 162: Juniper Networks Vendor-Specific RADIUS Attributes	1668
	Table 163: Juniper Networks Vendor-Specific TACACS+ Attributes	1670
	Table 164: Login Class Permission Flags	1672
	Table 165: Order of Authentication Attempts	1678
	Table 166: Predefined System Login Classes	1683
	Table 167: Configuration Mode Hierarchies—Common Regular Expression Operators	1686
	Table 168: Common Regular Expression Operators to Allow or Deny Operational Mode Commands	1686
	Table 169: Special Requirements for Plain-Text Passwords	1687
Chapter 18	Configuration	1691
	Table 170: Match Conditions	1725
	Table 171: Actions for VSAs	1726
Chapter 19	Administration	1831
	Table 172: show ethernet-switching interfaces Output Fields	1838
	Table 173: show lldp Output Fields	1843
	Table 174: show lldp local-information Output Fields	1847
	Table 175: show lldp neighbors Output Fields	1849
	Table 176: show lldp statistics Output Fields	1853
	Table 177: show route instance Output Fields	1855
	Table 178: show snmp statistics Output Fields	1859

Part 8	Ethernet Features	
Chapter 20	Overview	1867
	Table 179: Sample RVI Values	1875
	Table 180: Number of Supported RVIs by Platform	1876
	Table 181: PVLAN Requirements for 802.1Q Tags	1881
	Table 182: PVLAN Ports and Layer 2 Connectivity	1882
	Table 183: MVRP VLAN Requirements for Node Devices	1897
	Table 184: Protocol Destination MAC Addresses	1912
Chapter 21	Configuration	1917
	Table 185: Topology for Configuring RSTP on the QFX Series	1919
	Table 186: Topology for Configuring MSTP on the QFX Series	1933
	Table 187: Components of the Example Topology	1954
	Table 188: Components of the Topology for Configuring Reflective Relay	1959
	Table 189: Components of the Basic Bridging Configuration Topology	1962
	Table 190: Components of the Multiple VLAN Topology	1972
	Table 191: Components of the Multiple VLAN Topology	1977
	Table 192: Components of the Topology for Connecting an Access Switch to a Distribution Switch	1982
	Table 193: Components of the Topology for Configuring BPDU Protection on the QFX Series	1992
	Table 194: Topology for Configuring Root Protection on the QFX Series	1996
	Table 195: Topology for Configuring Loop Protection on the QFX Series	2000
	Table 196: Components of the Topology for Configuring a PVLAN	2006
	Table 197: Components of Switch 1 in the Topology for Configuring a PVLAN Spanning Multiple Devices	2013
	Table 198: Components of Switch 2 in the Topology for Configuring a PVLAN Spanning Multiple Devices	2013
	Table 199: Components of Switch 3 in the Topology for Configuring a PVLAN Spanning Multiple Devices	2014
	Table 200: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 1	2026
	Table 201: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 2	2027
	Table 202: Components of the Topology for Setting Up Q-in-Q Tunneling	2037
Chapter 22	Administration	2173
	Table 203: show ethernet-switching interfaces Output Fields	2191
	Table 204: show ethernet-switching layer2-protocol-tunneling interface Output Fields	2195
	Table 205: show ethernet-switching layer2-protocol-tunneling statistics Output Fields	2198
	Table 206: show ethernet-switching layer2-protocol-tunneling vlan Output Fields	2200
	Table 207: show ethernet-switching mac-learning-log Output Fields	2202
	Table 208: show ethernet-switching mac-notification Output Fields	2204
	Table 209: show ethernet-switching statistics aging Output Fields	2205
	Table 210: show ethernet-switching statistics mac-learning Output Fields	2208
	Table 211: show ethernet-switching table Output Fields	2211

	Table 212: show interfaces xe Output Fields	2218
	Table 213: show spanning-tree bridge Output Fields	2235
	Table 214: show spanning-tree Interface Output Fields	2240
	Table 215: show spanning-tree mstp configuration Output Fields	2246
	Table 216: show spanning-tree statistics Output Fields	2248
	Table 217: show vlans Output Fields	2252
Part 9	High Availability	
Chapter 24	Overview	2269
	Table 218: RVIs on QFabric systems in example VRRP configuration	2278
	Table 219: Sample VRRP configuration each RVI	2278
	Table 220: Interfaces on QFabric system A. All interfaces are members of VLAN 100.	2279
	Table 221: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).	2279
Chapter 25	Configuration	2281
	Table 222: Settings for VRRP Load-Sharing Example	2292
	Table 223: Interface State and Priority Cost Usage	2303
Chapter 26	Administration	2343
	Table 224: show bgp neighbor Output Fields	2347
	Table 225: show ospf overview Output Fields	2364
	Table 226: show chassis nonstop-upgrade node-group Output Fields	2399
	Table 227: show vrrp Output Fields	2400
Part 10	Interfaces	
Chapter 28	Overview	2415
	Table 228: Network Interface Types and Purposes	2415
	Table 229: Special Interface Types and Purposes	2416
	Table 230: Valid Port Ranges on QFX3500 Devices	2422
	Table 231: Valid Port Ranges on QFX3600 Standalone Switches	2426
	Table 232: Valid Port Ranges on QFX3600 Node Devices	2429
	Table 233: ICCP Failure Scenarios	2437
Chapter 29	Configuration	2449
	Table 234: Settings for Uplink Failure Protection Example	2451
	Table 235: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch	2455
	Table 236: Components of the Topology for Configuring a Multichassis LAG Between Two Switches	2463
	Table 237: Components of the Topology for Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP	2479
	Table 238: Components of the Topology for Configuring a Multichassis LAG Between Two Switches	2501
	Table 239: Components of the Topology for Configuring a Multichassis LAG Between Two Switches	2516
	Table 240: Protocol Families and Supported Interface Types	2570

Chapter 30	Administration	2627
	Table 241: Summary of System Process Information Output Fields	2627
	Table 242: Summary of Key System Properties Output Fields	2628
	Table 243: Output Control Keys for the monitor interface Command	2635
	Table 244: Output Control Keys for the monitor interface traffic Command	2636
	Table 245: monitor interface Output Fields	2637
	Table 246: show iccp	2643
	Table 247: show interfaces diagnostics optics Output Fields	2645
	Table 248: show interfaces fabric Output Fields	2652
	Table 249: show interfaces ge Output Fields	2673
	Table 250: show interfaces mc-ae Output Fields	2687
	Table 251: show interfaces statistics fabric Output Fields	2690
	Table 252: Layer 2 Overhead, Transmitted Packets/Bytes	2708
	Table 253: show interfaces queue Output Fields	2710
	Table 254: Byte Count by PIC Type	2713
	Table 255: show interfaces queue fabric Output Fields	2743
	Table 256: show interfaces xe Output Fields	2766
	Table 257: show interfaces xe Output Fields	2784
	Table 258: show lacp interfaces Output Fields	2802
	Table 259: show lacp statistics interfaces Output Fields	2806
	Table 260: show uplink-failure-detection Output Fields	2808
Part 11	Routing Options	
Chapter 34	Administration	2955
	Table 261: Filtering Route Messages	2955
	Table 262: Summary of Key Routing Information Output Fields	2956
	Table 263: show as-path Output Fields	2961
	Table 264: show as-path domain Output Fields	2964
	Table 265: show as-path summary Output Fields	2967
	Table 266: show ipv6 neighbors Output Fields	2969
	Table 267: show ipv6 router-advertisement Output Fields	2971
	Table 268: show route Output Fields	2975
	Table 269: show route damping Output Fields	2998
	Table 270: show route detail Output Fields	3002
	Table 271: Next-hop Types Output Field Values	3006
	Table 272: State Output Field Values	3008
	Table 273: Communities Output Field Values	3010
	Table 274: show route export Output Fields	3019
	Table 275: show route extensive Output Fields	3022
	Table 276: show route flow validation Output Fields	3037
	Table 277: show route forwarding-table Output Fields	3042
	Table 278: show route instance Output Fields	3058
	Table 279: show route martians Output Fields	3068
	Table 280: show route receive-protocol Output Fields	3095
	Table 281: show route resolution Output Fields	3104
	Table 282: show route summary Output Fields	3120
	Table 283: show route terse Output Fields	3135

Part 12	Border Gateway Protocol	
Chapter 37	Configuration	3149
	Table 284: MED Options for Routing Table Path Selection	3209
	Table 285: Default Route Preference Values	3326
	Table 286: Damping Parameters	3457
Chapter 38	Administration	3583
	Table 287: show bgp group Output Fields	3590
	Table 288: show bgp neighbor Output Fields	3597
	Table 289: show bgp summary Output Fields	3610
	Table 290: show policy damping Output Fields	3616
	Table 291: show route damping Output Fields	3618
Part 13	Intermediate System to Intermediate System	
Chapter 40	Configuration	3633
	Table 292: IPv4 Statements	3669
	Table 293: IPv6 Statements	3669
	Table 294: Default Metric Values for Routes Exported into IS-IS	3727
Chapter 41	Administration	3755
	Table 295: show isis adjacency Output Fields	3765
	Table 296: show isis authentication Output Fields	3768
	Table 297: show isis database Output Fields	3771
	Table 298: show isis hostname Output Fields	3777
	Table 299: show isis interface Output Fields	3780
	Table 300: show isis overview Output Fields	3784
	Table 301: show isis route Output Fields	3788
	Table 302: show isis statistics Output Fields	3792
Part 14	Open Shortest Path First	
Chapter 42	Overview	3797
	Table 303: Default Route Preference Values for OSPF	3800
Chapter 44	Administration	4031
	Table 304: show (ospf ospf3) backup coverage Output Fields	4043
	Table 305: show (ospf ospf3) backup neighbor Output Fields	4046
	Table 306: show ospf context-identifier Output Fields	4049
	Table 307: show ospf database Output Fields	4051
	Table 308: show ospf3 database Output Fields	4059
	Table 309: show (ospf ospf3) interface Output Fields	4070
	Table 310: show (ospf ospf3) io-statistics Output Fields	4075
	Table 311: show (ospf ospf3) log Output Fields	4077
	Table 312: show (ospf ospf3) neighbor Output Fields	4081
	Table 313: show ospf overview Output Fields	4087
	Table 314: show (ospf ospf3) route Output Fields	4092
	Table 315: show (ospf ospf3) statistics Output Fields	4097

Part 15	Routing Information Protocol	
Chapter 46	Configuration	4109
	Table 316: Configuring Simple RIP Authentication	4121
	Table 317: Configuring MD5 RIP Authentication	4122
Chapter 47	Administration	4197
	Table 318: show rip general-statistics Output Fields	4200
	Table 319: show rip neighbor Output Fields	4203
	Table 320: show rip statistics Output Fields	4205
Part 16	Multicast	
Chapter 48	Overview	4209
	Table 321: ASM and SSM Terminology	4239
Chapter 49	Configuration	4245
	Table 322: PIM Join Filter Match Conditions	4281
	Table 323: IGMP Event Messages	4312
	Table 324: Components of the IGMP Snooping Topology	4319
	Table 325: Source-Active Message Flooding Explanation	4334
Chapter 50	Administration	4487
	Table 326: Summary of IGMP Snooping Output Fields	4487
	Table 327: mtrace Output Fields	4512
	Table 328: mtrace from-source Output Fields	4516
	Table 329: mtrace monitor Output Fields	4518
	Table 330: mtrace to-gateway Output Fields	4521
	Table 331: show igmp group Output Fields	4523
	Table 332: show igmp group Output Fields	4525
	Table 333: show igmp interface Output Fields	4529
	Table 334: show igmp statistics Output Fields	4533
	Table 335: show igmp-snooping membership Output Fields	4536
	Table 336: show igmp-snooping route Output Fields	4539
	Table 337: show igmp-snooping statistics Output Fields	4541
	Table 338: show igmp-snooping vlans Output Fields	4543
	Table 339: show msdp Output Fields	4545
	Table 340: show msdp source Output Fields	4548
	Table 341: show msdp source-active Output Fields	4550
	Table 342: show msdp statistics Output Fields	4552
	Table 343: show multicast flow-map Output Fields	4556
	Table 344: show multicast interface Output Fields	4558
	Table 345: show multicast minfo Output Fields	4560
	Table 346: show multicast next-hops Output Fields	4563
	Table 347: show multicast pim-to-igmp-proxy Output Fields	4565
	Table 348: show multicast pim-to-mld-proxy Output Fields	4567
	Table 349: show multicast route Output Fields	4570
	Table 350: show multicast rpf Output Fields	4576
	Table 351: show multicast scope Output Fields	4579
	Table 352: show multicast sessions Output Fields	4581
	Table 353: show multicast usage Output Fields	4585

	Table 354: show pim bootstrap Output Fields	4587
	Table 355: show pim interfaces Output Fields	4589
	Table 356: show pim join Output Fields	4593
	Table 357: show pim neighbors Output Fields	4602
	Table 358: show pim rps Output Fields	4606
	Table 359: show pim source Output Fields	4613
	Table 360: show pim statistics Output Fields	4616
Part 17	Security	
Chapter 51	Overview	4635
	Table 361: Actions for Firewall Filters	4644
	Table 362: Supported Match Conditions for Firewall Filters	4645
	Table 363: Actions for Firewall Filters	4653
	Table 364: Action Modifiers for Firewall Filters	4654
	Table 365: Policer Actions	4666
	Table 366: Color-Blind Mode TCM Color-to-PLP Mapping	4669
	Table 367: Color-Aware Mode Single-Rate PLP Mapping	4670
	Table 368: Color-Blind Mode TCM Color-to-PLP Mapping	4671
	Table 369: Color-Aware Mode Two-Rate PLP Mapping	4672
Chapter 52	Configuration	4697
	Table 370: Servers Connected to Switch	4704
	Table 371: Components of the Port Security Topology	4708
	Table 372: Components of the Port Security Topology	4716
	Table 373: Components of the Port Security Topology	4720
	Table 374: Components of the Port Security Topology	4723
	Table 375: Components of Port Security Setup on Switch 1 with a DHCP Server Connected to Switch 2	4727
	Table 376: Components of the Port Security Topology	4734
	Table 377: Components of the Port Security Topology	4738
	Table 378: Unicast Forwarding Classes	4751
Chapter 53	Administration	4845
	Table 379: show arp inspection statistics Output Fields	4861
	Table 380: show dhcp snooping binding Output Fields	4862
	Table 381: show firewall Output Fields	4864
	Table 382: show firewall policer Output Fields	4868
	Table 383: show interfaces filters Output Fields	4870
Part 18	Services	
Chapter 55	Overview	4887
	Table 384: Port Mirroring Terms and Definitions	4888
Chapter 57	Administration	4947
	Table 385: show analyzer Output Fields	4948
Part 19	Storage	
Chapter 59	Overview	4955

	Table 386: Fibre Channel Protocol Layers	4957
	Table 387: Load-Balancing Algorithm Comparison	5012
	Table 388: Load-Balancing Triggers and Actions	5015
	Table 389: FC Interface Session-Based Load-Balancing Characteristics for Unequal Loads	5017
	Table 390: FC Interface Session-Based Load-Balancing Characteristics for Equal Loads	5018
	Table 391: VFP TCAM Entry Consumption Summary	5042
	Table 392: Summary of Differences Between IEEE DCBX and DCBX Version 1.01	5053
	Table 393: Asymmetric Ethernet PAUSE Flow Control Configuration	5066
	Table 394: Flow Control State Advertised to the Connected Peer (Autonegotiation)	5067
	Table 395: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces	5068
	Table 396: Default PFC Priority to Queue and Forwarding Class Mapping	5070
	Table 397: Fibre Channel Terms	5074
	Table 398: Default Fabric Forwarding Class Sets	5087
	Table 399: Default Forwarding Class to Fabric Forwarding Class Set Mapping	5088
	Table 400: Class Group Scheduling Properties and Membership	5093
	Table 401: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	5096
	Table 402: show class-of-service forwarding-class-set Command Output Fields	5097
	Table 403: Summary of Differences Between Fabric fc-sets and Node Device fc-sets	5098
Chapter 60	Configuration	5099
	Table 404: Components of the Fibre Channel Interface Configuration Topology	5101
	Table 405: Components of the PFC for FCoE Traffic Configuration Topology	5114
	Table 406: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology	5123
	Table 407: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)	5146
	Table 408: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)	5146
	Table 409: Components of DCBX Application Protocol Exchange Configuration Topology	5147
	Table 410: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)	5156
	Table 411: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to Different FCoE Transit Switches)	5161
	Table 412: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Indirectly Connected Across an Aggregation Layer FCoE Transit Switch)	5168
Chapter 61	Administration	5269

Table 413: Summary of Key FC Interface Load-Balancing Output Fields	5270
Table 414: show fibre-channel proxy fabric-state Output Fields	5271
Table 415: request fibre-channel proxy load-rebalance dry-run Output Fields . .	5287
Table 416: show dcbx output fields	5299
Table 417: show dcbx neighbors Output Fields	5300
Table 418: show fibre-channel fabric Output Fields	5322
Table 419: show fibre-channel fc2 sessions Output Fields	5324
Table 420: show fibre-channel fc2 statistics Output Fields	5326
Table 421: show fibre-channel fip Output Fields	5328
Table 422: show fibre-channel fip enode Output Fields	5333
Table 423: show fibre-channel fip fabric Output Fields	5337
Table 424: show fibre-channel fip fcf Output Fields	5340
Table 425: show fibre-channel fip interface Output Fields	5343
Table 426: show fibre-channel fip statistics Output Fields	5346
Table 427: show fibre-channel flogi fport Output Fields	5350
Table 428: show fibre-channel flogi nport Output Fields	5352
Table 429: show fibre-channel flogi statistics Output Fields	5354
Table 430: show fibre-channel interfaces Output Fields	5357
Table 431: show fibre-channel next-hops Output Fields	5360
Table 432: show fibre-channel routes Output Fields	5362
Table 433: show fibre-channel proxy fabric-state Output Fields	5364
Table 434: show fibre-channel proxy login-table Output Fields	5368
Table 435: show fibre-channel proxy np-port Output Fields	5371
Table 436: show fibre-channel proxy statistics Output Fields	5374
Table 437: show fip snooping Output Fields	5377
Table 438: show fip snooping enode Output Fields	5381
Table 439: show fip snooping fcf Output Fields	5385
Table 440: show fip snooping interface Output Fields	5388
Table 441: show fip snooping statistics Output Fields	5391
Table 442: show fip snooping vlan Output Fields	5394
Table 443: show fip vlan-discovery Output Fields	5398
Table 444: show route forwarding-table family fibre-channel Output Fields . .	5401

Part 20

Chapter 63

Traffic Management

Overview	5411
Table 445: Junos OS Release 11.1 and 11.2 Default Forwarding Classes and Queue Mapping	5421
Table 446: Junos OS Release 11.3 Default Forwarding Classes and Queue Mapping	5422
Table 447: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Unicast Classifiers	5422
Table 448: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers	5422
Table 449: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Multidestination Classifiers	5423
Table 450: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers	5423
Table 451: Junos OS Release 11.1 and 11.2 Default Schedulers	5424
Table 452: Default Schedulers	5425

Table 453: Policer Actions	5432
Table 454: CoS Mappings—Inputs and Outputs	5442
Table 455: Default Forwarding Classes and Queue Mapping	5443
Table 456: Default IEEE 802.1 Code-Point Aliases	5444
Table 457: Default DSCP and DCSP IPv6 Code-Point Aliases	5445
Table 458: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)	5446
Table 459: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)	5447
Table 460: Default IEEE 802.1 Multidestination Classifiers	5447
Table 461: Default DSCP IP and IPv6 Unicast Classifiers	5448
Table 462: Default Drop Profile	5449
Table 463: Default Schedulers	5449
Table 464: Default Scheduler Maps	5449
Table 465: Default Ingress Shared Buffer Configuration	5450
Table 466: Default Egress Shared Buffer Configuration	5450
Table 467: Routing Engine Protocol Default Queue Mapping	5452
Table 468: Default IEEE 802.1 Code-Point Aliases	5453
Table 469: Default DSCP and DSCP IPv6 Code-Point Aliases	5454
Table 470: Default BA Classification	5457
Table 471: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping	5458
Table 472: Supported Classifiers and Rewrite Rules	5461
Table 473: Ethernet Interface Support for Classifier and Rewrite Rule Configuration	5463
Table 474: Default Forwarding Classes for Unicast Packets	5470
Table 475: Default Forwarding Classes for Multicast Packets	5471
Table 476: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)	5476
Table 477: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)	5476
Table 478: Default IEEE 802.1 Multidestination Classifiers	5477
Table 479: Default DSCP IP and IPv6 Unicast Classifiers	5477
Table 480: Default Scheduler Configuration	5478
Table 481: Hierarchical Scheduling Tiers	5481
Table 482: Output Queue Scheduler Components	5486
Table 483: Other Scheduling Components	5487
Table 484: Default Schedulers	5487
Table 485: Priority Group Scheduler Components	5493
Table 486: Other Scheduling Components	5494
Table 487: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute	5509
Table 488: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3	5511
Table 489: Default Output Flow Control Profile	5516
Table 490: User-Configured Output Flow Control Profile	5517
Table 491: Results of Lossless Priority Configuration	5520
Table 492: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler)	5532

Table 493: Egress Queue Dedicated Buffer Allocation (Example 1)	5534
Table 494: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2)	5535
Table 495: Default Shared Ingress Buffer Values (KB)	5537
Table 496: Default Shared Ingress Buffer Values (Percentage)	5537
Table 497: Default Shared Egress Buffer Values (KB)	5538
Table 498: Default Shared Egress Buffer Values (Percentage)	5538
Table 499: Default Ingress Shared Buffer Configuration	5539
Table 500: Default Egress Shared Buffer Configuration	5539
Table 501: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic	5540
Table 502: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic	5540
Table 503: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled	5540
Table 504: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled	5541
Table 505: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic	5541
Table 506: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic	5541
Table 507: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic	5542
Table 508: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic	5542
Table 509: Asymmetric Ethernet PAUSE Flow Control Configuration	5556
Table 510: Flow Control State Advertised to the Connected Peer (Autonegotiation)	5557
Table 511: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces	5558
Table 512: Default PFC Priority to Queue and Forwarding Class Mapping	5560
Table 513: Summary of Differences Between IEEE DCBX and DCBX Version 1.01	5566
Table 514: Default Fabric Forwarding Class Sets	5579
Table 515: Default Forwarding Class to Fabric Forwarding Class Set Mapping	5580
Table 516: Class Group Scheduling Properties and Membership	5585
Table 517: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	5588
Table 518: show class-of-service forwarding-class-set Command Output Fields	5589
Table 519: Summary of Differences Between Fabric fc-sets and Node Device fc-sets	5590
Table 520: Class Group Default Scheduling Properties and Membership on Node Device Fabric Interfaces	5593
Table 521: Hierarchical Scheduler Architecture on Node Devices and Interconnect Devices	5595
Table 522: Class Group Default Scheduling Properties and Membership	5597
Table 523: Default Fabric FC-Set Scheduler Configuration	5598

	Table 524: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	5604
Chapter 64	Configuration	5605
	Table 525: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology	5608
	Table 526: Components of the PFC for FCoE Traffic Configuration Topology	5629
	Table 527: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology	5638
	Table 528: ba-ucast-classifier Loss Priority Assignments	5659
	Table 529: BA-mcast-classifier Loss Priority Assignments	5662
	Table 530: Forwarding-Class-to-Queue Example Configuration	5670
	Table 531: Components of the Forwarding Class Sets Configuration Example	5672
	Table 532: Components of the Queue Scheduler Configuration Example	5676
	Table 533: Components of the Queue Scheduler Priority Configuration Example	5680
	Table 534: Components of the Minimum Guaranteed Output Bandwidth Configuration Example	5686
	Table 535: Components of the Maximum Output Bandwidth Configuration Example	5690
	Table 536: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3	5694
	Table 537: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology	5702
	Table 538: Components of the Two Lossless FCoE Priorities Configuration Topology	5712
	Table 539: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology	5725
	Table 540: Components of the IEEE 802.1p Priority Remapping Configuration Topology	5741
	Table 541: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies	5750
	Table 542: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE	5756
	Table 543: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies	5761
	Table 544: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies	5767
	Table 545: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)	5772
	Table 546: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)	5772
	Table 547: Components of DCBX Application Protocol Exchange Configuration Topology	5773
	Table 548: Default Egress Shared Buffer Partitioning	5824
	Table 549: Default Ingress Shared Buffer Partitioning	5826
	Table 550: Default Output Queue Buffer Sizes	5829
Chapter 65	Administration	5915

	Table 551: Summary of Key CoS Classifier Output Fields	5915
	Table 552: Summary of Key CoS Forwarding Class Output Fields	5917
	Table 553: Summary of Key CoS Interfaces Output Fields	5918
	Table 554: Summary of Key CoS Rewrite Rule Output Fields	5919
	Table 555: Summary of Key CoS Scheduler Maps Output Fields	5919
	Table 556: Summary of Key CoS Value Alias Output Fields	5921
	Table 557: show class-of-service Output Fields	5923
	Table 558: show class-of-service classifier Output Fields	5927
	Table 559: show class-of-service code-point-aliases Output Fields	5929
	Table 560: show class-of-service congestion-notification Output Fields	5931
	Table 561: show class-of-service drop-profile Output Fields	5933
	Table 562: show class-of-service forwarding-class Output Fields	5936
	Table 563: show class-of-service forwarding-class-set Output Fields	5938
	Table 564: show class-of-service forwarding-table classifier Output Fields	5944
	Table 565: show class-of-service forwarding-table classifier mapping Output Fields	5946
	Table 566: show class-of-service forwarding-table drop-profile Output Fields	5948
	Table 567: show class-of-service forwarding-table rewrite-rule Output Fields	5950
	Table 568: show class-of-service forwarding-table rewrite-rule mapping Output Fields	5952
	Table 569: show class-of-service forwarding-table scheduler-map Output Fields	5953
	Table 570: show class-of-service interface Output Fields	5956
	Table 571: show class-of-service multi-destination Output Fields	5981
	Table 572: show class-of-service rewrite-rule Output Fields	5982
	Table 573: show class-of-service scheduler-map Output Fields	5984
	Table 574: show class-of-service shared-buffer Output Fields	5986
	Table 575: show class-of-service traffic-control-profile Output Fields	5988
	Table 576: show dcbx output fields	5992
	Table 577: show dcbx neighbors Output Fields	5993
	Table 578: Layer 2 Overhead, Transmitted Packets/Bytes	6016
	Table 579: show interfaces queue Output Fields	6018
	Table 580: Byte Count by PIC Type	6021
Chapter 66	Troubleshooting	6065
	Table 581: Components of the Rate Shaping Troubleshooting Example	6073
Part 21	Network Management and Monitoring	
Chapter 67	Overview	6077
	Table 582: Device and Network Management Features on the QFX Series	6077
	Table 583: Differences Between Persistent and Transient Changes	6092
	Table 584: Device and Network Management Tasks in Junos Space	6102
	Table 585: Fabric Chassis MIB Tables and Objects	6113
	Table 586: Fabric Chassis MIB SNMPv2 Traps	6115
	Table 587: RMON Event Table	6121
	Table 588: RMON Alarm Table	6121
	Table 589: jnxRmon Alarm Table	6122

	Table 590: RMON History Control Table	6122
	Table 591: Monitored Object Instances	6124
	Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches . .	6125
	Table 593: Juniper Networks Enterprise-Specific MIBs Supported on QFX3500 and QFX3600 Switches	6130
	Table 594: Standard MIBs Supported on QFabric Systems	6133
	Table 595: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems	6137
	Table 596: Standard SNMP Version 1 Traps Supported on QFX3500 and QFX3600 Switches	6140
	Table 597: Enterprise-Specific SNMPv1 Traps Supported on QFX3500 and QFX3600 Switches	6142
	Table 598: Standard SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches	6144
	Table 599: Enterprise-Specific SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches	6147
	Table 600: Standard SNMPv2 Traps Supported on QFabric Systems	6149
	Table 601: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems	6150
Chapter 68	Configuration	6159
	Table 602: Fabric OAM Configuration Elements	6173
	Table 603: Minimum Configuration Statements for System Logging	6207
	Table 604: Fields in Structured-Data Messages	6217
	Table 605: Facility and Severity Codes in the priority-code Field	6218
	Table 606: Fields in Standard-Format Messages	6220
	Table 607: Junos OS System Logging Facilities	6222
	Table 608: System Log Message Severity Levels	6223
	Table 609: Default Facilities for Messages Directed to a Remote Destination . .	6224
	Table 610: Facilities for the facility-override Statement	6224
	Table 611: Regular Expression Operators for the match Statement	6228
Chapter 69	Administration	6385
	Table 612: Output Control Keys for the monitor interface Command	6388
	Table 613: SNMP Tracing Flags	6395
	Table 614: Commit Script Configuration and Operational Mode Commands . .	6398
	Table 615: Match Conditions for the monitor traffic Command	6407
	Table 616: Logical Operators for the monitor traffic Command	6408
	Table 617: Arithmetic and Relational Operators for the monitor traffic Command	6410
	Table 618: ping fabric multicast-flow Output Fields	6419
	Table 619: ping fabric unicast-flow Output Fields	6422
	Table 620: show oam fabric flow-specification Output Fields	6432
	Table 621: show oam fabric interfaces Output Fields	6435
	Table 622: show sflow Output Fields	6440
	Table 623: show sflow collector Output Fields	6442
	Table 624: show sflow interface Output Fields	6443
	Table 625: show snmp health-monitor Output Fields	6445
	Table 626: show snmp inform-statistics Output Fields	6450
	Table 627: show snmp mib Output Fields	6453

	Table 628: show snmp rmon Output Fields	6455
	Table 629: show snmp statistics Output Fields	6460
	Table 630: show snmp v3 Output Fields	6465
	Table 631: traceroute fabric unicast-flow Output Fields	6468
Chapter 70	Troubleshooting	6471
	Table 632: Troubleshooting Resources on the QFX Series	6471
	Table 633: Troubleshooting on the QFX Series	6473
Part 22	Troubleshooting	
Chapter 71	Overview	6483
	Table 634: Troubleshooting Resources on the QFX Series	6483
	Table 635: Troubleshooting on the QFX Series	6485
	Table 636: Alarm Terms and Definitions	6488
	Table 637: QFX3500 Chassis Alarm Messages	6489
Chapter 72	Administration	6493
	Table 638: SNMP Tracing Flags	6497
	Table 639: Output Control Keys for the monitor interface Command	6503
Chapter 73	Troubleshooting	6505
	Table 640: Components of the Rate Shaping Troubleshooting Example	6560

About the Documentation

- Documentation and Release Notes on page cxxi
- Supported Platforms on page cxxi
- Using the Examples in This Manual on page cxxi
- Documentation Conventions on page cxxiii
- Documentation Feedback on page cxxv
- Requesting Technical Support on page cxxv

Documentation and Release Notes

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

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

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

Supported Platforms

For the features described in this document, the following platforms are supported:

- QFabric System

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

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

Documentation Conventions

Table 1 on page cxxiii defines notice icons used in this guide.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name domain-name
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options { static { route default { nexthop address; retain; } } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

QFX3500 Switch Overview

- [QFX3500 Switch Overview on page 3](#)

CHAPTER 1

QFX3500 Switch Overview

- [QFX3500 Device Overview on page 3](#)

QFX3500 Device Overview

The Juniper Networks QFX3500 device is a high-speed, multipurpose switch especially designed for next-generation data centers. Forty-eight 10-Gbps access ports in the device use small form-factor pluggable plus (SFP+) transceivers and operate by default as 10-Gigabit Ethernet interfaces. Optionally, you can choose to configure up to 12 of the ports as 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel (FC) interfaces, and up to 36 of the ports as 1-Gigabit Ethernet interfaces. When used as a standalone switch, four 40-Gbps uplink ports in the device use quad small form-factor pluggable plus (QSFP+) to four SFP+ copper breakout cables to support an additional 15 10-Gigabit Ethernet interfaces.

QFX3500 devices can function as a Fibre Channel over Ethernet (FCoE)-FC gateway or as an FCoE transit switch. FCoE is a method of supporting converged FC and Ethernet traffic on a data center bridging (DCB) network by encapsulating unmodified FC frames in Ethernet to transport the FC frames over the physical Ethernet network.

In a QFabric system, a QFX3500 device functions as a Node device, connected to a QFabric system through 40-Gbps uplink ports to a Juniper Networks QFX3008-I or QFX3600-I Interconnect device. Together, the QFX3500 Node devices and QFX3008-I or QFX3600-I Interconnect devices form a multistage, nonblocking switch fabric that provides a high-performance, low-latency, unified interconnect solution for next-generation data centers.

The QFX3500 Node devices and Interconnect devices are connected to Juniper Networks QFX3100 Director devices in an out-of-band management network through Juniper Networks EX4200 Ethernet Switches. The QFX3100 Director devices present the QFabric system devices as a single network entity, which enables simplified management of your data center using the Junos OS command-line interface (CLI).

- [Software on page 3](#)
- [Hardware on page 4](#)

Software

QFX Series devices use the Junos operating system (OS), which provides Layer 2 and Layer 3 switching, routing, and security services. Junos OS is installed on the QFX3500

device's 8-gigabyte (GB) internal flash drive. The same Junos OS code base that runs on QFX3500 devices also runs on all Juniper Networks EX Series switches, and J Series, M Series, MX Series, and T Series routers.

For more information about which features are supported on QFX Series devices, see [“QFX Series Software Features Overview” on page 15](#).

When the QFX3500 device is operating as a standalone switch, you manage the switch using the Junos OS command-line interface (CLI), accessible through the console and out-of-band management ports on the device.

When a QFX Series device operates as part of a QFabric system, all the devices in the data center fabric are managed through the Administrator software installed on the QFX3100 Director devices. Each device in a QFabric system is interconnected in a single control plane and management network, using the redundant management ports on each device.

Hardware

The compact QFX3500 device is 1 rack unit (1 U) in size and designed to fit in industry-standard 19-inch rack-mount enclosures. See [Figure 1 on page 4](#) and [Figure 2 on page 4](#) and *Chassis Physical Specifications for a QFX3500 Chassis*.

Figure 1: QFX3500 Device Front





Figure 2: QFX3500 Device Rear



The front panel of the QFX3500 chassis has an LCD panel that displays the device hostname and the number of active alarms. See *Front Panel of a QFX3500 Device*. The rear panel has 48 10-Gbps access ports and 4 40-Gbps uplink ports. See *Rear Panel of a QFX3500 Device*.

SFP+ Access Ports

The QFX3500 device has 48 access ports (0 through 47) that support small form-factor pluggable plus (SFP+) and small form-factor pluggable (SFP) transceivers, as well as SFP+ direct attach copper cables, also known as Twinax cables. See *Interface Support for the QFX3500 Device*.

- Up to 48 of the access ports can be used for SFP+ transceivers or SFP+ direct attach copper cables. You can use 10-Gigabit Ethernet SFP+ transceivers and SFP+ direct attach copper cables in any access port. You can use 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel SFP+ transceivers in ports 0 through 5 and ports 42 through 47.



NOTE: If you use Fibre Channel SFP+ transceivers in ports 0 through 5 or ports 42 through 47, you must configure the entire block of ports as Fibre Channel ports. For example, if you use a Fibre Channel SFP+ transceiver in any of the ports 0 through 5, then ports 0 through 5 must be configured as Fibre Channel ports. If you use a Fibre Channel SFP+ transceiver in any of the ports 42 through 47, then ports 42 through 47 must be configured as Fibre Channel ports. You then cannot use 10-Gigabit Ethernet SFP+ transceivers in these ports.


- Up to 36 of the access ports can be used for SFP transceivers. Gigabit Ethernet SFP transceivers can be used in ports 6 through 41.



CAUTION: Do not place a copper transceiver in an access port directly above or below another copper transceiver. Internal damage to the access ports and device can occur. Because of this limitation, a maximum of 18 copper transceivers can be installed in ports 6 through 41. We recommend using only the top row of access ports for copper transceivers.

Figure 3 on page 5 shows the location of the SFP+ access ports, including the ports that can be used with Fibre Channel SFP+ transceivers and Gigabit Ethernet SFP transceivers.

Figure 3: SFP+ Access Port Locations



QSFP+ Uplink Ports

The QFX3500 device has four uplink ports (**Q0** through **Q3**) that support up to four QSFP+ transceivers, as well as QSFP+ DAC or DAC breakout cables. See *Interface Support for the QFX3500 Device*.


When the QFX3500 device is used as part of a QFabric system, these uplink ports are used to connect the QFX3500 Node device to QFX3008-I or QFX3600-I Interconnect devices. See *Connecting a QFX3500 Node Device to a QFX3008-I Interconnect Device* or *Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device*.

When the QFX3500 device is used as a standalone switch, these uplink ports are configured by default as 15 10-Gigabit Ethernet interfaces. Port **Q0** supports only three 10-Gigabit Ethernet interfaces. Ports **Q1** through **Q3** support four 10-Gigabit Ethernet interfaces. Optionally, you can configure one or more of the ports as 40-Gigabit Ethernet interfaces. [Figure 4 on page 6](#) shows the location of the QSFP+ uplink ports and the default 10-Gigabit Ethernet interface numbering.



NOTE: The QSFP+ uplink ports are not supported in Junos OS Release 11.1. The QSFP+ uplink ports are supported in Junos OS Release 11.2 and later. To configure the ports as 40-Gigabit Ethernet interfaces, you must be using Junos OS Release 12.2X50-D20 or later.

Figure 4: QSFP+ Uplink Port Locations



*Port availability is release dependent.
See the topic on Channelizing Interfaces for your Junos Release.

Related Documentation

- *Field-Replaceable Units in a QFX3500 Device*
- *Site Preparation Checklist for a QFX3500 Device*
- *Access Port and Uplink Port LEDs on a QFX3500 Device*
- *Installing and Removing QFX3500 Device Hardware Components*

PART 2

QFX3600 Switch Overview

- [QFX3600 Switch Overview on page 9](#)

CHAPTER 2

QFX3600 Switch Overview

- [QFX3600 Device Overview on page 9](#)

QFX3600 Device Overview

The Juniper Networks QFX3600 device is a high-speed, multipurpose switch especially designed for next-generation data centers. Sixteen 40-Gbps ports in the device use quad small form-factor pluggable plus (QSFP+) transceivers. The small form-factor and front facing ports in the switch make it suitable for deployment in high-density server racks and container-based data center deployments.

The QFX3600 device runs in three modes: standalone switch, a QFX Node in a QFabric system, and as a interconnect device in QFabric system. You can configure the device to run in any of these modes by using the **request chassis device-mode** commands.

In a QFabric system, a QFX3600 device functions as a Node device, connected to a QFabric system through 40-Gbps uplink ports to a Juniper Networks QFX3008-I or QFX3600-I Interconnect device. Together, the QFX3600 Node devices and QFX3008-I or QFX3600-I Interconnect devices form a multistage, nonblocking switch fabric that provides a high-performance, low-latency, unified interconnect solution for next-generation data centers.

The QFX3600 Node devices and QFX3008-I or QFX3600-I Interconnect devices are connected to Juniper Networks QFX3100 Director devices in a control plane and management network. The QFX3100 Director device presents the QFabric system devices as a single network entity, allowing for simplified management of your data center using the Junos OS command-line interface (CLI).

- [Software on page 9](#)
- [Hardware on page 10](#)

Software

QFX Series devices use the Junos operating system (OS), which provides Layer 2 and Layer 3 switching, routing, and security services. Junos OS is installed on the QFX3600 device's 8-gigabyte (GB) internal flash drive. The same Junos OS code base that runs on QFX Series devices also runs on all Juniper Networks EX Series, J Series, M Series, MX Series, and T Series devices.

For more information about which features are supported on QFX Series devices, see [“QFX Series Software Features Overview” on page 15](#).

When the QFX3600 device is operating as a standalone switch, you manage the switch using the Junos OS command-line interface (CLI), accessible through the console and out-of-band management ports on the device.

When a QFX Series device operates as part of a QFabric system, all the devices in the data center fabric are managed through the Administrator software installed on the QFX3100 Director devices. Each device in a QFabric system is interconnected in a single control plane and management network, using the redundant management ports on each device.

Hardware

The compact QFX3600 chassis is 1 rack unit (1 U) in size and designed to fit in industry-standard 19-inch rack-mount enclosures, as well as high-density server racks and container-based data center deployments. See [Figure 5 on page 10](#) and [Figure 6 on page 10](#) and *Chassis Physical Specifications for QFX3600 and QFX3600-I Devices*.

Figure 5: QFX3600 Chassis Front





Figure 6: QFX3600 Chassis Rear



The front panel of the QFX3600 chassis has an LCD panel that displays the device hostname and the number of active alarms. It also has sixteen 40-Gbps ports labeled

Q0 through **Q15** that support quad small form-factor pluggable plus (QSFP+) transceivers. See *Front Panel of a QFX3600 Device*.

If you are using the QFX3600 device as a Node device in a QFabric system, by default, four ports (labeled **Q0** through **Q3**) are configured for uplink connections between your QFX3600 Node device and your Interconnect device, and twelve ports (labeled **Q4** through **Q15**) support 48 10-Gigabit Ethernet or 12 40-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (labeled **Q0** through **Q7**) for uplink connections between your QFX3600 Node device and your Interconnect device and ports **Q2** through **Q15** for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks. See [“Configuring the Port Type on QFX3600 Node Devices” on page 1367](#) and *Interface Support for the QFX3600 Device*.

If you are using the QFX3600 device as a standalone switch, by default, all 16 QSFP+ ports (**Q0** through **Q15**) are configured as 40-Gigabit Ethernet (xle) ports. Optionally, you can choose to configure each port to operate as 10-Gigabit Ethernet (xe) ports to support up to 63 10-Gigabit Ethernet ports. See [“Configuring the Port Type on QFX3600 Standalone Switches” on page 2545](#) and *Interface Support for the QFX3600 Device*.

The rear panel of the QFX3600 chassis has two redundant power supplies and three redundant fan trays that are field-replaceable and hot-swappable. See *Rear Panel of QFX3600 and QFX3600-I Devices*.

Related Documentation

- *Field-Replaceable Units for QFX3600 and QFX3600-I Devices*
- *Site Preparation Checklist for a QFX3600 or QFX3600-I Device*
- *Installing and Removing QFX3600 or QFX3600-I Device Hardware Components*
- [Converting the Device Mode for a QFabric System Component on page 1258](#)

PART 3

Software Feature Support

- [Software Feature Support for the QFX Series on page 15](#)

CHAPTER 3

Software Feature Support for the QFX Series

- [QFX Series Software Features Overview on page 15](#)

QFX Series Software Features Overview

The following tables list the Juniper Networks QFX Series software features and the Juniper Networks Junos operating system (Junos OS) release in which they were introduced:

- [Table 3 on page 16](#)—Access Control Features
- [Table 4 on page 16](#)—Administration Features
- [Table 5 on page 16](#)—Class-of-Service (CoS) Features
- [Table 6 on page 17](#)—High Availability and Resiliency Features
- [Table 7 on page 19](#)—Interface Features
- [Table 8 on page 19](#)—IP Address Management Features
- [Table 9 on page 20](#)—Layer 2 Network Protocol Features
- [Table 10 on page 22](#)—Layer 3 Protocol Features
- [Table 11 on page 23](#)—Multicast Protocol Features
- [Table 12 on page 24](#)—Network Management and Monitoring Features
- [Table 13 on page 26](#)—Port Security Features
- [Table 14 on page 26](#)—QFabric System-Specific Features
- [Table 15 on page 27](#)—Security
- [Table 16 on page 28](#)—Storage and Fibre Channel Features
- [Table 17 on page 29](#)—System Management Features



NOTE: The command-line interface (CLI) on the QFX Series might display configuration statements that are not supported. However, configuring an unsupported statement on a device has no effect on the operation of the device. See [“QFX Series CLI Hierarchy” on page 192](#) to see which statements are supported on the QFX Series.

Table 3: Access Control Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
MAC RADIUS authentication	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
TACACS+	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
SSH authentication keys	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15

Table 4: Administration Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
System logging (syslog) over IPv4	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Licensing	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10

Table 5: CoS Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Class of service (CoS)—Class-based queuing with prioritization	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS—Multidestination	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS support on link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 5: CoS Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Enhanced transmission selection (ETS)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Interface-specific CoS rewrite rules	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Port shaping and queue shaping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control (PFC)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control enhancements	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Re-marking of bridged packets	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority remapping	Junos OS 12.3X50-D10	Not supported	Not supported	Not supported
Weighted random early detection (WRED) tail-drop profiles	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
802.3X Ethernet PAUSE autonegotiation enhancements	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Layer 3 ingress packet classification and egress rewrite rule class-of-service features	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Software buffer configurability	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 6: High Availability and Resiliency Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Graceful protocol restart for BGP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10

Table 6: High Availability and Resiliency Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Graceful protocol restart for OSPF	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Support for 32 members in a link aggregation group (LAG)	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Multichassis link aggregation	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10	Not supported	Not supported
Virtual Router Redundancy Protocol (VRRP)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Graceful restart for FIP snooping to recover gracefully from Ethernet protocol restarts and crashes	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Nonstop software upgrade (NSSU)	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Layer 3 unicast and multicast support for multichassis link aggregation	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 7: Interface Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Interface ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
VLAN-tagged Layer 3 logical interfaces	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 8: IP Address Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Static addresses	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IPv4 support for telnet	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081 .	Junos OS 12.2X50-D10. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081 .
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, virtual routing instances, firewall filters, RADIUS, NTP, and SNMP, on network Node group	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, firewall filters, and SNMP on management interfaces	Junos OS 12.2X50-D20	Not supported	Not applicable	Not applicable

Table 8: IP Address Management Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IPv6 support for neighbor discovery, router advertisements, stateless autoconfiguration, link aggregation, SSH, Telnet, ping, traceroute, path MTU, static routing, dynamic routing (BGPv6, IS-ISv6, OSPFv3, RIPv6), graceful restart, BFD, virtual routers, firewall filters, SNMP, CoS, VRRPv3, Radius, TACACS+, AAA, NTP, and syslog, on network interfaces	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not applicable	Not applicable

Table 9: Layer 2 Network Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
802.1Q VLAN tagging	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
BPDU protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)
Jumbo frames on routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Link Layer Discovery Protocol (LLDP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Loop protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 9: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Root protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Spanning tree:	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)
<ul style="list-style-type: none"> Spanning Tree Protocol (STP) Rapid Spanning Tree Protocol (RSTP) Multiple Spanning Tree Protocol (MSTP) VLAN Spanning Tree Protocol (VSTP) RSTP and VSTP concurrent configuration 				
Multiple VLAN Registration Protocol (MVRP)	Not supported	Not supported	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15
VLAN ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Private VLANs	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Secondary VLAN trunk ports and promiscuous access ports for private VLANs	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Q-in-Q tunneling	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15

Table 9: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
VLAN translation	Junos OS 12.1X49-D1	Junos OS 12.1X49-D1	Not supported	Not supported
Layer 2 Tunneling Protocol	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Reflective relay	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Proxy ARP	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported

Table 10: Layer 3 Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Bidirectional Forwarding Detection (BFD)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Border Gateway Protocol (BGP) A separate software license is required for BGP. See “Software Features That Require Licenses on the QFX Series” on page 87.	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
BGPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Intermediate System-to-Intermediate System (IS-IS) A separate software license is required for IS-IS. See “Software Features That Require Licenses on the QFX Series” on page 87.	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IS-ISv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Open Shortest Path First (OSPF) v2	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
OSPFv3	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 10: Layer 3 Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Per-packet load balancing (ECMP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Routing Information Protocol versions 1 and 2 (RIPv1 and RIPv2)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
RIPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Routing options	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Static routes	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Virtual router routing instances for unicast protocols	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Virtual router routing instances for multicast protocols	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 11: Multicast Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IGMPv1 and v2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
IGMPv3	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IGMPv1 and v2 snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP snooping with routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP querier	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
IGMP filtering	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 11: Multicast Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Protocol Independent Multicast sparse mode (PIM SM)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Protocol Independent Multicast source-specific multicast (PIM SSM)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Bidirectional Forwarding Detection (BFD) for PIM	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static RP	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Anycast RP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast Source Discovery Protocol (MSDP)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast VLAN registration	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 12: Network Management and Monitoring Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Junos Space	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Local port mirroring	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Remote port mirroring	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Remote port mirroring to IP address (GRE encapsulation)	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 12: Network Management and Monitoring Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Increased numbers of port-mirroring sessions	Not supported	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
RMON	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
sFlow monitoring technology	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Simple Network Management Protocol version 1 (SNMPv1) and SNMPv2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Simple Network Management Protocol version 3 (SNMPv3)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Juniper enterprise-specific SNMP Utility MIB	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Juniper enterprise-specific SNMP Fabric Chassis MIB	Not applicable	Not applicable	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)
Juniper Class-of-Service MIB	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Uplink failure detection	Junos OS 11.3R4	Junos OS 12.2X50-D20	Not supported	Not supported
Junos OS automation script support	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system control plane debugging and diagnostics support	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Internal fabric OAM monitoring	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DHCP smart relay	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 13: Port Security Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Automatic recovery for port error disable conditions	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC move limiting	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Persistent MAC learning (sticky MAC)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static ARP support	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Storm control (broadcast, unicast, and multicast)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DHCP snooping	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Dynamic ARP inspection (DAI)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 14: QFabric System-Specific Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Aliasing of Node devices	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
High availability of QFabric system components	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Node groups	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 14: QFabric System-Specific Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Aliases for Director devices and Interconnect devices	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system backup and recovery (software and configuration)	Not applicable	Not applicable	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 15: Security

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Firewall filters and rate limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on LAGs	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on loopback interfaces	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on management interfaces	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported
Policing	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increased numbers of supported filters and policers	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Filter-based forwarding	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Not supported	Not supported

(For a list of supported firewall filter match conditions and actions, see [“Firewall Filter Match Conditions and Actions” on page 4644.](#))

Table 16: Storage and Fibre Channel Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Data center bridging technologies	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Data Center Bridging Capability Exchange protocol (DCBX)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
FCoE Initialization Protocol (FIP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Fibre Channel fabrics	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel interfaces	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel over Ethernet (FCoE) to Fibre Channel gateway	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
FCoE-FC gateway link automated load balancing	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
FCoE OxID hash control	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
FIP snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increase to 2500 FIP snooping sessions	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
Increase to 2500 FIP snooping sessions when QFX3500 switch acts as an FCoE-FC gateway.	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
Transit switch	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DCBX application protocol TLV exchange	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3R1	Junos OS 12.2X50-D10

Table 16: Storage and Fibre Channel Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
User-Configurable DCBX application protocol TLV exchange	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DCBX version support for IEEE DCBX and DCBX version 1.01	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Graceful restart for FCoE-FC gateway	Junos OS 12.1X49-D1	Not supported	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel load balancing, maximum session limit, disabling fabric WWN verification check	Junos OS 12.1X49-D1	Not supported	Junos OS 12.1X49-D1 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
VN_Port to VN_Port FIP snooping to enable configuring virtual links between VN_Ports without sending traffic through an FCoE forwarder (FCF)	Junos OS 12.2X50-D10	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 17: System Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Configuration rollback	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IP directed broadcast	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Online insertion and removal (OIR)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 17: System Management Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
J-Web interface, for switch configuration and management	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Zero Touch Provisioning	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported

Related Documentation

- *QFX3000-G QFabric System Hardware Documentation*
- *QFX3000-M QFabric System Hardware Documentation*
- *QFX3500 Device Hardware Documentation*
- *QFX3600 Device Hardware Documentation*

PART 4

Junos OS Basics

- [Overview on page 33](#)
- [Installation on page 99](#)
- [Configuration on page 141](#)
- [Administration on page 315](#)
- [Troubleshooting on page 1157](#)

CHAPTER 4

Overview

- [Software Overview on page 33](#)
- [User Interfaces on page 70](#)
- [Licenses on page 86](#)

Software Overview

- [QFX Series Software Features Overview on page 33](#)
- [Configuration File Terms on page 48](#)
- [Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements on page 49](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 50](#)
- [Junos OS Package Names on page 51](#)
- [NTP Time Server and Time Services Overview on page 52](#)
- [NTP Time Server and Time Services Overview \(QFabric System\) on page 53](#)
- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 53](#)
- [Understanding Autoinstallation of Configuration Files on page 55](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [Understanding Zero Touch Provisioning on page 61](#)
- [Understanding DHCP Services for Switches on page 63](#)
- [Understanding Software Infrastructure and Processes on page 67](#)

QFX Series Software Features Overview

The following tables list the Juniper Networks QFX Series software features and the Juniper Networks Junos operating system (Junos OS) release in which they were introduced:

- [Table 3 on page 16](#)—Access Control Features
- [Table 4 on page 16](#)—Administration Features
- [Table 5 on page 16](#)—Class-of-Service (CoS) Features
- [Table 6 on page 17](#)—High Availability and Resiliency Features

- [Table 7 on page 19](#)—Interface Features
- [Table 8 on page 19](#)—IP Address Management Features
- [Table 9 on page 20](#)—Layer 2 Network Protocol Features
- [Table 10 on page 22](#)—Layer 3 Protocol Features
- [Table 11 on page 23](#)—Multicast Protocol Features
- [Table 12 on page 24](#)—Network Management and Monitoring Features
- [Table 13 on page 26](#)—Port Security Features
- [Table 14 on page 26](#)—QFabric System-Specific Features
- [Table 15 on page 27](#)—Security
- [Table 16 on page 28](#)—Storage and Fibre Channel Features
- [Table 17 on page 29](#)—System Management Features



NOTE: The command-line interface (CLI) on the QFX Series might display configuration statements that are not supported. However, configuring an unsupported statement on a device has no effect on the operation of the device. See “[QFX Series CLI Hierarchy](#)” on page 192 to see which statements are supported on the QFX Series.

Table 18: Access Control Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
MAC RADIUS authentication	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
TACACS+	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
SSH authentication keys	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15

Table 19: Administration Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
System logging (syslog) over IPv4	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Licensing	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10

Table 20: CoS Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Class of service (CoS)—Class-based queuing with prioritization	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS—Multidestination	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS support on link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Enhanced transmission selection (ETS)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Interface-specific CoS rewrite rules	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Port shaping and queue shaping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control (PFC)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control enhancements	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Re-marking of bridged packets	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority remapping	Junos OS 12.3X50-D10	Not supported	Not supported	Not supported
Weighted random early detection (WRED) tail-drop profiles	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 20: CoS Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
802.3X Ethernet PAUSE autonegotiation enhancements	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Layer 3 ingress packet classification and egress rewrite rule class-of-service features	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Software buffer configurability	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 21: High Availability and Resiliency Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Graceful protocol restart for BGP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Graceful protocol restart for OSPF	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Support for 32 members in a link aggregation group (LAG)	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Multichassis link aggregation	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10	Not supported	Not supported
Virtual Router Redundancy Protocol (VRRP)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 21: High Availability and Resiliency Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Graceful restart for FIP snooping to recover gracefully from Ethernet protocol restarts and crashes	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Nonstop software upgrade (NSSU)	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Layer 3 unicast and multicast support for multichassis link aggregation	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 22: Interface Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Interface ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
VLAN-tagged Layer 3 logical interfaces	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 23: IP Address Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Static addresses	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IPv4 support for telnet	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081.	Junos OS 12.2X50-D10. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081.

Table 23: IP Address Management Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, virtual routing instances, firewall filters, RADIUS, NTP, and SNMP, on network Node group	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, firewall filters, and SNMP on management interfaces	Junos OS 12.2X50-D20	Not supported	Not applicable	Not applicable
IPv6 support for neighbor discovery, router advertisements, stateless autoconfiguration, link aggregation, SSH, Telnet, ping, traceroute, path MTU, static routing, dynamic routing (BGPv6, IS-ISv6, OSPFv3, RIPv6), graceful restart, BFD, virtual routers, firewall filters, SNMP, CoS, VRRPv3, Radius, TACACS+, AAA, NTP, and syslog, on network interfaces	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not applicable	Not applicable

Table 24: Layer 2 Network Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
802.1Q VLAN tagging	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 24: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
BPDU protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)
Jumbo frames on routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Link Layer Discovery Protocol (LLDP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Loop protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Root protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Spanning tree: <ul style="list-style-type: none"> • Spanning Tree Protocol (STP) • Rapid Spanning Tree Protocol (RSTP) • Multiple Spanning Tree Protocol (MSTP) • VLAN Spanning Tree Protocol (VSTP) • RSTP and VSTP concurrent configuration 	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)

Table 24: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Multiple VLAN Registration Protocol (MVRP)	Not supported	Not supported	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15
VLAN ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Private VLANs	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Secondary VLAN trunk ports and promiscuous access ports for private VLANs	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Q-in-Q tunneling	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15
VLAN translation	Junos OS 12.1X49-D1	Junos OS 12.1X49-D1	Not supported	Not supported
Layer 2 Tunneling Protocol	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Reflective relay	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Proxy ARP	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported

Table 25: Layer 3 Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Bidirectional Forwarding Detection (BFD)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Border Gateway Protocol (BGP)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
A separate software license is required for BGP. See “Software Features That Require Licenses on the QFX Series” on page 87.				

Table 25: Layer 3 Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
BGPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Intermediate System-to-Intermediate System (IS-IS) A separate software license is required for IS-IS. See “Software Features That Require Licenses on the QFX Series” on page 87.	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IS-ISv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Open Shortest Path First (OSPF) v2	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
OSPFv3	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Per-packet load balancing (ECMP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Routing Information Protocol versions 1 and 2 (RIPv1 and RIPv2)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
RIPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Routing options	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Static routes	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Virtual router routing instances for unicast protocols	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Virtual router routing instances for multicast protocols	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 26: Multicast Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IGMPv1 and v2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
IGMPv3	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IGMPv1 and v2 snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP snooping with routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP querier	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
IGMP filtering	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Protocol Independent Multicast sparse mode (PIM SM)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Protocol Independent Multicast source-specific multicast (PIM SSM)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Bidirectional Forwarding Detection (BFD) for PIM	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static RP	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Anycast RP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast Source Discovery Protocol (MSDP)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast VLAN registration	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 27: Network Management and Monitoring Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Junos Space	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Local port mirroring	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Remote port mirroring	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Remote port mirroring to IP address (GRE encapsulation)	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Increased numbers of port-mirroring sessions	Not supported	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
RMON	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
sFlow monitoring technology	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Simple Network Management Protocol version 1 (SNMPv1) and SNMPv2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Simple Network Management Protocol version 3 (SNMPv3)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Juniper enterprise-specific SNMP Utility MIB	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Juniper enterprise-specific SNMP Fabric Chassis MIB	Not applicable	Not applicable	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)
Juniper Class-of-Service MIB	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Uplink failure detection	Junos OS 11.3R4	Junos OS 12.2X50-D20	Not supported	Not supported

Table 27: Network Management and Monitoring Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Junos OS automation script support	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system control plane debugging and diagnostics support	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Internal fabric OAM monitoring	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DHCP smart relay	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 28: Port Security Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Automatic recovery for port error disable conditions	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC move limiting	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Persistent MAC learning (sticky MAC)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static ARP support	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Storm control (broadcast, unicast, and multicast)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DHCP snooping	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 28: Port Security Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Dynamic ARP inspection (DAI)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 29: QFabric System-Specific Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Aliasing of Node devices	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
High availability of QFabric system components	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Node groups	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Aliases for Director devices and Interconnect devices	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system backup and recovery (software and configuration)	Not applicable	Not applicable	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 30: Security

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Firewall filters and rate limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on LAGs	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 30: Security (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Firewall filters on loopback interfaces	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on management interfaces	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported
Policing	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increased numbers of supported filters and policers	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Filter-based forwarding	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Not supported	Not supported

(For a list of supported firewall filter match conditions and actions, see [“Firewall Filter Match Conditions and Actions” on page 4644.](#))

Table 31: Storage and Fibre Channel Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Data center bridging technologies	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Data Center Bridging Capability Exchange protocol (DCBX)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
FCoE Initialization Protocol (FIP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Fibre Channel fabrics	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel interfaces	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)

Table 31: Storage and Fibre Channel Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Fibre Channel over Ethernet (FCoE) to Fibre Channel gateway	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
FCoE-FC gateway link automated load balancing	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
FCoE OxID hash control	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
FIP snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increase to 2500 FIP snooping sessions	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
Increase to 2500 FIP snooping sessions when QFX3500 switch acts as an FCoE-FC gateway.	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
Transit switch	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DCBX application protocol TLV exchange	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3R1	Junos OS 12.2X50-D10
User-Configurable DCBX application protocol TLV exchange	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DCBX version support for IEEE DCBX and DCBX version 1.01	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Graceful restart for FCoE-FC gateway	Junos OS 12.1X49-D1	Not supported	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel load balancing, maximum session limit, disabling fabric WWN verification check	Junos OS 12.1X49-D1	Not supported	Junos OS 12.1X49-D1 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)

Table 31: Storage and Fibre Channel Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
VN_Port to VN_Port FIP snooping to enable configuring virtual links between VN_Ports without sending traffic through an FCoE forwarder (FCF)	Junos OS 12.2X50-D10	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 32: System Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Configuration rollback	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IP directed broadcast	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Online insertion and removal (OIR)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
J-Web interface, for switch configuration and management	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Zero Touch Provisioning	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported

Related Documentation

- [QFX3000-G QFabric System Hardware Documentation](#)
- [QFX3000-M QFabric System Hardware Documentation](#)
- [QFX3500 Device Hardware Documentation](#)
- [QFX3600 Device Hardware Documentation](#)

Configuration File Terms

Table 33 on page 49 lists the various configuration file terms used for the QFX Series and their definitions.

Table 33: Configuration File Terms

Term	Definition
active configuration	Current committed configuration of a switch.
candidate configuration	Working copy of the configuration that allows users to make configurational changes without causing any operational changes until this copy is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Check configuration for proper syntax, activate and mark as the current configuration file running on the switching platform.
configuration hierarchy	Junos OS configuration consists of a hierarchy of statements. There are two types of statements: container statements, which contain other statements, and leaf statements, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
default configuration	Default configuration contains the initial values set for each configuration parameter when a switch is shipped.
rescue configuration	Well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the CLI.
roll back a configuration	Return to a previously committed configuration.

**Related
Documentation**

- [Loading a Previous Configuration File on page 1600](#)
- [Reverting to the Rescue Configuration on page 175](#)
- [Understanding Configuration Files on page 1590](#)

Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements

Many statements in the Junos OS configuration include an option to specify an IP address or route prefix. This option is represented in one of the following ways:

- **network/prefix-length**—Network portion of the IP address, followed by a slash and the destination prefix length (previously called the subnet mask). For example, 10.0.0.1/8.
- **network**—IP address. For example, 10.0.0.2.
- **destination-prefix/prefix-length**—Route prefix, followed by a slash and the destination prefix length. For example, 192.168.1.10/32.

You enter all IP addresses in classless mode. You can enter the IP address with or without a prefix length, in standard dotted notation (for example, 1.2.3.4), or hexadecimal notation as a 32-bit number in network-byte order (for example, 0x01020304). If you omit any octets, they are assumed to be zero. Specify the prefix length as a decimal number from 1 through 32.

- Related Documentation**
- [Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 73](#)

Junos OS Commit Model for Router or Switch Configuration

The router or switch configuration is saved using a commit model: that is, a candidate configuration is modified as desired and then committed to the system. Once a configuration has been committed, the router or switch checks the configuration for syntax errors, and if no errors are found, the configuration is saved as **juniper.conf.gz** and activated. The former active configuration file is saved as the first rollback configuration file (**juniper.conf.1.gz**), and all other rollback configuration files are incremented by 1. For example, **juniper.conf.1.gz** is incremented to **juniper.conf.2.gz**, making it the second rollback configuration file. The router or switch can have a maximum of 49 rollback configurations (1–49) saved on the system.

On the router or switch, the active configuration file and the first three rollback files (**juniper.conf.gz.1**, **juniper.conf.gz.2**, **juniper.conf.gz.3**) are located in the **/config** directory. If the file **rescue.conf.gz** is saved on the system, this file should also be saved in the **/config** directory. The factory default files are located in the **/etc/config** directory.

There are two mechanisms used to propagate the configurations between Routing Engines within a router or switch:

- Synchronization—Propagates a configuration from one Routing Engine to a second Routing Engine within the same router or switch chassis.



NOTE: The QFX3500 switch has only one Routing Engine.

To synchronize configurations, use the **commit synchronize** CLI command. If one of the Routing Engines is locked, the synchronization fails. If synchronization fails because of a locked configuration file, you can use the **commit synchronize force** command. This command overrides the lock and synchronizes the configuration files.

- Distribution—Propagates a configuration across the routing plane on a multichassis router or switch. Distribution occurs automatically. There is no user command available to control the distribution process. If a configuration is locked during a distribution of a configuration, the locked configuration does not receive the distributed configuration file, so the synchronization fails. You need to clear the lock before the configuration and resynchronize the routing planes.



NOTE: When you use the **commit synchronize force** CLI command on a multichassis platform, the forced synchronization of the configuration files does not affect the distribution of the configuration file across the routing plane. If a configuration file is locked on a router or switch remote from the router or switch where the command was issued, the synchronization fails on the remote router or switch. You need to clear the lock and reissue the **synchronize** command.

**Related
Documentation**

- [Configuring Junos OS for the First Time on a Router or Switch with a Single Routing Engine](#)

Junos OS Package Names

You upgrade the Juniper Networks Junos OS on the QFX Series by copying a software package to your switch or another system on your local network and then installing the new software package on the switch.

A software package name is in the following format:



NOTE: A signed domestic package is used as an example only. Other types of software packages might be available in future releases.

package-name-m.nZx.y-domestic-signed.tgz

where:

- ***package-name*** is the name of the package—for example, ***jinstall-qfx***.
- ***m.n*** is the software release, with ***m*** representing the major release number and ***n*** representing the minor release number—for example, ***11.1***.
- ***Z*** indicates the type of software release, where ***R*** indicates released software and ***B*** indicates beta-level software.
- ***x.y*** represents the maintenance software release, with ***x*** representing the maintenance software release number and ***y*** representing the maintenance software spin number—for example, ***1.5***.

A sample switch software package name is:

jinstall-qfx-11.1R1.5-domestic-signed.tgz

**Related
Documentation**

- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Upgrading Software on a QFabric System on page 134](#)
- [Software Installation Overview on page 120](#)

NTP Time Server and Time Services Overview

When configuring the Network Time Protocol (NTP), you can specify which system on the network is the authoritative time source, or time server, and how time is synchronized between systems on the network. To do this, you configure the router or switch to operate in one of the following modes:

- Client mode—In this mode, the local router or switch can be synchronized with the remote system, but the remote system can never be synchronized with the local router or switch.
- Symmetric active mode—In this mode, the local router or switch and the remote system can synchronize with each other. You use this mode in a network in which either the local router or switch or the remote system might be a better source of time.



NOTE: Symmetric active mode can be initiated by either the local or the remote system. Only one system needs to be configured to do so. This means that the local system can synchronize with any system that offers symmetric active mode without any configuration whatsoever. However, we strongly encourage you to configure authentication to ensure that the local system synchronizes only with known time servers.

- Broadcast mode—In this mode, the local router or switch sends periodic broadcast messages to a client population at the specified broadcast or multicast address. Normally, you include this statement only when the local router or switch is operating as a transmitter.
- Server mode—In this mode, the local router or switch operates as an NTP server.



NOTE: In NTP server mode, the Junos OS supports authentication as follows:

- If the NTP request from the client comes with an authentication key (such as a key ID and message digest sent with the packet), the request is processed and answered based on the authentication key match.
- If the NTP request from the client comes without any authentication key, the request is processed and answered without authentication.

Related Documentation

- [Configuring the NTP Time Server and Time Services on page 153](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)

NTP Time Server and Time Services Overview (QFabric System)

Network Time Protocol (NTP) synchronizes the time among all of the various devices in the QFabric system to an external time server. This ensures that time-stamped events, like log entries and database transactions, can be correlated between multiple devices in the QFabric system. The QFabric system can only be in client mode.

Additionally, you can authenticate time synchronization to ensure that the QFabric system obtains its time services only from known sources. By default, network time synchronization is unauthenticated. We strongly encourage you to configure authentication of network time services.

Related Documentation

- [Configuring NTP Authentication Keys \(QFabric System\) on page 153](#)
- [authentication-key on page 240](#)
- [ntp on page 277](#)
- [server on page 289](#)

Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)

Before you upgrade to Junos OS Release 11.3, you must deactivate the CoS configuration if the CoS configuration includes any of the following features:

- **excess-rate** option
- **strict-high** or **high** priority queues
- Any of the Junos OS Release 11.1 or 11.2 default multidestination forwarding classes



CAUTION: If your CoS configuration contains any of the features listed above and you attempt to upgrade from Junos OS Release 11.1 or 11.2 to a later version without first editing the configuration, the Junos OS might not restart.

Junos OS Release 11.3 and later for QFX Series no longer supports the **excess-rate** statement, the **strict** priority option, or the default multidestination forwarding classes used in Junos OS Release 11.1 and 11.2. In addition, Junos OS Release 11.3 introduces new restrictions on how to configure and use **strict-high** priority queues.

This topic does not describe how to perform the software upgrade procedure. It describes how to deactivate your CoS configuration, edit your CoS configuration, and reactivate your CoS configuration at the appropriate times.

Use the following procedure to upgrade safely from Junos OS Release 11.1 or 11.2 to a later release:

1. Deactivate the CoS configuration *before* you upgrade the software:

```
user@switch# deactivate class-of-service
```
2. Follow the upgrade procedure to Junos OS Release 11.3 or later software.

3. Make the following changes to the CoS configuration while the CoS configuration is still deactivated:
 - Remove the **excess-rate** statement from the CoS configuration if you have used it at the **[edit class-of-service schedulers]** or **[edit class-of-service traffic-control-profiles]** hierarchy level.
 - Remove the **strict-high** and **strict** priority queue configurations if you have used them at the **[edit class-of-service schedulers]** hierarchy level.
 - Remove the default multdestination forwarding classes (**mcast-be**, **mcast-af**, **mcast-ef**, and **mcast-nc**) if you have used them at the **[edit class-of-service schedulers]**, **[edit class-of-service rewrite-rules]**, **[edit class-of-service classifiers]**, **[edit class-of-service scheduler-maps]**, or **[edit class-of-service forwarding-class-sets]** hierarchy level. Alternatively, you can change the mapping of the multdestination traffic to use the new default multdestination forwarding class (**mcast**).
4. If desired, configure **strict-high** priority queues in accordance with the Junos OS Release 11.3 or later configuration rules, and map multdestination traffic to the default multdestination forwarding class (**mcast**).
5. Activate the CoS configuration:

```
user@switch# activate class-of-service
```
6. Commit the CoS configuration:

```
user@switch# commit
```



NOTE: If you configured the **transmit-rate** option for any queues under the **[edit class-of-service schedulers]** hierarchy level, if the rate is configured as an exact rate in Mbps, we recommend that you reconfigure the **transmit-rate** option as a percentage. This is because the scheduler converts exact rates to percentages, and when the exact rate is below 1 Gbps, some granularity may be lost in the conversion. You can avoid this potential issue by specifying the **transmit-rate** option as a percentage.

Related Documentation

- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 5416](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 5418](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)

Understanding Autoinstallation of Configuration Files

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

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

This topic describes:

- [Typical Uses for Autoinstallation on page 55](#)
- [Autoinstallation Configuration Files and IP Addresses on page 55](#)
- [Typical Autoinstallation Process on a New Switch on page 56](#)

Typical Uses for Autoinstallation

Typical uses for autoinstallation of the software include:

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

Autoinstallation Configuration Files and IP Addresses

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

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

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

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

Typical Autoinstallation Process on a New Switch

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

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

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

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

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

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

- Related Documentation**
- *Configuring Autoinstallation of Configuration Files (CLI Procedure)*
 - *Connecting and Configuring an EX Series Switch (CLI Procedure)*
 - *Connecting and Configuring an EX Series Switch (J-Web Procedure)*
 - *Configuration Files Terms*

Understanding Nonstop Software Upgrade for QFabric Systems

The framework that underlies a nonstop software upgrade in a QFabric system enables you to upgrade the system in a step-by-step manner and minimize the impact to the continuous operation of the system. This topic explains how a nonstop software upgrade works in a QFabric system, the steps that are involved, and the procedures that you need to implement to experience the benefits of this style of software upgrade.

Nonstop software upgrade enables some QFabric system components to continue operating while similar components in the system are being upgraded. In general, the QFabric system upgrades redundant components in stages so that some components remain operational and continue forwarding traffic while their equivalent counterparts upgrade to a new version of software.



TIP: Use the following guidelines to decide when to implement a nonstop software upgrade:

- If you need to upgrade all components of the system in the shortest amount of time (approximately one hour) and you do not need to retain the forwarding resiliency of the data plane, issue the `request system software add component all` command to perform a standard software upgrade. All components of the QFabric system upgrade simultaneously and expediently, but this type of upgrade does not provide resiliency or switchover capabilities.
- If you need to minimize service impact, preserve the forwarding operations of the data plane during the upgrade, and are willing to take the extra time required for component switchovers (in many cases, several hours), issue the three nonstop software upgrade commands (`request system software nonstop-upgrade (director-group | fabric | node-group)`) described in this topic in the correct order.



NOTE:

- Before you begin a nonstop software upgrade, issue the **request system software download** command to copy the software to the QFabric system.
 - Each of the 3 nonstop software upgrade steps must be considered parts of the whole process. You must complete all 3 steps of a nonstop software upgrade in the correct order to ensure the proper operation of the QFabric system.
 - Open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected. For more information on verification of the upgrade, see [“Verifying Nonstop Software Upgrade for QFabric Systems” on page 1410](#).
 - Issue the **show fabric administration inventory** command to verify that all upgraded components are operational at the end of a step before beginning the next step.
 - Once you start the nonstop software upgrade process, we strongly recommend that you complete all 3 steps within 12 hours.
-

The three steps to a successful nonstop software upgrade must be performed in the following order:

- Director group—The first step upgrades the Director devices, the fabric manager Routing Engine, and the diagnostic Routing Engine. To perform the first step, issue the **request system software nonstop-upgrade director-group** command. The key actions that occur during a Director group upgrade are:
 1. Connecting to the QFabric system by way of an SSH connection. This action establishes a load-balanced CLI session on one of the Director devices in the Director group.
 2. The QFabric system downloads and installs the new software in both Director devices.
 3. The Director device hosting the CLI session becomes the master for all QFabric system processes running on the Director group, such as the fabric manager and network Node group Routing Engines.
 4. The QFabric system installs the new software for the backup fabric manager Routing Engine on the backup Director device.
 5. The backup Director device reboots to activate the new software.
 6. The master Director device begins a 15 minute sequence that includes a temporary suspension of QFabric services and a QFabric database transfer. You cannot issue operational mode commands in the QFabric CLI during this period.
 7. The QFabric system installs the new software for the fabric manager and diagnostic Routing Engines on the Director group master.

8. The QFabric system switches mastership of all QFabric processes from the master Director device to the backup Director device.
9. The master Director device reboots to activate the new software.
10. The CLI session terminates, and logging back in to the QFabric system with a new SSH connection establishes the session on the new master Director device (the original backup).
11. The previous master Director device resumes operation as a backup and the associated processes (such as the fabric manager and network Node group Routing Engines) become backup as well. The fabric control Routing Engine associated with this Director device returns to active status.



NOTE: After the Director group nonstop software upgrade completes, any Interconnect device or Node device that reboots will automatically download the new software, install it, and reboot again. As a result, try not to restart any QFabric system devices before you complete the rest of the nonstop software upgrade steps.



TIP:

- To enable BGP and OSPF to continue operating on the network Node group during a Director group nonstop service upgrade, we recommend that you configure graceful restart for these routing protocols. For more information on graceful restart, see [“Configuring Graceful Restart for QFabric Systems” on page 1375](#).
 - Wait 15 minutes after the second Director device returns to service and hosts Routing Engine processes before proceeding to step 2—the fabric upgrade. You can verify the operational status of both Director devices by issuing the `show fabric administration inventory director-group status` command. Also, issue the `show fabric administration inventory infrastructure` command to verify when the Routing Engine processes become load balanced (typically, there will be three to four Routing Engines running on each Director device).
-
- Fabric—The second step upgrades the Interconnect devices and the fabric control Routing Engines. To perform the second step, issue the **request system software nonstop-upgrade fabric** command. The key actions that occur during a fabric upgrade are:
 1. The QFabric system downloads, validates, and installs the new software in all Interconnect devices and fabric control Routing Engines (FC-0 and FC-1).
 2. One fabric control Routing Engine reboots and comes back online.
 3. The other fabric control Routing Engine reboots and comes back online.

4. The first Interconnect device reboots, comes back online, and resumes the forwarding of traffic.
5. Subsequent Interconnect devices reboot one at a time, come back online, and return to service.

**NOTE:**

- If the software does not load properly on any one of the fabric components, all components revert back to the original software version.
- If one of the components in a fabric upgrade does not reboot successfully, issue the **request system reboot fabric** command to reattempt the rebooting process for this fabric component and activate the new software.

- **Node group**—The third and final step upgrades Node groups. You can choose to upgrade a network Node group, a redundant server Node group, or individual server Node groups. You can upgrade the Node groups one at a time or in groups (known as upgrade groups). However, you must upgrade all Node groups in your QFabric system before you can complete the nonstop software upgrade process. To perform the third step, issue the **request system software nonstop-upgrade node-group** command.

The key actions that occur during a network Node group upgrade are:

1. The QFabric system copies the new software to each Node device one at a time.
2. The QFabric system validates and then installs the new software in all Node devices simultaneously.
3. The system copies the software to the network Node group Routing Engines.
4. The QFabric system validates and then installs the software in the network Node group Routing Engines one at a time -- first the backup, then the master.
5. The backup network Node group Routing Engine reboots and comes back online.
6. The supporting Node devices reboot and come back online one at a time.



NOTE: To reduce the total upgrade duration, configure an upgrade group. All Node devices within the upgrade group reboot at the same time.

7. The master network Node group Routing Engine relinquishes mastership to the backup, reboots, and comes back online.

The key actions that occur during a redundant server Node group upgrade are:

1. The QFabric system copies the new software to the backup Node device, then the master Node device.
2. The QFabric system validates and then installs the new software on the backup Node device, then the master Node device.

3. The backup Node device reboots, comes back online, and becomes the master Node device.
4. The previous master Node device reboots and comes back online as a backup Node device.



NOTE: For redundant server Node groups, both Node devices must be online before the upgrade will proceed. If one of the devices is no longer available, remove the Node device from the Node group configuration before you issue the nonstop software upgrade command.

The key actions that occur during a server Node group upgrade for a Node group that contains one member are:

1. The Node device downloads the software package and validates the software.
2. The Node device installs the software and reboots.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [request system software add on page 394](#)
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)

Understanding Zero Touch Provisioning



NOTE: To see which platforms support Zero Touch Provisioning, in a browser, go to [Feature Explorer](#). In the Explore Features section of the Feature Explorer page, select **All Features**. In the Features Grouped by Feature Family box, select Zero Touch Provisioning. You can also type the name of the feature in the Search for Features edit box. In previous Junos OS releases on EX Series switches, Zero Touch Provisioning was called EZ Touchless Provisioning. Search for that feature name if you want to know if this feature is supported on EX Series switches.

Zero Touch Provisioning allows you to provision new Juniper Networks switches in your network automatically, without manual intervention. When you physically connect a switch to the network and boot it with a default factory configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from

the network. To make sure you have the default factory configuration loaded on the switch, issue the **request system zeroize** command on the switch you want to provision.

The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default factory configuration.

The Zero Touch Provisioning process will either upgrade or downgrade the Junos OS version. During an downgrade:

- On an EX Series switch, If you downgrade to a software version earlier than Junos OS Release 12.2, in which Zero Touch Provisioning is not supported, the configuration file autoinstall phase of the Zero Touch Provisioning process does not happen.
- On an EX Series switch, to downgrade to a software version that does not support resilient dual-root partitions (Junos OS Release 10.4R2 or earlier), you must perform some manual work on the switch. For more information, see *Understanding Resilient Dual-Root Partitions on Switches*.

When you boot a switch with the default factory configuration, the following process happens:

1. If DHCP option 43, suboption 00 (the name of the software image file on the FTP, HTTP, or TFTP server) is configured, the switch compares the version of the provided software image to the version of the software installed on the switch.



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

2. If DHCP option 43, suboption 02 (a symbolic link to the software image file on the FTP, HTTP, or TFTP server), the switch compares the version of the provided software image to the version of the software installed on the switch.
 - If the Junos OS versions are different, the switch downloads the software image from the FTP, HTTP, or TFTP server, installs the Junos OS, and reboots using the default factory configuration.
 - If the software versions are the same, the switch does not upgrade the software.
3. If DHCP option 43, suboption 01 (the name of the configuration file on the FTP, TFTP, or HTTP server is configured, the switch compares the version of the provided configuration file to the version of the configuration file on the switch.

If DHCP option 43 suboption 01 is not specified, the switch uses the default factory configuration.

If the configuration file version on the FTP, HTTP, or TFTP server is newer than the configuration file on the switch, the configuration file is updated on the switch.

If both DHCP option 43 suboption 01 and suboption 2 are specified, suboption 01 is processed before suboption 02. The Junos OS is upgraded, and then the configuration file is applied.

4. If DHCP option 43, suboption 03 (the transfer mode setting) is configured, the switch accesses the FTP, HTTP, or TFTP server using the specified transfer mode setting—for example, FTP.

If DHCP option 43, suboption 03, is not configured, TFTP becomes the transfer mode automatically.

5. If DHCP option 43, suboption 04 (the name of the software image file on the FTP, HTTP, or TFTP server) is configured, the switch compares the version of the provided software image to the version of the software installed on the switch.



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

6. If DHCP option 150 or option 66 is specified, the IP address of the FTP, HTTP, or TFTP server is configured.



NOTE: You must configure either option 150 or option 66. If you configure both option 150 and option 66, option 150 takes precedence, and option 66 is ignored. Also, make sure you specify an IP address, not a hostname, because name resolution is not supported.

7. (Optional) If DHCP option 7 is specified, you can configure one or more syslog servers.
8. (Optional) If DHCP option 42 is specified, you can configure one or more NTP servers.
9. (Optional) If DHCP option 12 is specified, you can configure the hostname of the switch.

Related Documentation

- [Configuring Zero Touch Provisioning on page 101](#)

Understanding DHCP Services for Switches

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

This topic describes:


- [DHCP Client/Server Model on page 64](#)
- [Using DHCP on page 64](#)
- [DHCP Relay Servers and DHCP Servers on page 65](#)

- [Legacy DHCP and Extended DHCP for Server Versions on page 65](#)
- [Configuring DHCP on a Switch on page 66](#)
- [How DHCP Works on page 66](#)

DHCP Client/Server Model

DHCP IP address allocation works on a client/server model in which the server, in this case a switch, assigns the client reusable IP information from an address pool. A DHCP client might receive offer messages from multiple DHCP servers and can accept any one of the offers; however, the client usually accepts the first offer it receives. See [Figure 7 on page 64](#).

Figure 7: DHCP Client/Server Model



Using DHCP

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

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

Another valuable DHCP feature is automatic software download for installation of software packages on switches. DHCP clients configured for automatic software download receive messages as part of the DHCP message exchange process—when the software package name in the DHCP server message is different from that of the software

package that booted the DHCP client switch, the new software is downloaded and installed. See *Upgrading Software Using Automatic Software Download*.

DHCP Relay Servers and DHCP Servers

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

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

Legacy DHCP and Extended DHCP for Server Versions

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

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

The extended DHCP server version has the following added features:

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



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

Legacy DHCP and extended DHCP servers can be configured at the hierarchy levels shown in [Table 34 on page 65](#):

Table 34: Legacy DHCP and Extended DHCP Server Hierarchy Levels

DHCP Service	Hierarchy
Extended DHCP server	<code>edit system services dhcp-local-server</code>
Extended DHCP address pool	<code>edit access address-assignment pool</code>

Table 34: Legacy DHCP and Extended DHCP Server Hierarchy Levels (*continued*)

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

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

Configuring DHCP on a Switch

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

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

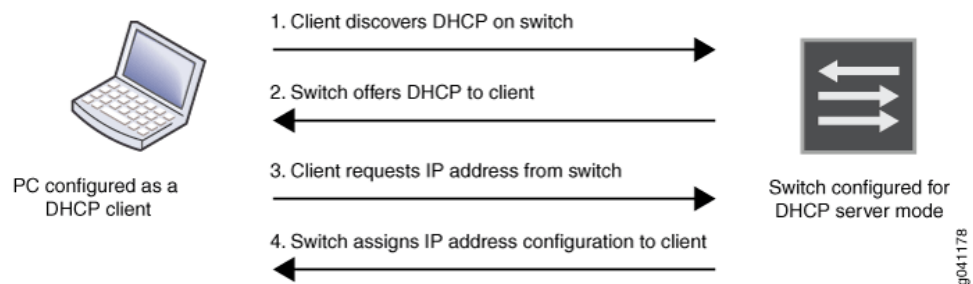
For directions for configuring either a legacy DHCP server or an extended DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 143](#).

To configure a DHCP client, set the client’s DHCP interface address in the `[edit interfaces interface-name unit 0 family inet dhcp]` hierarchy. For directions for configuring a DHCP client on a switch, see [“Configuring a DHCP Client \(CLI Procedure\)” on page 142](#).

How DHCP Works

DHCP consists of a four-step transfer process beginning with a broadcast DHCP discovery message from the client. As the second step, the client receives a DHCP offer message from the server. This message includes the IP address and mask, and some other specific parameters. The client then sends a DHCP request message to accept the IP address and other parameters that it received from the server in the previous step. The DHCP server sends a DHCP response message and removes the now-allocated address from the DHCP address pool. See [Figure 8 on page 67](#).

Figure 8: DHCP Four-Step Transfer



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

Related Documentation

- [Configuring a DHCP Client \(CLI Procedure\) on page 142](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 143](#)
- [Configuring an Extended DHCP Relay Server on EX Series Switches \(CLI Procedure\)](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\)](#)
- [Upgrading Software Using Automatic Software Download](#)
- [Monitoring DHCP Services](#)

Understanding Software Infrastructure and Processes

The QFX Series products run the Juniper Networks Junos OS. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the switch.

Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

Using the Junos OS command-line interface (CLI), you configure switching features and set the properties of network interfaces. After activating a software configuration, use either the Junos Space or CLI user interface to monitor, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 68](#)
- [Junos OS Processes on page 68](#)

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- Packet Forwarding Engine—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- Routing Engine—Provides three main functions:
 - Creates the packet forwarding switch, which provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network.
 - Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
 - Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of the Junos OS for added flexibility.

Table 35 on page 68 describes the primary Junos OS processes.

Table 35: Junos OS Processes

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
DNS Server process	named-service	Resolves hostnames into addresses.
Dynamic Host Configuration Protocol (DHCP) process	dhcp-service	Enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree Protocol, and access port security.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>

Table 35: Junos OS Processes (*continued*)

Process	Name	Description
Firewall management process	firewall	Manages the firewall configuration and helps accept or reject packets that are transiting an interface on a switch.
Forwarding process	pfem	Defines how routing protocols operate on the partition. The overall performance of the partition is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Integrated Local Management Interface (ILMI) process	ilmi	Provides bidirectional exchange of management information between two ATM interfaces across a physical connection.
Link Management Protocol (LMP) process	linkmanagement	Establishes and maintains LMP control channels.
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the partition.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Multicast snooping process	multicast-snooping	Makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
Secure Neighbor Discovery (SEND) Protocol process	send	Protects Neighbor Discovery Protocol (NDP) messages.
Simple Network Management Protocol (SNMP) process	snmp	Enables the monitoring of network devices from a central location and provides the switch's SNMP master agent.
Tunnel OAM process	tunnel-oamd	Enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
Virtual Router Redundancy Protocol (VRRP) process	vrrp	Enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

- Related Documentation**
- *Junos OS Baseline Network Operations Guide*
 - *Junos OS System Basics Configuration Guide*

User Interfaces

- [CLI User Interface Overview on page 70](#)
- [Configuring CLI Tips on page 72](#)
- [Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 73](#)
- [Junos OS Operational Mode Commands That Combine Other Commands on page 74](#)
- [Overview of Junos OS CLI Operational Mode Commands on page 75](#)
- [Overview of Navigating the CLI on page 78](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 79](#)
- [Understanding Junos OS CLI Configuration Mode on page 80](#)

CLI User Interface Overview

- [CLI Overview on page 70](#)
- [CLI Key Features on page 71](#)
- [CLI Command Modes on page 71](#)

CLI Overview

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

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

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

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

CLI Key Features

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

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

CLI Command Modes

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

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

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

In configuration mode, you are actually viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the current operating configuration, called the active configuration. When you commit the changes you added to the candidate configuration, the system updates the active configuration. Candidate configurations enable you to alter your configuration without causing potential damage to your current network operations.

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

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

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

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

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

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

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

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

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

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

Related Documentation

- [Configuring CLI Tips on page 72](#)
- [Overview of Navigating the CLI on page 78](#)
- *CLI User Guide*
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 318](#)

Configuring CLI Tips

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

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

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

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

Related Documentation

- [CLI User Interface Overview on page 70](#)
- [Defining Junos OS Login Classes](#)
- [login-tip on page 265](#)

Format for Specifying Filenames and URLs in Junos OS CLI Commands

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

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

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

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



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

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

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

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

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

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

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

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



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

**Related
Documentation**

- [Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements on page 49](#)
- [Default Directories for Junos OS File Storage on the Router or Switch](#)

Junos OS Operational Mode Commands That Combine Other Commands

In some cases, some Junos OS operational commands are created from a combination of other operational commands. These commands can be useful shortcuts for collecting information about the device, as shown in [Figure 9 on page 75](#).

Figure 9: Commands That Combine Other Commands

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

```

user@host> request support information

root@host> show system uptime

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

root@host> show version detail

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

root@host> show system core-dumps

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

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

root@host> show chassis hardware detail

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

```

Related Documentation

- [Overview of Junos OS CLI Operational Mode Commands on page 75](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 79](#)

Overview of Junos OS CLI Operational Mode Commands

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

- [CLI Command Categories on page 75](#)
- [Commonly Used Operational Mode Commands on page 76](#)

CLI Command Categories

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

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

- **monitor**—Perform real-time debugging of various software components, including the routing protocols and interfaces.
- **ping**—Determine the reachability of a remote network host.
- **show**—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
- **test**—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
- **traceroute**—Trace the route to a remote network host.
- Commands for connecting to other network systems—The **ssh** command opens Secure Shell connections, and the **telnet** command opens telnet sessions to other hosts on the network. For information about these commands, see the [CLI Explorer](#).
- Commands for copying files—The **copy** command copies files from one location on the router or switch to another, from the router or switch to a remote system, or from a remote system to the router or switch. For information about these commands, see the [CLI Explorer](#).
- Commands for restarting software processes—The commands in the **restart** hierarchy restart the various Junos OS processes, including the routing protocol, interface, and SNMP. For information about these commands, see the [CLI Explorer](#).
- A command—**request**—for performing system-level operations, including stopping and rebooting the router or switch and loading Junos OS images. For information about this command, see the [CLI Explorer](#).
- A command—**start**—to exit the CLI and start a UNIX shell. For information about this command, see the [CLI Explorer](#).
- A command—**configure**—for entering configuration mode, which provides a series of commands that configure Junos OS, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see “[Understanding Junos OS CLI Configuration Mode](#)” on page 80.
- A command—**quit**—to exit the CLI. For information about this command, see the [CLI Explorer](#).
- For more information about the CLI operational mode commands, see the [CLI Explorer](#) and the [CLI Explorer](#).

Commonly Used Operational Mode Commands

Table 36 on page 77 lists some operational commands you may find useful for monitoring router or switch operation. For a complete description of operational commands, see the Junos OS command references.



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

Table 36: Commonly Used Operational Mode Commands

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

Table 36: Commonly Used Operational Mode Commands (*continued*)

Items to Check	Description	Command
RSVP	Status of interfaces on which RSVP is running	show rsvp interface
	Currently active RSVP sessions	show rsvp session
	RSVP packet and error counters	show rsvp statistics

Related Documentation

- [Junos OS Operational Mode Commands That Combine Other Commands on page 74](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 79](#)

Overview of Navigating the CLI


This topic describes how to navigate the CLI.

- [CLI Command Hierarchy on page 78](#)
- [CLI Configuration Statements on page 78](#)
- [Moving Among Hierarchy Levels on page 79](#)

CLI Command Hierarchy

CLI commands are organized in a hierarchy. Commands that perform a similar function are grouped together under the same level of the hierarchy. For example, all commands that display information about the system and the system software are grouped under the **show system** command, and all commands that display information about the routing table are grouped under the **show route** command. [Figure 10 on page 78](#) illustrates a portion of the **show** command hierarchy.

Figure 10: CLI Command Hierarchy



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

CLI Configuration Statements

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

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

Moving Among Hierarchy Levels

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

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

Related Documentation

- [CLI User Interface Overview on page 70](#)
- *CLI User Guide*

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

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

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

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

2. Choose the option you wish to use with the command. (See [Figure 11 on page 80](#).)

Figure 11: Command Output Options

Command output with the brief option.	<pre> user@host> show interfaces fe-1/1/1 brief Physical interface: fe-1/1/1, Enabled, Physical link is Down Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Down Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000 Link flags : None </pre>
Command output with the terse option.	<pre> user@host> show interfaces fe-1/1/1 terse Interface Admin Link Proto Local Remote fe-1/1/1 up down </pre>
Command output with the extensive option.	<pre> user@host> show interfaces fe-1/1/1 extensive Physical interface: fe-1/1/1, Enabled, Physical link is Down Interface index: 141, SNMP ifIndex: 33, Generation: 24 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Down Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000 Link flags : None CoS queues : 4 supported, 4 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:90:69:d0:f8:9e, Hardware address: 00:90:69:d0:f8:9e Last flapped : 2007-02-02 09:26:25 PST (2w0d 03:40 ago) Statistics last cleared: Never Traffic statistics: Input bytes : 0 0 bps Output bytes : 0 0 bps Input packets: 0 0 pps Output packets: 0 0 pps --(more)-- </pre>

- Related Documentation**
- [Overview of Junos OS CLI Operational Mode Commands on page 75](#)
 - [Controlling the Scope of an Operational Mode Command](#)

Understanding Junos OS CLI Configuration Mode

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

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

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

This topic covers:

- [Configuration Mode Commands on page 81](#)
- [Configuration Statements and Identifiers on page 82](#)
- [Configuration Statement Hierarchy on page 84](#)

Configuration Mode Commands

Table 37 on page 81 summarizes each CLI configuration mode command. The commands are organized alphabetically.

Table 37: Summary of Configuration Mode Commands

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

Table 37: Summary of Configuration Mode Commands (*continued*)

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

Configuration Statements and Identifiers

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

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

Table 38 on page 83 describes top-level CLI configuration mode statements.



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

Table 38: Configuration Mode Top-Level Statements

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

Table 38: Configuration Mode Top-Level Statements (*continued*)

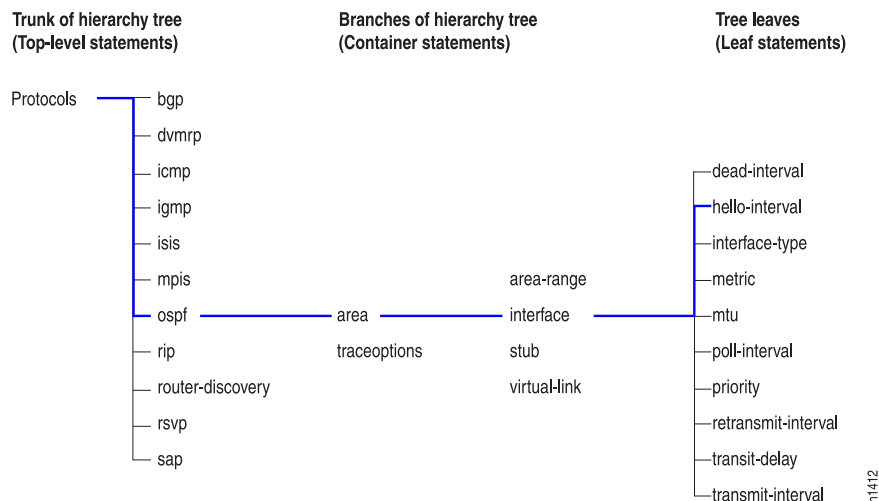
Statement	Description
routing-instances	Configure multiple routing instances. For information about the statements in this hierarchy, see the <i>Junos OS Routing Protocols Configuration Guide</i> .
routing-options	Configure protocol-independent routing options, such as static routes, autonomous system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

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

Configuration Statement Hierarchy

The Junos OS configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which are statements that contain other statements, and *leaf statements*, which do not contain other statements (see [Figure 12 on page 84](#)). All of the container and leaf statements together form the *configuration hierarchy*.

Figure 12: Configuration Mode Hierarchy of Statements



Each statement at the top level of the configuration hierarchy resides at the trunk (or root level) of a hierarchy tree. The top-level statements are container statements, containing other statements that form the tree branches. The leaf statements are the leaves of the hierarchy tree. An individual hierarchy of statements, which starts at the trunk of the hierarchy tree, is called a *statement path*. [Figure 12 on page 84](#) illustrates the hierarchy tree, showing a statement path for the portion of the protocol configuration hierarchy that configures the hello interval on an interface in an OSPF area.

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

The CLI represents the statement path shown in [Figure 12 on page 84](#) as **[edit protocols ospf area area-number interface interface-name]** and displays the configuration as follows:

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

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

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

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

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

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

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

Related Documentation

- [Entering and Exiting the Junos OS CLI Configuration Mode](#)

Licenses

- [Junos OS Feature Licenses on page 86](#)
- [Software Features That Require Licenses on the QFX Series on page 87](#)
- [Junos OS License Keys on page 88](#)
- [Generating the License Keys for a Standalone QFX Series Device on page 90](#)
- [Generating the License Keys for a QFabric System on page 91](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Deleting a License \(CLI Procedure\) on page 94](#)
- [Saving License Keys on page 95](#)
- [Verifying Junos OS License Installation on page 96](#)

Junos OS Feature Licenses

Some Junos OS software features require a license to activate the feature. To enable a licensed feature, you need to purchase, install, manage, and verify a license key that corresponds to each licensed feature. To conform to Junos OS feature licensing requirements, you must purchase one license per feature per device. The presence of the appropriate software license key on your device determines whether you are eligible to configure and use the licensed feature.

To speed deployment of licensed features, Junos OS software implements an honor-based licensing structure and provides you with a 30-day grace period to use a licensed feature without a license key installed. The grace period begins when you configure the feature and your device uses the licensed feature for the first time, but not necessarily when you install the license. After the grace period expires, the system generates system log messages saying that the feature requires a license. To clear the error message and use the licensed feature properly, you must install and verify the required license.

For information about how to purchase software licenses, contact your Juniper Networks sales representative.

Related Documentation

- [License Enforcement](#)
- [Junos OS License Keys on page 88](#)

- *Software Feature Licenses*
- [Verifying Junos OS License Installation on page 96](#)

Software Features That Require Licenses on the QFX Series

The following Junos OS features require an Advanced Feature License (AFL) on QFX Series devices:



NOTE: If you try to configure a feature that is not licensed, you will receive syslog messages saying that you are using a feature that is licensable and that you do not possess a license for the feature. If you try to commit configuration changes for a feature that is not licensed, you will receive a commit warning saying that you have exceeded the allowed license limit for the feature.

- Border Gateway Protocol (BGP)
- Intermediate System-to-Intermediate System (IS-IS)
- Multiprotocol Label Switching (MPLS)
- Fibre Channel support

[Table 39 on page 87](#) lists the licenses you can purchase for each QFX Series software feature.

For information about how to purchase a software license, contact your Juniper Networks sales representative.

Table 39: Junos OS Feature Licenses and Model Numbers for QFX Series Devices

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Multiprotocol Label Switching (MPLS)	QFX3500, QFX3600, and QFX5100 switch	One per switch	QFX-JSL-EDGE-ADV1
QFX Series advanced feature license for Border Gateway Protocol (BGP)	QFX3100 Director device	One per Node device in a network Node group	QFX-JSL-DRCTR-ADV1
QFX Series advanced feature license for Fibre Channel	QFX3500 switch	One per switch on which fibre channel ports are configured	QFX-JSL-EDGE-FC

Table 39: Junos OS Feature Licenses and Model Numbers for QFX Series Devices (*continued*)

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Fibre Channel	QFX3100 Director device	One per QFX3500 Node device on which fibre channel ports are configured	QFX-JSL-DRCTR-FC
QFX Series advanced feature license for Fibre Channel - Capacity 16	QFX3100 Director device	One for up to 16 QFX3500 Node devices on which fibre channel ports are configured	QFX-JSL-DRCTR-FC-C16
QFX Series feature license for enabling fabric mode	QFX3500 and QFX3600 device	One per device	QFX3000-JSL-EDGE-FAB
QFX Series feature license for base software for QFX3000-G QFabric system	QFX3100 Director device	One per QFX3000-G QFabric system	QFX3008-JSL-DRCTR-FAB
QFX Series feature license for base software for QFX3000-M QFabric system	QFX3100 Director device	One per QFX3000-M QFabric system	QFX3000M-JSL-DRCTR-FAB

Related Documentation

- [Junos OS Feature Licenses on page 86](#)
- [Junos OS License Keys on page 88](#)
- [Generating the License Keys for a Standalone QFX Series Device on page 90](#)
- [Generating the License Keys for a QFabric System on page 91](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Deleting a License \(CLI Procedure\) on page 94](#)
- [Saving License Keys on page 95](#)
- [Verifying Junos OS License Installation on page 96](#)

Junos OS License Keys

Licensable Ports on MX5, MX10, and MX40 Routers

Starting with Junos OS Release 12.2, license keys are available to enhance the port capacity on MX5, MX10, and MX40 routers up to the port capacity of an MX80 router. The MX5, MX10, and MX40 routers are derived from the modular MX80 chassis with similar slot and port assignments, and provide all functionality available on an MX80 router, but at a lower capacity. Restricting port capacity is achieved by making a set of MIC slots and ports licensable. MICs without a license are locked, and are unlocked or made usable by installing appropriate upgrade licenses.

The base capacity of a router is identified by the Ideeprom assembly ID (I2C ID), which defines the board type. However, the Junos OS licensing infrastructure allows the use of

restricted ports without a license for a grace period of 30 days. After the grace period expires, the router reverts back to the base capacity if no upgrade license is purchased and installed for the locked ports. The I2C ID along with an upgrade license determine the final capacity of an MX5, MX10, or MX40 router.

The MX5, MX10, MX40, and MX80 routers support the following types of MICs:

- A built-in 10-Gigabit Ethernet MIC with four 10-Gigabit Ethernet ports
- Two front-pluggable MICs

A feature ID is assigned to every license upgrade for enhancing port capacity.

[Table 40 on page 89](#) displays the chassis types and their associated port capacity, I2C ID, base capacity, feature ID, feature name, and the final capacity after a license upgrade.

Table 40: Upgrade Licenses for Enhancing Port Capacity

Chassis Type	Port Capacity	I2C ID	Base Capacity	Feature ID and Feature Name	Upgrade Capacity
MX5	20G	0x556	Slot 1 <ul style="list-style-type: none"> • 1/MIC0 	f1—MX5 to MX10 upgrade	Slot 1 and 2 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1
MX10	40G	0x555	Slot 1 and 2 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1 	f2—MX10 to MX40 upgrade	Slot 2 and first 2 ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC1 • First 2 ports on 0/MIC0
MX40	60G	0x554	Slot 1, Slot 2 and first 2 ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1 • First 2 ports on 0/MIC0 	f3—MX40 to MX80 upgrade	Slot 2 and all ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC1 • All 4 ports on 0/MIC0

When installing an upgrade license for enhancing port capacity on MX5, MX10 and MX40 routers, consider the following:

- To upgrade an MX5 router to MX80 router capacity, licenses for all three features (f1, f2, f3) must be installed. All three features can be provided in a single license key.
- To upgrade an MX10 router to MX40 router capacity, installing a license key with f2 feature is sufficient.
- Non-applicable feature IDs in a license key reject the upgrade license. For example:
 - An f1 feature ID on an MX10 upgrade license key rejects the license.
 - Feature IDs f1 and f2 on an MX40 upgrade license key reject the entire license.

- Related Documentation**
- [Junos OS Feature Licenses on page 86](#)
 - *License Enforcement*
 - *Software Feature Licenses*
 - [Verifying Junos OS License Installation on page 96](#)

Generating the License Keys for a Standalone QFX Series Device

When you purchase a Junos OS software feature license for a QFX Series device, you receive an e-mail containing an authorization code for the feature license from Juniper Networks. You can use the authorization code to generate a unique license key (a combination of the authorization code and the device's serial number) for the device, and then add the license key on the device.

Before generating the license keys for a device:

- Purchase the required licenses for the device. See "[Software Features That Require Licenses on the QFX Series](#)" on page 87.
- Note down the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
- Determine the serial number of the device. For instructions, see *Locating the Serial Number on a QFX3500 Device or Component*.

To generate the license keys for a device:

1. In a browser, log in to the Juniper Networks License Management System at <https://www.juniper.net/lcrs/license.do>.

The Manage Product Licenses page appears.



NOTE: To access the licensing site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. On the Generate Licenses tab, select **QFX Series Product** from the drop-down list, and click **Go**.

The Generate Licenses - QFX Series Product page appears.

3. Select the **QFX Series Product Device** option button, and click **Continue**.

The Generate Licenses - QFX Series Product Devices page appears.

4. In the **Device Serial Number** field, enter the serial number for the device.
5. In the **Authorization Code** field, enter the authorization code in the e-mail you received from Juniper Networks when you purchased the license.

6. (Optional) If you want to enter another authorization code for the same device, click **Enter More Authorization Codes** to display a new authorization code field. Enter the authorization code in this field.

7. Click **Confirm**.

The Confirm License Information page appears, displaying a summary of the information you submitted to the License Management System.

8. Review the information to ensure everything is correct and then click **Generate License**.

The Generate Licenses - QFX Series Product Devices page appears, displaying a summary of your license keys, including a link that displays the details of your new license keys.

9. Select the file format in which you want to obtain your new license keys.
10. Select the delivery method you want to use to obtain your new license keys.

To download the license keys:

- Select the **Download to this computer** option button, and click **OK**.

To e-mail the license keys:

- Select the **Send e-mail to e-mail ID** option button, and click **OK**.

Related Documentation

- [Software Features That Require Licenses on the QFX Series on page 87](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Locating the Serial Number on a QFX3500 Device or Component](#)

Generating the License Keys for a QFabric System

When you purchase a Junos OS software feature license for a QFabric system, you receive an e-mail containing an authorization code for the feature license from Juniper Networks. You can use the authorization code to generate a unique license key (a combination of the authorization code and the QFabric system ID) for the QFabric system, and then add the license key on the QFabric system.

Before generating the license keys for a QFabric system:

- Purchase the required licenses for the QFabric system. See [“Software Features That Require Licenses on the QFX Series” on page 87](#).
- Note down the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
- Perform the initial setup of the QFabric system on the Director group. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#).
- Log in to the QFabric system, issue the `show version` command, and note down the software serial number and QFabric system ID for the QFabric system.

```
user@qfabric> show version
```

```
Hostname: qfabric
Model: qfx3000-g
Serial Number: qfsn-0123456789
QFabric System ID: f158527a-f99e-11e0-9fbd-00e081c57cda
JUNOS Base Version [12.2I20111018_0215_dc-builder]
```

To generate the license keys for a QFabric system:

1. In a browser, log in to the Juniper Networks License Management System at <https://www.juniper.net/lcrs/license.do>.

The Manage Product Licenses page appears.



NOTE: To access the licensing site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. On the Generate Licenses tab, select **QFX Series Product** from the drop-down list, and click **Go**.

The Generate Licenses - QFX Series Product page appears.

3. Select the **QFX Series Product Fabric** option button, and then click **Continue**.

The Generate Licenses - QFX Series Product Fabrics page appears.

4. In the **Software Serial No** field, enter the software serial number for the QFabric system.
5. In the **QFabric System ID** field, enter the QFabric system ID for the QFabric system.
6. In the **Authorization Code** field, enter the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
7. (Optional) If you want to enter another authorization code for the same device, click **Enter More Authorization Codes** to display a new authorization code field. Enter the authorization code in this field.

8. Click **Confirm**.

The Confirm License Information page appears, displaying a summary of the information you submitted to the License Management System.

9. Review the information to ensure everything is correct and then click **Generate License**.

The Generate Licenses - QFX Series Product Fabrics page appears, displaying a summary of your license keys, including a link that displays the details of your new license keys.

10. Select the file format in which you want to obtain your new license keys.
11. Select the delivery method you want to use to obtain your new license keys.

To download the license keys:

- Select the **Download to this computer** option button, and click **OK**.

To e-mail the license keys:

- Select the **Send e-mail to e-mail ID** option button, and click **OK**.

Related Documentation

- [Software Features That Require Licenses on the QFX Series on page 87](#)
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [show version on page 1135](#)

Adding New Licenses (CLI Procedure)

Before adding new licenses, complete the following tasks:

- Purchase the required licenses.
- Establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your device.



NOTE: On QFabric systems, install your licenses in the default partition of the QFabric system and not on the individual components (Node devices and Interconnect devices).

To add a new license key to the device using the CLI:

1. From the CLI operational mode, enter one of the following CLI commands:
 - To add a license key from a file or URL, enter the following command, specifying the filename or the URL where the key is located:


```
user@host> request system license add filename | url
```
 - To add a license key from the terminal, enter the following command:


```
user@host> request system license add terminal
```
2. When prompted, enter the license key, separating multiple license keys with a blank line.

If the license key you enter is invalid, an error appears in the CLI output when you press Ctrl+d to exit license entry mode.



NOTE: If the router has two Routing Engines, add the new license on each Routing Engine separately. This ensures that the license key is enabled on the backup Routing Engine during changeover of mastership between the Routing Engines.

To add a new license key to a router with dual Routing Engines:

1. After adding the new license key on the master Routing Engine, use the **request chassis routing-engine master switch** command to have the backup Routing Engine become the master Routing Engine.
2. Log in to the active Routing Engine and add the new license key, repeating the same process.

3. Go on to [“Verifying Junos OS License Installation” on page 96](#).



NOTE: Adding a license key to the router or switch might be delayed if a kernel resynchronization operation is in progress at that time. The following message is displayed on the CLI when the license-adding operation is about to be delayed:

A kernel re-sync operation is in progress. License update may take several minutes to complete.

Related Documentation

- [Deleting a License \(CLI Procedure\) on page 94](#)
- [Junos OS Feature Licenses on page 86](#)
- [Verifying Junos OS License Installation on page 96](#)
- [request system license add on page 381](#)

Deleting a License (CLI Procedure)

Before deleting a license, establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your router or switch.

You have the options to delete a single license, delete all licenses, or delete a list of licenses enclosed in brackets.

To delete a license key from a device using the CLI:

1. Display the licenses available to be deleted.

```
user@host> request system license delete license-identifier-list ?
```

Possible completions:

E00468XXX4	License key identifier
JUNOS10XXX1	License key identifier
JUNOS10XXX2	License key identifier
JUNOS10XXX3	License key identifier

```
JUNOS10XXX4      License key identifier
[                Open a set of values
```

2. To delete a license key or keys from a device using the CLI operational mode, select one of the following methods:

- Delete a single license by specifying the license ID. Using this option, you can delete only one license at a time.

```
user@host> request system license delete license-identifier
```

- Delete all license keys from the current device.

```
user@host> request system license delete all
```

- Delete multiple license keys from the current device. Specify the license identifier for each key and enclose the list of identifiers in brackets.

```
user@host> request system license delete license-identifier-list [JUNOS10XXX1
JUNOS10XXX3 JUNOS10XXX4 ...]
```

```
Delete license(s) ?
[yes,no] (no) yes
```

3. Go on to “[Verifying Junos OS License Installation](#)” on page 96.



NOTE: Deleting a license key from the router or switch might be delayed if a kernel resynchronization operation is in progress at that time. The following message is displayed on the CLI when the license-deleting operation is about to be delayed:

A kernel re-sync operation is in progress. License update may take several minutes to complete.

Related Documentation

- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Saving License Keys on page 95](#)
- [Junos OS Feature Licenses on page 86](#)
- [Verifying Junos OS License Installation on page 96](#)
- [request system license delete on page 382](#)

Saving License Keys

Before saving a license, establish basic network connectivity with the router or switch. For instructions on establishing basic connectivity, see the *Getting Started Guide* or *Quick Start Guide* for your router or switch.

To save the licenses installed on a device to a file using the CLI:

1. From the CLI operational mode, enter one of the following CLI commands:

- To save the installed license keys to a file or URL, enter the following command:

```
user@host> request system license save filename | url
```

For example, the following command saves the installed license keys to a file named **license.config**:

- To save a license key from the terminal, enter the following command:

```
user@host> request system license save ftp://user@host/license.config
```

2. Go on to [“Verifying Junos OS License Installation” on page 96](#).

Related Documentation

- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Deleting a License \(CLI Procedure\) on page 94](#)
- [Junos OS Feature Licenses on page 86](#)
- [Verifying Junos OS License Installation on page 96](#)

Verifying Junos OS License Installation

To verify Junos OS license management, perform the following tasks:

- [Displaying Installed Licenses on page 96](#)
- [Displaying License Usage on page 97](#)

Displaying Installed Licenses

Purpose Verify that the expected licenses are installed and active on the router or switch.

Action From the CLI, enter the **show system license** command.

Sample Output

```
user@host> show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-acct	0	1	0	permanent
subscriber-auth	0	1	0	permanent
subscriber-addr	0	1	0	permanent
subscriber-vlan	0	1	0	permanent
subscriber-ip	0	1	0	permanent
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent

Licenses installed:

License identifier: E000185416

License version: 2

Features:

```
subscriber-acct - Per Subscriber Radius Accounting
permanent
subscriber-auth - Per Subscriber Radius Authentication
permanent
subscriber-addr - Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
```

```
subscriber-ip    - Dynamic and Static IP
permanent
```

Meaning The output shows a list of the license usage and a list of the licenses installed on the router or switch. Verify the following information:

- Each license is present. Licenses are listed in ascending alphanumeric order by license ID.
- The state of each license is **permanent**.



NOTE: A state of **invalid** indicates that the license key is not a valid license key. Either it was entered incorrectly or it is not valid for the specific device.

- The feature for each license is the expected feature. The features enabled are listed by license. An all-inclusive license has all features listed.
- All configured features have the required licenses installed. The Licenses needed column must show that no licenses are required.

Displaying License Usage

Purpose Verify that the licenses fully cover the feature configuration on the router or switch.

Action From the CLI, enter the **show system license usage** command.

Sample Output

```
user@host> show system license usage
```

	Licenses	Licenses	Licenses	Expiry
Feature name	used	installed	needed	
subscriber-addr	1	0	1	29 days
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent

Meaning The output shows any licenses installed on the router or switch and how they are used. Verify the following information:

- Any configured licenses appear in the output. The output lists features in ascending alphabetical order by license name. The number of licenses appears in the third column. Verify that you have installed the appropriate number of licenses.
- The number of licenses used matches the number of configured features. If a licensed feature is configured, the feature is considered used. The sample output shows that the subscriber address pooling feature is configured.
- A license is installed on the router or switch for each configured feature. For every feature configured that does not have a license, one license is needed.

For example, the sample output shows that the subscriber address feature is configured but that the license for the feature has not yet been installed. The license must be installed within the remaining grace period to be in compliance.

CHAPTER 5

Installation

- [Software Installation on page 99](#)

Software Installation

- [Junos OS Package Names on page 99](#)
- [Configuring Zero Touch Provisioning on page 101](#)
- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)
- [Recovering from a Failed Software Installation on page 119](#)
- [Software Installation Overview on page 120](#)
- [Upgrading Jloader Software on QFX Series Devices on page 120](#)
- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Upgrading Software on a QFabric System on page 134](#)

Junos OS Package Names

You upgrade the Juniper Networks Junos OS on the QFX Series by copying a software package to your switch or another system on your local network and then installing the new software package on the switch.

A software package name is in the following format:



NOTE: A signed domestic package is used as an example only. Other types of software packages might be available in future releases.

package-name-m.nZx.y-domestic-signed.tgz

where:

- ***package-name*** is the name of the package—for example, ***jinstall-qfx***.
- ***m.n*** is the software release, with ***m*** representing the major release number and ***n*** representing the minor release number—for example, ***11.1***.

- **Z** indicates the type of software release, where **R** indicates released software and **B** indicates beta-level software.
- **x.y** represents the maintenance software release, with **x** representing the maintenance software release number and **y** representing the maintenance software spin number—for example, **1.5**.

A sample switch software package name is:

`jinstall-qfx-11.1R1.5-domestic-signed.tgz`

**Related
Documentation**

- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Upgrading Software on a QFabric System on page 134](#)
- [Software Installation Overview on page 120](#)

Configuring Zero Touch Provisioning



NOTE: To see which platforms support Zero Touch Provisioning, in a browser, go to [Feature Explorer](#). In the Explore Features section of the Feature Explorer page, select All Features. In the Features Grouped by Feature Family box, select Zero Touch Provisioning. You can also type the name of the feature in the Search for Features edit box. In previous Junos OS releases on EX Series switches, Zero Touch Provisioning was called EZ Touchless Provisioning. Search for that feature name if you want to know if this feature is supported on EX Series switches.

Zero Touch Provisioning allows you to provision new switches in your network automatically, without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

The switch uses information that you configure on a Dynamic Host Control Protocol (DHCP) server to determine whether to perform these actions and to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration.



NOTE: If you have both DHCP and ZTP enabled, the switch broadcasts a DHCP DISCOVER packet every six minutes. If a DHCP server on the network responds with a DHCP ACK packet with DHCP vendor options set with the necessary values to initiate ZTP, then ZTP proceeds.

To disable broadcasting the DHCP DISCOVER packet every six minutes, without performing the ZTP process, manually delete the `auto-image-upgrade` statement located in the `[edit chassis]` hierarchy. If ZTP completes without errors, the `auto-image-upgrade` statement is automatically deleted.



NOTE: For detailed information regarding the DHCP and DHCP options, refer to RFC2131 (<http://www.ietf.org/rfc/rfc2131.txt>) and RFC2132 (www.ietf.org/rfc/rfc2132.txt). Also, this document refers to Internet Systems Consortium (ISC) DHCP version 4.2. For more information regarding this version, refer to <http://www.isc.org/software/dhcp/documentation>.

Before you begin:

- Ensure that the switch has access to the following network resources:

- The DHCP server provides the location of the software image and configuration files on the network

Refer to your DHCP server documentation for configuration instructions.

- The File Transfer Protocol (anonymous FTP), Hypertext Transfer Protocol (HTTP), Trivial File Transfer Protocol (TFTP) server on which the software image and configuration files are stored



NOTE: Although TFTP is supported, we recommend that you use FTP or HTTP instead, because these transport protocols are more reliable.

- A Domain Name System (DNS) server to perform reverse DNS lookup
 - (Optional) An NTP server to perform time synchronization on the network
 - (Optional) A system log (syslog) server to manage system log messages and alerts
- Locate and record the MAC address printed on the switch chassis.



CAUTION: You cannot commit a configuration while the switch is performing the software update process. If you commit a configuration while the switch is performing the configuration file autoinstallation process, the process stops, and the configuration file is not downloaded from the network.

To configure Zero Touch Provisioning for a switch:

1. Make sure the switch has the default factory configuration installed.

Issue the **request system zeroize** command on the switch that you want to provision.

2. Download the software image file and the configuration file to the FTP, HTTP, TFTP, server that the switch will download these files from.

You can download either one or both of these files.

3. Configure the DHCP server to provide the necessary information to the switch.

Configure IP address assignment.

You can configure dynamic or static IP address assignment for the switch's management address. To determine the switch's management MAC address for static IP address mapping, add 1 to the last byte of the switch's MAC address, which you noted before you began this procedure.

4. Define the format of the vendor-specific information for DHCP option 43 in the dhcpd.conf file.

Here is an example of an ISC DHCP 4.2 server dhcpd.conf file:

```
option space NEW_OP; option;  
option NEW_OP.config-file-name code 1 = text;  
option NEW_OP.image-file-type code 2 = text;  
option NEW_OP.transfer-mode code 3 = text;
```

```
option NEW_OP.alt-image-file-name code 4= text;
option NEW_OP-encapsulation code 43 = encapsulate NEW_OP;
```

5. Configure the following DHCP option 43 suboptions:

- Suboption 00: The name of the software image file to install



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

```
option NEW_OP.image-file-name
"/dist/images/jinstall-ex-4200-13.2R1.1-domestic-signed.tgz";
```

- Suboption 01: The name of the configuration file to install

```
option NEW_OP.config-file-name "/dist/config/jn-switch35.config";
```

- Suboption 02: The symbolic link to the software image file to install

```
option NEW_OP.image-file-type "symlink";
```



NOTE: If you do not specify suboption 2, the Zero Touch Provisioning process handles the software image as a filename, not a symbolic link.

- Suboption 03: The transfer mode that the switch uses to access the TFTP/FTP/HTTP server

```
option NEW_OP.transfer-mode "ftp";
```



NOTE: If suboption 03 is not configured, TFTP becomes the transfer mode by default.

- Suboption 04: The name of the software image file to install



NOTE: When the DHCP server cannot use suboption 00, configure the image file using suboption 04. If both suboption 00 and suboption 4 are defined, suboption 04 is ignored.

```
option NEW_OP.alt-image-file-name
"/dist/images/jinstall-ex-4200-13.2R1.1-domestic-signed.tgz";
```

6.



NOTE: You must configure either option 150 or option 66. If you configure both option 150 and option 66, option 150 takes precedence, and option 66 is ignored. Also, make sure you specify an IP address, not a hostname, because name resolution is not supported.

Configure DHCP option 150 to specify the IP address of the FTP, HTTP, or TFTP server.

```
option option-150 code 150 "10.100.31.71";
```

7. Configure DHCP option 66 to specify the IP address of the FTP, HTTP, or TFTP server.

```
option tftp-server-name "10.100.31.71";
```

8. (Optional) Configure DHCP option 7 to specify one or more system log (syslog) servers.

```
option log-servers 10.100.31.72;
```

9. (Optional) Configure DHCP option 42 to specify one or more NTP servers.

```
option ntp-servers 10.100.31.73;
```

10. (Optional) Configure DHCP option 12 to specify the hostname of the switch.

```
option hostname "jn-switch35";
```

The following sample configuration shows the DHCP options you just configured:

```
host jn-switch35 {  
  hardware ethernet ac:4b:c8:29:5d:02;  
  fixed-address 10.100.31.36;  
  option tftp-server-name "10.100.31.71";  
  option host-name "jn-switch35";  
  option log-servers 10.100.31.72;  
  option ntp-servers 10.100.31.73;  
  option NEW_OP.image-file-name  
    "/dist/images/jinstall-ex-4200-13.2R1.1-domestic-signed.tgz";  
  option NEW_OP.transfer-mode "ftp";  
  option NEW_OP.config-file-name "/dist/config/jn-switch35.config";  
}
```

Based on the DHCP options you just configured, the following statements are appended to the Junos OS configuration file (for example, **jn-switch35.config**):

```
system {  
  host-name jn-switch35;  
  syslog {  
    host 10.100.31.72 {  
      any any;  
    }  
  }  
  ntp {  
    server 10.100.31.73;  
  }  
}
```

11. Connect the switch to the network that includes the DHCP server and the FTP, HTTP, or TFTP,server.
12. Boot the switch with the default configuration.
13. Monitor the ZTP process by looking at the following log files.



NOTE: When SLAX scripts are issued, the **op-script.log** and **event-script.log** files are produced.

- /var/log/dhcp_logfile
- /var/log/image_load_log
- /var/log/op-script.log
- /var/log/event-script.log

Related Documentation

- [Understanding Zero Touch Provisioning on page 61](#)
- [NTP Time Server and Time Services Overview on page 52](#)
- [Op Script Overview](#)
- [Understanding DHCP Services for Switches on page 63](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 174](#)

Performing a Nonstop Software Upgrade on the QFabric System



NOTE: Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the `request system software add component all` command) does not require that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a nonstop software upgrade to Junos OS Release 12.2X50-D42.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device or one fabric control Routing Engine is always operational
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device

- Specifying an upgrade group if you want all Node devices in a Node group to be upgraded in parallel (which shortens the time of the upgrade)
- Rebooting devices automatically as part of the nonstop upgrade process

When performing a nonstop upgrade, start with the Director group upgrade, then issue the fabric upgrade, and end with the Node group upgrades.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade. For node-groups defined with two node-devices, both must be online in order for upgrade to succeed.



NOTE: Before you install the software, we recommend that you back up your current configuration files by issuing the `request system software configuration-backup` command.



NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the `request system software add component all` command.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 106](#)
- [Downloading Software Files Using a Browser on page 107](#)
- [Retrieving Software Files for Download on page 108](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group on page 108](#)
- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components on page 108](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups on page 109](#)
- [Performing a Nonstop Software Upgrade on a Node Group on page 110](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@qfabric> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup/media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
 2. Click **Download Software**.
 3. In the **Switching** box, click **Junos OS Platforms**.
 4. In the **QFX Series** section, click the name of the platform for which you want to download software.
 5. Click the **Software** tab and select the release number from the **Release** drop-down list.
 6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device” on page 111](#).
- A login screen appears.
7. Enter your user ID and password and press **Enter**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@qfabric> request system software download /path/package-name
```

For example:

```
user@qfabric> request system software download  
ftp://server/files/jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group



NOTE: If you reboot any Node groups or Interconnect devices after you perform a nonstop upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group  
jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To verify that both Director devices are online, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric  
jinstall-qfabric-12.2X50-D10.3.rpm
```


(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node group, create an upgrade group and include the devices you want to reboot at the same time.



NOTE:

- Upgrade groups work best when you build in redundancy so that not all the Node devices within the Node group restart at the same time. We suggest using the upgrade group feature for Node groups that have at least 2 pairs of Node devices (4 or more members), such as the network Node group.
- If you add Node devices that have links to the same link aggregation group (LAG), there might be traffic loss.

- Create the upgrade group by issuing the **set chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name* node-devices** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade1 node-devices
[ node1 node2 ]
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade2 node-devices
[ node3 node4 ]
```

Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.



NOTE: You can upgrade multiple Node groups with this command. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To install software on a Node group:

- Issue the **request system software nonstop-upgrade node-group *node-group-name* *package-name*** command.

To perform a nonstop upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1  
jinstall-qfabric-12.2X50-D10.3.rpm
```

To perform a nonstop upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1  
nodegroup2 nodegroup3] jinstall-qfabric-12.2X50-D10.3.rpm
```

Related Documentation

- [Configuring Graceful Restart for QFabric Systems on page 1375](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [show system software upgrade status on page 1570](#)

Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device for a QFX Series Device” on page 165](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



WARNING: The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the QFX Series device.
2. Reboot the QFX Series device.



NOTE: Do not power off the device if it is already on.

```
[edit system]
user@device> request system reboot
```

If you do not have access to the CLI, power cycle the QFX Series device.

The emergency boot device (external USB install media) is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following options:

```
External USB install media detected.
You can load Junos from this media onto an internal drive.
Press 'y' to proceed, 'f' to format and install, or 'n' to abort.
Do you wish to continue ([y]/f/n)? f
```

4. Type **f** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.

If you do not want to format the internal flash storage, type **y**.

The following messages are displayed:

```
Installing packages from external USB drive da1
Packages will be installed to da0, media size: 8G
```

```
Processing format options
Fri September  4 01:18:44 UTC 2012
```

```
-- IMPORTANT INFORMATION --
Installer has detected settings to format system boot media.
This operation will erase all data from your system.
```

```
Formatting installation disk .. this will take a while, please wait
Disabling platform watchdog - threshold 12 mins
```

```
Determining installation slice
Fri September  4 01:27:07 UTC 2012
```

5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the device is finished copying the software, you are presented with the following prompt:

```
*** Fri September  4 01:19:00 UTC 2012***
Installation successful..
Please select one of the following options:
Reboot to installed Junos after removing install media (default) ... 1
Reboot to installed Junos by disabling install media ..... 2
Exit to installer debug shell ..... 3
Install Junos to alternate slice ..... 4
Your choice: 4
NOTE: System installer will now install Junos to alternate slice
Do not power off or remove the external installer media or
interrupt the installation mechanism.
```

6. Select **4** to install Junos OS to the alternate slice of the partition, and then press Enter.
7. Remove the emergency boot device when prompted and then press Enter. The device then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the login prompt.
8. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.

**Related
Documentation**

- [Creating an Emergency Boot Device for a QFX Series Device on page 165](#)

Performing a QFabric System Recovery Installation on the Director Group

If the software on your QFabric system is damaged in some way that prevents the software from loading correctly, or you need to upgrade the software on your QFabric system, you may need to perform a recovery installation on the Director group.

If possible, perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device (for example, an external USB flash drive) for each of your Director devices to use during the recovery installation.

You can either use the external USB flash drive containing the software supplied by Juniper Networks, or you can use an external USB flash drive supplied by Juniper Networks on which you install the QFabric system install media.

2. Because the recovery installation process completely overwrites the entire contents of the Director device, make sure you back up any configuration files and initial setup information on a different external USB flash drive before you begin a recovery installation. You will need to restore this information as part of recovery process.

Use the **request system software configuration-backup** command to back up your configuration files and initial setup information:

```
user@switch> request system software configuration-backup path
```



NOTE: To recover the Director group, you must upgrade both Director devices in parallel. If you are recovering only one Director device in a Director group, and the software version will remain the same between the two Director devices, make sure that the other Director device is powered on and operational. If the software version of the Director device you are recovering will be different, make sure that the other Director device is powered off and is not operational.

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive on page 113
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software on page 115

(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive

If you do not have an external USB flash drive preloaded with the software from Juniper Networks to use as an emergency boot device, you can create your own, using a blank external USB flash drive provided by Juniper Networks. Download the install media from the Juniper Networks Support website onto your UNIX workstation, uncompress and untar the software, and then burn the software image onto your Juniper Networks external USB (4-gigabyte) flash drive. Make sure you create two emergency boot devices, one for each Director device, so you can perform a recovery installation in parallel.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.
4. In the **QFX Series** section, click the name of the platform for which you want to download software.

5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. Select the complete install media you want to download in the **QFabric System Install Media** section.
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Log in and save the install media file to your UNIX workstation.
10. Use FTP to access the UNIX workstation where the install media resides.

```
ftp ftp://hostname/pathname install-media-qfabric-<version>.img.tgz
```
11. When prompted, enter your username and password.
12. Make sure you are in binary mode by entering **binary** at the prompt.

```
binary
```
13. Use the **get** command to transfer the installation package from the FTP host to your UNIX workstation.

```
get install-media-qfabric-<version>.img.tgz
```
14. Close the FTP session:

```
bye
```
15. Untar the *install-media-qfabric-<version>.img.tgz* file on your UNIX workstation.

```
tar -xvzf install-media-qfabric-11.3X30.6.img.tgz
```
16. Insert a blank external USB (4-gigabyte) flash drive supplied by Juniper Networks into your UNIX workstation.
17. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive using the **dd** command:

```
dd if=install-media-qfabric-11.3X30.6.img of=/dev/sdb bs=16k
250880+0 records in
250880+0 records out
4110417920 bytes (4.1 GB) copied, 5.10768 seconds, 805 MB/s
```
18. Perform the steps in [“Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software” on page 115](#) to continue with the recovery installation.

Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software

This procedure describes how to perform a recovery installation using an external USB flash drive that contains Junos OS software.



NOTE: Since the recovery installation process completely overwrites the entire contents of the Director device, you will need to restore the required configuration files and initial setup information. The following procedure assumes you previously saved these backup files with the **request system software configuration-backup** command. Ensure that you have these backup files available on an external USB flash drive before you perform the following steps.

1. Insert the external USB flash drive into the Director device.
2. Perform one of the following tasks:
 - If you have access to the default partition, reboot the Director device by issuing the **request system reboot director-group** command.
 - If you do not have access to the default partition, power cycle the Director device.

The following menu appears on the Director device console when the Director device boots up:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

3. Type **install** and then press **Enter** to install the software on the Director device.

Once the installation process is complete, the Director device reboots, and the following menu appears on the Director device console:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

4. Press **Enter**.

The Director device reboots from the local disk on which the software was just installed.

5. Log in as root on the Director device.

The following menu appears on the Director device console:

```
Before you can access the QFabric system, you must complete the initial setup
of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of the
Director device and then log back in to restart
this setup menu.
```

```
Continue?[y/n]
```

6. Enter **n** to bypass the initial setup script and enter the Director device root directory, where you can mount the external USB flash drive containing the configuration files and initial setup information.

7. Issue the **ls /mnt** command to list the **mount** directory.

```
root@dg0 ~]# ls /mnt
```

8. Issue the **mkdir** command to create a directory within the mount directory.

```
root@dg0 ~]# mkdir /mnt/myusb
```

9. Issue the **mount /dev/sdb2 /mnt/myusb/** command to mount the external USB flash drive to the local drive of the Director device.

```
root@dg0 ~]# mount /dev/sdb2 /mnt/myusb/
```

10. Issue the **ls -la /mnt/myusb/** command to verify the contents of your mounted external USB flashdrive.

```
root@dg0 ~]# ls -la /mnt/myusb/
total 1770884
drwxr-xr-x 2 root root      4096 Sep  7 05:16 .
drwxr-xr-x 3 root root      4096 Sep  7 10:15 ..
-rw-r--r-- 1 root root    4249 Sep  7 03:52 mybackup-20110907
```

11. Exit the Director device and log back in as root on the Director device.

The following menu appears:

Before you can access the QFabric system, you must complete the initial setup of the Director group by using the steps that follow.

If the initial setup procedure does not complete successfully, log out of the Director device and then log back in to restart this setup menu.

```
Continue?[y/n] y
Initial Configuration
```

You may enter the configuration manually or restore from a backup.

```
Specify a backup file? [y/n] : y
Please specify the full path of the configuration backup file. :
/mnt/myusb/mybackup-20110907
```

12. Enter **y** to continue.

13. Enter **y** and specify the path to the backup configuration file located on the external USB flash drive.

```
/mnt/myusb/mybackup-20110907
```

The following messages appear:

```
Saving temporary configuration...
Configuring peer...
connect error for 1.1.1.2:9001
Configuring local interfaces...
Configuring interface eth0 with [10.49.213.163/24:10.49.213.254]
Configured interface eth0 with [10.49.213.163/24:10.49.213.254]
Configuring QFabric software with initial pool of 4000 MAC addresses
[00:10:00:00:00:00 - 00:10:00:00:0f:3b]
Configuring QFabric address [10.49.213.50]
Reconfiguring QFabric software static configuration
Applying the new Director Device password
Applying the QFabric component password
```


First install initial configuration, generating and sharing SSH keys.
 First install initial configuration, generating SSH keys.
 connect error for 1.1.1.2:9001
 Shared SSH keys.
 Configuration complete. Director Group services will auto start within 30 seconds.

The Director device reboots from the local disk on which the software was just installed.
 Exit the Director device session and log in to the QFabric default partition CLI.

14. Issue the **request system software configuration-restore** command and specify the path to the backup configuration file located on the external USB flash drive to load the previously saved QFabric system configuration.

15. From the default partition, issue the **request system reboot node-group all** command to reboot all of the Node groups in the QFabric system to ensure that all Node devices are running the same version of software as the Director-group.

```
user@switch> request system reboot node-group all
```

16. From the default partition, issue the **request system reboot fabric** command to reboot the Interconnect devices and the other components in the fabric in the QFabric system to ensure that Interconnect devices are running the same version of software as the Director group.

```
user@switch> request system reboot fabric
```

17. Log in to the default partition and issue the **show version component all** command to verify that all components are running the same version of software.

```
user@switch> show version component all
```

```
dg1:
```

```
-
```

```
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]
```

```
dg0:
```

```
-
```

```
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]
```

```
NW-NG-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

```
FC-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
```

JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-1:

Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

DRE-0:

-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FM-0:

-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

nodedevice1:

-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

interconnectdevice1:

-
Hostname: qfabric

```

Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
warning: from interconnectdevice0: Disconnected

```

- Related Documentation**
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
 - [Upgrading Software on a QFabric System on page 134](#)
 - [request system software configuration-backup on page 402](#)
 - [request system software configuration-restore on page 403](#)

Recovering from a Failed Software Installation

Problem **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

Solution If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

Hit [Enter] to boot immediately, or space bar for command prompt.

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [– –format] [– –external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:

- Network address of the server and the path on the server; for example, `tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz`
- Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, `file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz`.

The installation now proceeds normally and ends with a login prompt.

Software Installation Overview

A QFX Series product is delivered with the Junos OS preinstalled. As new features and software fixes become available, you can upgrade your software to use them.

When you power on the switch, it starts (boots) using the installed software.

You upgrade the Junos OS on a switch by copying a software package to a switch or other system on your local network and then using the CLI to install the new software on the switch. You then reboot the switch, which boots from the upgraded software. After a successful upgrade, you should back up the new current configuration to a secondary device.

During a successful upgrade, the installation package removes all files from the `/var/tmp` directory of the switch and completely reinstalls the existing software. It retains configuration files, and similar information, such as secure shell and host keys, from the previous version. The previous software package is preserved in a separate disk partition, and you can manually revert to it if necessary. If the software installation fails for any reason, such as loss of power during the installation process, the system returns to the originally active installation when you reboot.

If you encounter any difficulties during software installation or an upgrade, you can use the recovery installation procedure to install the Junos OS on the switch.

Related Documentation

- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Upgrading Software on a QFabric System on page 134](#)
- [Recovering from a Failed Software Installation on page 119](#)
- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111](#)

Upgrading Jloader Software on QFX Series Devices

Jloader software contains a boot loader (Uboot), which is used to bring up QFX Series devices and load the Junos OS from the flash memory of these devices. You can upgrade Jloader software on QFX3500 switches, QFX3500 and QFX3600 Node devices, and QFX3600-I and QFX3008-I Interconnect devices.



NOTE: Before you upgrade the Jloader software, see [Table 41 on page 121](#), [Table 42 on page 121](#), and [Table 43 on page 122](#) to make sure that you are upgrading to the right version of Jloader software for the Junos OS software release running on your QFX3500 and QFX3600 switches, or Node devices and Interconnect devices in your QFabric system.

See [Table 44 on page 122](#) to see which Uboot software versions are available and the filenames of the Jloader software packages.

Table 41: Junos OS and Jloader Software Compatibility Matrix for the QFX3500 Switch and QFX3500 Node Device

Junos OS Software Version	Jloader Software Version				
	1.0.73	1.1.1	1.1.2	1.1.4	1.1.5
11.3R1 and later (QFX3500 switch)	Not supported	Supported	Supported	Supported and recommended	Not supported
11.3X30.6 and later (QFX3500 Node device)	Not supported	Supported	Supported	Supported and recommended	Not supported
12.1X49-D1 and later (QFX3500 switch)	Not supported	Supported	Supported	Supported and recommended	Not supported
12.2X50-D1 and later (QFX3500 switch and QFX3500 Node device)	Not supported	Supported	Supported	Supported and recommended	Not supported



NOTE: An en dash means that the item is not applicable.

Table 42: Junos OS and Jloader Software Compatibility Matrix for the QFX3008-I Interconnect Device

Junos OS Software Version	Jloader Software Version				
	1.0.73	1.1.1	1.1.2	1.1.4	1.1.5
11.3X30.9 and later (QFX3008-I Interconnect device)	-	Supported	Supported	Supported and recommended	Not supported
11.3X30.6 and later (QFX3008-I Interconnect device)	-	Supported	Supported	Supported and recommended	Not supported
12.2X50-D1 and later (QFX3008-I Interconnect device)	-	Supported	Supported	Supported and recommended	Not supported



NOTE: An en dash means that the item is not applicable.

Table 43: Junos OS and Jloader Software Compatibility Matrix for the QFX3600-I Interconnect Device and QFX3600 Node Device

Junos OS Software Version	Jloader Software Version				
	1.0.73	1.1.1	1.1.2	1.1.4	1.1.5
12.2X50-D1 and later (QFX3600-I Interconnect Device and QFX3600 Node Device)	-	-	-	-	Supported
12.2X50-D20 and later (QFX3600 switch)	-	-	-	-	Supported

Table 44: Uboot Software Release and Jloader Software Compatibility Matrix

Uboot Software Release Number	Jloader Package Name
1.0.73	jloader-qfx-1_0_73.tgz
1.1.1	jloader-qfx-11.3-20110510.0-signed.tgz
1.1.2	jloader-qfx-11.3X30.9-signed.tgz
1.1.4 (11.3R3 and 11.3R2 releases only. Not supported on 11.3R1)	jloader-qfx-11.3I20120127_0733_dc-builder-signed.tgz
1.1.4 (12.1R1 release and later)	jloader-qfx-12.1-20120125_pr.0-signed.tgz
1.1.5 (12.2X50-D1 and later)	jloader-qfx-12.2X50.D10.3-signed.tgz
1.1.8 (13.1X50-D15 and later)	jloader-qfx-13.3-20130831_pr_branch_qfd.0.tgz

Jloader Software Version 1.1.4 Guidelines

Jloader software version 1.1.4 is compatible with Junos OS Release 11.3R3 and 11.3R2, and Junos OS Release 12.1R1 and later. Jloader software version 1.1.4 is not compatible with Junos OS

Release 11.3R1. The Jloader package names are different for versions 1.1.4 (Junos OS 11.3R3 and 11.3R2) and 1.1.4 (Junos OS 12.2R1 release and later), but the binaries are the same. Because the binaries are the same, you can upgrade or downgrade to any Junos OS release.

- If you have Junos OS Release 11.3 installed and want to upgrade the Jloader software from version 1.1.2 to version 1.1.4, you need to upgrade using the **jloader-qfx-11.3I20120127_0733_dc-builder-signed.tgz** package.
- If you have Junos OS Release 11.3R2 installed and want to upgrade to Junos OS Release 12.1, you do not need to upgrade the Jloader software version and can continue to use Jloader version 1.1.2.

- If you have Junos OS Release 12.1 installed and want to upgrade the Jloader software from version 1.1.2 to version 1.1.4, you need to upgrade using the **jloader-qfx-12.1-20120125_pr.0-signed.tgz** package.
- If you upgrade to Junos OS Release 12.1, you can upgrade to Jloader version 1.1.4 using the **jloader-qfx-12.1-20120125_pr.0-signed.tgz** package.

Upgrading Jloader Software on a QFX3500 Switch

The Jloader software for a QFX3500 switch resides in two flash memory banks. At any time, one bank acts as the primary bank, and the QFX3500 switch boots from it. The other bank is the backup bank—if the QFX3500 switch cannot boot from the primary bank, it boots from the backup bank. When you upgrade the Jloader software, the upgraded software is installed in the backup bank, which then becomes the new primary bank. Thus the primary and backup banks alternate each time you upgrade the Jloader software, with the primary bank containing the most recently installed version of the software, and the backup bank containing the previous version. To upgrade the Jloader software on a QFX3500 switch, you must perform the upgrade twice: once for each bank. Each upgrade requires that you to reboot the QFX3500 switch.



NOTE: If you are running Junos OS Release 11.3R1 or Junos OS Release 11.3R2, you must use the **no-validate** option when you issue the **request system software add** command to upgrade the Jloader software. Otherwise, the installation will fail and you receive a configuration error. The **no-validate** option is not required for Junos OS Release 11.3R3 and later.



NOTE: After you upgrade the Jloader software on the first bank, the package is deleted after you reboot. Make sure that you have either downloaded the Jloader software to either a remote site or in a local directory on the switch, such as the **/var/tmp** directory on the QFX3500 device.

1. In a browser, go to <http://www.juniper.net/support/downloads/junos.html> .
The Junos Platforms Download Software page appears.
2. In the QFX Series section of the Junos Platforms Download Software download page, select the QFX Series platform software you want to download.
3. Select the number of the software version that you want to download in the Release: pull-down window to the right of the tabs on the Download Software page.
4. Select the Software tab and then select the install package you want to download in the Install Package section.
5. In the pop-up Alert box, click the link to the Product Support Notification (PSN) document.
6. Enter your name and password and press **Enter**.

7. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
8. Open or save the **jloader-qfx-version-signed.tgz** file either to a local system or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.
9. Log in to the QFX3500 switch and enter the shell. We recommend using a console connection.
10. Determine the version of the Jloader software package installed on the switch.

For example:

```
root@switch% ls
gres-tp krt_gencfg_filter.txt
jloader-qfx-11.3-20110510.0-signed.tgz
```

11. Determine the version of the Uboot software that is running in the bank:

For example:

```
root@switch% kenv | grep boot.version
boot.version="1.0.7"
```

12. Enter the CLI and install the Jloader package.

- To install a Jloader package that is located in the **/var/tmp** directory, issue the **request system software add /var/tmp/jloader-qfx-version.tgz no-validate** command:

For example:

```
user@switch> request system software add
/var/tmp/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```

You see the following messages during the installation:

```
Verified jloader-qfx-11.3-20110510.0.tgz signed by PackageProduction_11_3_0
Adding jloader-qfx...
Installation in progress, please wait...
Mounted jloader-qfx package on /dev/md8...
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3-20110510.0 signed by PackageProduction_11_3_0
Registering jloader-qfx as unsupported
```

```
Installation finished successfully.
Please reboot to activate the package
Saving package file in /var/sw/pkg/jloader-qfx-11.3-20110510.0-signed.tgz
...
Saving state for rollback ...
```

```
juniper@qfx3500>
```

- To install a Jloader package located on a remote server using FTP, issue the **request system software add /ftp://hostname/pathname/jloader-qfx-version-signed.tgz no-validate** command.

For example:

```
user@switch> request system software add
/ftp://hostname/pathname/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```


- To install a Jloader package located on a remote server using HTTP, issue the **request system software add /http://hostname/pathname/jloader-qfx-version-signed.tgz no-validate** command.

For example:

```
user@switch> request system software add  
/http://hostname/pathname/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```

13. When prompted, reboot the Control Board by issuing the **request system reboot** command.

For example:

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

14. Enter the shell and verify that the version of the Uboot software in the primary bank is the version you just installed.

For example:

```
root@switch% kenv | grep boot.version  
boot.version="1.1.1"
```

15. To install the Jloader software package on the current backup bank, repeat Step 10 through Step 14.

Upgrading Jloader Software on a QFabric System

This procedure explains how to upgrade the Jloader software on your Node devices and Interconnect devices. The example shows how to upgrade the Jloader software from version 1.1.1 to 1.1.2 on a Node device with the serial number BBAK1186.



NOTE: Before you upgrade the Jloader software, make sure you have the serial numbers of the Node devices, Interconnect devices, and Control Boards in the Interconnect devices you want to upgrade.

1. Issue the **show chassis hardware node-device ?** command to view the serial numbers of the Node devices.

For example:

```
user@qfabric> show chassis hardware node-device ?
<node-device>      Node device identifier
BBAK1186            Node device
BBAK3149            Node device
BBAK3177            Node device
BBAK8063            Node device
BBAK8799            Node device
P2443-C             Node device
P2515-C             Node device
P3708-C             Node device
P3885-C             Node device
P3916-C             Node device
node0               Node device
node1               Node device
node2               Node device
node3               Node device
node4               Node device
node5               Node device
node6               Node device
node7               Node device
node8               Node device
```

An example of a Node device serial number is BBAK1186.

2. Issue the **show chassis hardware interconnect-device ?** command to view the serial numbers of the Interconnect devices.

For example:

```
user@qfabric> show chassis hardware interconnect-device ?
Possible completions:
interconnect-device  Interconnect device identifier
IC-F1052             Interconnect device
IC-F3947             Interconnect device
```

The Interconnect device serial numbers are IC-F1052 and IC-F3947.

3. Issue the **show chassis hardware interconnect-device name** command to view the serial numbers of the Control Boards in the Interconnect device.

For example:

```
user@qfabric> show chassis hardware interconnect-device IC-F3947
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis	REV 10		F3947	QFXC08-3008
Midplane	REV 10	750-035835	F3947-C	QFX Midplane
CB 0 Board	REV 14	750-035855	ZJ9432	QFX Chassis Control
Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
CB 1 Board	REV 14	750-035855	ZJ9404	QFX Chassis Control

The Control Board serial numbers are ZJ9432 and ZJ9404.

4. Issue the **show chassis firmware node-device *name*** command to see which version of Uboot software you have installed on your Node device.

For example:

```
user@qfabric> show chassis firmware node-device BBAK1186
```

Part	Type	Version
node4	U-Boot	1.1.6 (May 10 2011 - 04:52:59) 1.1.1
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1

The Uboot software version is 1.1.1. The loader software version appears after the timestamp for U-Boot 1.1.6.

5. Issue the **show chassis firmware interconnect-device *name*** command to see which version of Uboot software you have installed on the Routing Engines located on the Control Boards of the Interconnect device.

For example:

```
user@qfabric> show chassis firmware interconnect-device IC-F3947
```

Part	Type	Version
Routing Engine 0	U-Boot	U-Boot 1.1.6 (Jan 27 2012 - 03:24:34) 1.1.4
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1	U-Boot	U-Boot 1.1.6 (Jan 27 2012 - 03:24:34) 1.1.4
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1

The Uboot software version is 1.1.4. The loader software version appears after the timestamp for U-Boot 1.1.6.

6. In a browser, go to <http://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

7. In the QFX Series section of the Junos Platforms Download Software download page, select the QFX Series platform software you want to download.
8. Select the number of the software version that you want to download in the Release: pull-down window to the right of the tabs on the Download Software page.
9. Select the **Software** tab and then select the install package you want to download in the Install Package section.

10. In the pop-up Alert box, click the link to the Product Support Notification (PSN) document.
11. Enter your username and password, and press **Enter**.
12. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
13. Open or save the **jloader-qfx-version-signed.tgz** file either to a local system or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.
14. Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download /path/package-name** command.

For example:

```
user@qfabric> request system software download
ftp://server/files/jloader-qfx-11.3X30.9-signed.tgz
```

15. Log in to the Director device as root and enter the shell to verify that you have downloaded the Jloader software package. We recommend using a console connection. The software package is copied from where you downloaded it and is placed locally on the QFabric system in the **/pbdata/packages** directory.

For example:

```
[root@dg0] # pwd
/pbdata/packages

[root@dg0] # ls
jloader-qfx-11.3X30.9-signed.tgz
```

16. Before you copy over the Jloader software package to the Node device or Interconnect device, determine the directory that matches the serial number of the Node device or Interconnect device that you want to upgrade. View the remote logs and the Node device and Interconnect device serial numbers by issuing the **ls /pbdata/export/rlogs** command at the command line of the Director device before you copy the package over to the device.



NOTE: The **/pbdata/export/rlogs/node-device-serial-ID** and **/pbdata/export/rlogs/interconnect-device-serial-ID** directories on the Director device are NFS mounted as the **/tftpboot/logfiles** directories on the Node device and Interconnect device. These directories are created for all Node devices and Interconnect devices in a QFabric system. The Jloader files are stored in the **/tftpboot/logfiles** directories for each Node device and Interconnect device.

For example:

```
[root@dg0 tmp] # ls /pbdata/export/rlogs
02de4930-828b-11e1-a319-00e081c57938 c9898afe-828b-11e1-956c-00e081c57938
04103b2a-29d5-e011-bf8a-0e6bdf3aa1e6 eeba4aac-828b-11e1-85e2-00e081c57938
1e2739e0-828b-11e1-bf74-00e081c57938 F1052
8d8a978c-828b-11e1-a833-00e081c57938 F3947
ad55b89e-828b-11e1-b70e-00e081c57938 P2443-C
BBAK1186                             P2515-C
```

BBAK3149	P3708-C
BBAK3177	P3885-C
BBAK8063	P3916-C
BBAK8799	

BBAK1186 is the serial number of the Node device that needs to be upgraded.

17. Copy the Jloader software package from the `/var/tmp` directory to the `/pbdata/export/rlogs/BBAK1186` directory.

For example:

```
[root@dg0 tmp] # cp jloader-qfx-11.3X30.9-signed.tgz /pbdata/export/rlogs/BBAK1186
```

18. Confirm that the Jloader software package you copied over is in the `/pbdata/export/rlogs/BBAK1186` directory.

For example:

```
[root@dg0 tmp] # ls /pbdata/export/rlogs/BBAK1186
jloader-qfx-11.3X30.9-signed.tgz
```

19. Issue the `/root/dns.dump` command to find out the internal IP addresses of the Node device or Interconnect device.

```
[root@dg0 tmp] # /root/dns.dump
; <<>> DiG 9.3.6-P1-RedHat-9.3.6-4.P1.e15 <<>> -t axfr pkg.dcbg.juniper.net
@169.254.0.1
;; global options: printcmd
pkg.dcbg.juniper.net. 600 IN SOA ns.pkg.dcbg.juniper.net.
mail.pkg.dcbg.juniper.net. 152 3600 600 7200 3600
pkg.dcbg.juniper.net. 600 IN NS ns.pkg.dcbg.juniper.net.
pkg.dcbg.juniper.net. 600 IN A 169.254.0.1
pkg.dcbg.juniper.net. 600 IN MX 1 mail.pkg.dcbg.juniper.net.
dcfnode---DCF-ROOT.pkg.dcbg.juniper.net. 45 IN A 169.254.192.17
dcfnode---DRE-0.pkg.dcbg.juniper.net. 45 IN A 169.254.3.3
dcfnode-8d8a978c-828b-11e1-a833-00e081c57938.pkg.dcbg.juniper.net. 45 IN A
169.254.128.19
dcfnode-ad55b89e-828b-11e1-b70e-00e081c57938.pkg.dcbg.juniper.net. 45 IN A
169.254.128.20
dcfnode-BBAK1186.pkg.dcbg.juniper.net. 45 IN A 169.254.128.14
```

The internal IP address for BBAK1186 is 169.254.128.14.

20. Upgrade the Jloader software on the Node device or Interconnect device.

Before you can upgrade the Jloader software, you need to use SSH to log in to the Node device or Interconnect device and verify that the software is in the `/tftpboot/logfiles` directory.

- a. Use SSH to log in to the Node device or Interconnect device.

For example:

```
[root@dg0 tmp] # ssh 160.254.128.14
root@169.254.128.14's password:
--- JUNOS 11.3X30.10 built 2012-03-11 22:55:43 UTC
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@sng3%
```

- b. Verify that the Jloader software package is in the `tftpboot/logfiles` directory of the Node device or Interconnect device.

For example:

```
root@sng3% ls /tftpboot/logfiles
.index                               jloader-qfx-11.3X30.9-signed.tgz
```

- c. Copy the Jloader software package from the **/tftpboot/logfiles** directory to the **/var/tmp** directory of the Node device or Interconnect device.

For example:

```
root@sng3% cp /tftpboot/logfiles/jloader-qfx-11.3X30.9-signed.tgz /var/tmp
```

- d. Verify that the Jloader software package is in the **/var/tmp** directory of the Node device or Interconnect device.

For example:

```
root@sng3% ls /var/tmp
.snap                               jloader-qfx-11.3X30.9-signed.tgz
    tmp
gres-tp                             krt_gencfg_filter.txt
    vc-autoupgrade
if-rtbdb                             rtsdb
```

- e. Enter CLI mode and issue the **request system software add /var/tmp/jloader-qfx-version-signed.tgz** command.

For example:

```
root@sng3% cli
root@sng3> request system software add /var/tmp/jloader-qfx-11.3X30.9-signed.tgz
Validating on fpc0
Checking compatibility with configuration
Initializing...
Using jbase-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jbase-11.3X30.10 signed by PackageProduction_11_3_0
Using /var/tmp/jloader-qfx-11.3X30.9-signed.tgz
Verified jloader-qfx-11.3X30.9.tgz signed by PackageProduction_11_3_0
Using jloader-qfx-11.3X30.9.tgz
Checking jloader-qfx requirements on /
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3X30.9 signed by PackageProduction_11_3_0
Using jkernel-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jkernel-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jroute-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jroute-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jcrypto-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jcrypto-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jweb-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jweb-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jswitch-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jswitch-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
```

Done with validate on all chassis

```
fpc0:
Verified jloader-qfx-11.3X30.9.tgz signed by PackageProduction_11_3_0
Adding jloader-qfx...
Installation in progress, please wait...
Mounted jloader-qfx package on /dev/md10...
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3X30.9 signed by PackageProduction_11_3_0
#####
#####
Installation finished successfully.
Please reboot to activate the package
Saving package file in /var/sw/pkg/jloader-qfx-11.3X30.9-signed.tgz ...
Saving state for rollback ...
```

Upgrade has completed successfully.
Reboot is now required.

- f. Reboot both the Node device and Interconnect device twice, because they each contain two partitions.

For example:

```
root@sng3> request system reboot
Reboot the system ? [yes,no] (no) yes
Shutdown NOW!
[pid 37663]
```

```
root@sng3>
```

```
*** FINAL System shutdown message from root@sng3 ***
```

```
System going down IMMEDIATELY
```

- g. Verify that the Uboot software on the Node device or Interconnect device has been upgraded to the new Uboot software by logging in to the QFabric CLI and issuing either the **show chassis firmware node-device *name*** command or the **show chassis firmware interconnect-device *name*** command.

For example:

```
user@qfabric> show chassis firmware node-device BBAK1186
Part          Type      Version
node4         U-Boot    1.1.6 (Nov 19 2011 - 11:42:07) 1.1.2
                                loader    FreeBSD/MIPS U-Boot bootstrap loader
0.1
```

The Uboot software version is now 1.1.2. The loader software version appears after the timestamp for U-Boot 1.1.6.

Upgrading Software on QFX3500 and QFX3600 Standalone Switches

To upgrade Junos OS, you need to install the appropriate upgrade package on the QFX Series. Upgrading involves these tasks:

1. [Downloading Software Files with a Browser on page 132](#)
2. [Accessing Software Downloaded to a Remote Location on page 133](#)
3. [Connecting to the Console Port on page 133](#)
4. [Backing Up the Current Configuration Files on page 133](#)
5. [Installing the Software Package on page 133](#)

Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website, go to <http://www.juniper.net/support/>.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.
4. In the **QFX Series** section, click the name of the platform for which you want to download software.
5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. In the **Install Package** section of the **Software** tab, select the **Install Package** for the release.
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Save the **jinstall-qfx-<version>-domestic-signed.tgz** file on your computer.
10. Open or save the installation package either to the local system in the **var/tmp** directory or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.

Accessing Software Downloaded to a Remote Location

To access the installation package if you downloaded it to a remote location (for example, any system other than the switch):

1. From the command line, make sure you are in the `/var/tmp` directory of the switch.

2. Start the shell interface:

```
user@switch> start shell
```

3. Initiate an FTP, TFTP, or scp session.

In this example, FTP is used.

```
>ftp
```

4. Use FTP to access the remote location where the installation package resides.

```
ftp ftp://<hostname>/<pathname>/<package-name-m.mZx-distribution>.tgz.
```

where `<package-name-m.mZx-distribution>.tgz` is

```
jinstall-qfx-11.1R1.5-domestic-signed.tgz
```

5. When prompted, enter your username and password.

6. Use the **get** command to transfer the installation package from the remote location to your `/var/tmp` directory on your switch.

```
get <package-name-m.mZx-distribution>.tgz
```

7. Close the FTP session:

```
bye
```

Connecting to the Console Port

We recommend that you connect to the console port while installing the installation package so you can respond to any required user input and detect any errors that may occur.

Backing Up the Current Configuration Files

Before you install the new installation package, we strongly recommend that you back up your current configuration files because the upgrade process removes all of the stored files on the switch.

To back up your current configuration files, enter the **save** command:

```
user@switch> save filename
```

Executing this command saves a copy of your configuration files to a remote location such as an external USB device.

Installing the Software Package



NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <http://www.juniper.net/support>.

Install the software in one of two ways:

If the installation package resides locally on the switch, execute the **request system software add validate <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add validate  
/var/tmp/jinstall-qfx-11.1R1.5-domestic-signed.tgz reboot
```

If the Install Package resides remotely, execute the **request system software add validate <pathname> <source> reboot** command.

For example:

```
user@switch# request system software add validate  
ftp://ftpserver/directory/jinstall-qfx-11.1R1.5-domestic-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 53](#)
- [Software Installation Overview on page 120](#)
- [Recovering from a Failed Software Installation on page 119](#)
- [Upgrading Jloader Software on QFX Series Devices on page 120](#)
- [request system software add on page 394](#)
- [Junos® OS Installation and Upgrade Guide](#)

Upgrading Software on a QFabric System

The QFabric system software package contains software for all of the different components in the QFabric system, such as the Director group, Interconnect devices, Node devices, and other QFabric system components. You can upgrade the software on all of the QFabric components at the same time using the **request system software add package-name component all reboot** command.



NOTE: Downgrading software on a QFabric system is not supported.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 135](#)
- [Downloading Software Files Using a Browser on page 135](#)
- [Retrieving Software Files for Download on page 136](#)
- [Installing the Software Package on the Entire QFabric System on page 136](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@switch> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@switch> request system software configuration-backup /media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
 2. Click **Download Software**.
 3. In the **Switching** box, click **Junos OS Platforms**.
 4. In the **QFX Series** section, click the name of the platform for which you want to download software.
 5. Click the **Software** tab and select the release number from the **Release** drop-down list.
 6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device” on page 111](#).
- A login screen appears.
7. Enter your user ID and password and click **Login**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the `jinstall-qfabric-version.rpm` file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@switch> request system software download /path/package-name
```

For example:

```
user@switch> request system software download  
ftp://server/files/jinstall-qfabric-11.3X30.6.rpm
```

Installing the Software Package on the Entire QFabric System



NOTE: On a QFabric system, a QFX3500 Node device or QFX3600 Node device might not be able to participate as a Node device in the QFabric system if the Node device is running a different version of software from that of the Director group. This mismatch of software versions between the Node device and the Director group can occur when the Node device is introduced into the setup, and both Director devices go offline before the Node device completes its auto-upgrade process to upgrade its software version to the same software version running on the Director group. The workaround is to reboot the QFX3500 or QFX3600 Node device once the Director group comes back online. The QFX3500 or QFX3600 Node device will initiate auto-upgrade and upgrade its software version from the Director group.

1. Issue the **request system software add package-name component all reboot** command.

For example:

```
user@switch> request system software add jinstall-qfabric-11.3X30.6.rpm component all  
reboot
```



NOTE: If you receive an error message after issuing the **request system software add package-name component all reboot** command that says that the configuration file cannot be loaded as is, you will need to enter configuration mode, make any necessary changes to the configuration file, and then commit the changes.



NOTE: The default value for a QFabric system software upgrade is `validate`. The validation step adds up to 10 minutes to the overall software upgrade. If the validation fails, the upgrade does not proceed and the QFabric system automatically issues the request system software rollback command to restore the current software image. If you upgrade more than one component (for example, by issuing the `component all` option), validation failure on one device stops the upgrade process for the other devices. If you do not want to validate the software package against the current configuration, issue the `no-validate` option.

2. After the reboot has finished, verify that the new version of software has been properly installed by issuing the `show version component all` command.

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
```

```
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

DRE-0:

```
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

FM-0:

```
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

nodedevice1:

```
-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

interconnectdevice1:

```
-
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

- Related Documentation**
- [Software Installation Overview on page 120](#)
 - [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)
 - [Upgrading Jloader Software on QFX Series Devices on page 120](#)
 - [request system software add on page 394](#)
 - *Junos® OS Installation and Upgrade Guide*

CHAPTER 6

Configuration

- [Initial Configuration on page 141](#)
- [Configuration Examples on page 182](#)
- [Configuration Statements on page 190](#)

Initial Configuration

- [Configuring a DHCP Client \(CLI Procedure\) on page 142](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 143](#)
- [Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 146](#)
- [Configuring the Domain Name for the Router or Switch on page 146](#)
- [Configuring the Domains to Search When a Router or Switch Is Included in Multiple Domains on page 147](#)
- [Configuring the Hostname of the Router or Switch on page 147](#)
- [Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types on page 147](#)
- [Configuring the Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 148](#)
- [Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 148](#)
- [Configuring the Junos OS to Display a System Login Announcement on page 149](#)
- [Configuring the Junos OS to Display a System Login Message on page 149](#)
- [Configuring the Junos OS to Extend the Default Port Address Range on page 151](#)
- [Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 151](#)
- [Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 151](#)
- [Configuring NTP Authentication Keys on page 152](#)
- [Configuring NTP Authentication Keys \(QFabric System\) on page 153](#)
- [Configuring the NTP Time Server and Time Services on page 153](#)
- [Configuring the NTP Time Server and Time Services \(QFabric System\) on page 156](#)

- [Specifying the Physical Location of the Switch on page 156](#)
- [Configuring the Root Password on page 158](#)
- [Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 160](#)
- [Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 160](#)
- [Configuring System Alarms to Appear Automatically Upon Login on page 161](#)
- [Configuring Time-Based User Access on page 161](#)
- [Configuring the Timeout Value for Idle Login Sessions on page 162](#)
- [Configuring a QFX3500 Device as a Standalone Switch on page 163](#)
- [Creating an Emergency Boot Device for a QFX Series Device on page 165](#)
- [Creating a Snapshot and Using It to Boot a QFX Series Switch on page 166](#)
- [Including the Year or Millisecond in Timestamps on page 167](#)
- [Mapping the Hostname of the Switch to IP Addresses on page 168](#)
- [Methods for Configuring Junos OS on page 169](#)
- [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 172](#)
- [Rebooting and Halting a QFX Series Product on page 172](#)
- [Reverting to the Default Factory Configuration on page 173](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 174](#)
- [Reverting to the Rescue Configuration on page 175](#)
- [Saving Core Files Generated by Junos OS Processes on page 175](#)
- [Setting a Custom Time Zone on Routers or Switches Running Junos OS on page 175](#)
- [Setting the Date and Time on page 177](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 178](#)
- [Synchronizing and Coordinating Time Distribution Using NTP on page 180](#)
- [Viewing Core Files from Junos OS Processes on page 181](#)

Configuring a DHCP Client (CLI Procedure)

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

DHCP configuration consists of two components, configuration of DHCP clients and configuration of a DHCP server. Client configuration determines how clients send a message requesting an IP address, whereas a DHCP server configuration enables the server to send an IP address configuration back to the client. This topic describes configuring a DHCP client. For directions for configuring a DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 143](#).

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

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

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

The options that you can configure are listed in [Table 45 on page 143](#). Replace the variable *configuration-statement* with one or more of the statements listed in this table. If you do not explicitly configure these options, the switch uses default values for them.

Table 45: DHCP Client Settings

Configuration Statement	Description
client-identifier	Unique client ID—By default this consists of the hardware type (01 for Ethernet) and the MAC address (a.b.c.d). For this example, the value would be 01abcd.
lease-time	Ttime in seconds that a client holds the lease for an IP address assigned by a DHCP server. If a client does not request a specific lease time, then the server sends the default lease time. The default lease time on a Junos OS DHCP server is 1 day.
retransmission-attempt	Number of times the client attempts to retransmit a DHCP packet.
retransmission-interval	Time between transmission attempts.
server-address	IP address of the server that the client queries for an IP address.
update-server	TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch are propagated.
vendor-option	Vendor class ID (CPU's manufacturer ID string) for the DHCP client.

- Related Documentation**
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 143](#)
 - [Understanding DHCP Services for Switches on page 63](#)

Configuring a DHCP Server on Switches (CLI Procedure)

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

DHCP configuration consists of two components, optional reconfiguration of default settings on DHCP clients and configuration of a DHCP server. This topic covers configuration of the DHCP server. For directions for reconfiguring a DHCP client, see [“Configuring a DHCP Client \(CLI Procedure\)” on page 142](#).

You can configure either of two versions of a DHCP server on a switch—either the extended server version or the legacy server version. We recommend that you use the extended

server configuration unless you need to keep your DHCP server configuration backward-compatible with the legacy server version.

This topic includes the following tasks:

1. [Configuring an Extended DHCP Server on a Switch on page 144](#)
2. [Configuring a Legacy DHCP Server on a Switch \(CLI Procedure\) on page 145](#)

Configuring an Extended DHCP Server on a Switch

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

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

1. Create an address pool for DHCP IP addresses:

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

2. Configure addresses for DHCP dynamic assignment:

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

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

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

4. Optionally, process the information protocol data units (PDUs):

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

5. Optionally, redefine the order of attribute matching for pool selection:

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

6. Optionally, enable dynamic reconfiguration triggered by the DHCP extended server of all DHCP clients or only the DHCP clients serviced by the specified group of interfaces:

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

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

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

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

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

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



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

2. Optionally, change the domain search list used to resolve hostnames:

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

3. Optionally, change the domain name server (DNS) name that the DHCP server advertises to clients:

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

4. Optionally, change the DHCP options:

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

5. Optionally, change the devices advertised to clients:

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

6. Optionally, change the SIP server:

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

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

7. Optionally, change the DHCP client's hardware address:

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

8. Optionally, change the NetBIOS name server:

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

Related Documentation

- [Configuring a DHCP Client \(CLI Procedure\) on page 142](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\)](#)

- [Understanding DHCP Services for Switches on page 63](#)

Configuring a DNS Name Server for Resolving a Hostname into Addresses

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

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

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

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

Related Documentation

- [name-server on page 273](#)

Configuring the Domain Name for the Router or Switch

For each router or switch, you should configure the name of the domain in which the router or switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.

To configure the domain name, include the **domain-name** statement at the **[edit system]** hierarchy level:

```
[edit system]
domain-name domain-name;
```

The following example shows how to configure the domain name:

```
[edit]
user@host# set system domain-name company.net
[edit]
user@host# show
system {
  domain-name company.net;
}
```

Related Documentation

- [domain-name](#)

- [domain-name on page 257](#)
- [Example: Configuring the Domain Name for the Router or Switch on page 184](#)

Configuring the Domains to Search When a Router or Switch Is Included in Multiple Domains

If your router or switch is included in several different domains, you can configure those domain names to be searched.

To configure more than one domain to be searched, include the **domain-search** statement at the **[edit system]** hierarchy level:

```
[edit system]
domain-search [ domain-list ];
```

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

The following example shows how to configure two domains to be searched:

```
[edit system]
domain-search [ domainone.net domainonealternate.com ]
```

Related Documentation

- [Example: Configuring the Domain Name for the Router or Switch on page 184](#)
- [Configuring a DNS Name Server for Resolving a Hostname into Addresses](#)
- [Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 146](#)

Configuring the Hostname of the Router or Switch

To configure the name of the router or switch, include the **host-name** statement at the **[edit system]** hierarchy level:

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

The name value must be less than 256 characters.

Related Documentation

- [Example: Configuring the Name of the Router, IP Address, and System ID](#)
- [Example: Configuring the Name of the Switch, IP Address, and System ID on page 184](#)
- [Configuring Basic Router or Switch Properties](#)
- [Mapping the Hostname of the Switch to IP Addresses on page 168](#)
- [host-name on page 260](#)

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

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



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

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

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

alarm-name is the name of an alarm.

**Related
Documentation**

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

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

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

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

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

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

**Related
Documentation**

- *Configuring the Junos OS to Ignore ICMP Source Quench Messages*
- *Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 151*
- *Junos® OS Network Interfaces*

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

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

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

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

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

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

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

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

**Related
Documentation**

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

Configuring the Junos OS to Display a System Login Announcement

By default, no login announcement is displayed. To configure a system login announcement, include the **announcement** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
announcement text;
```

If the announcement text contains any spaces, enclose the text in quotation marks.

A system login *announcement* appears after the user logs in. A system login *message* appears before the user logs in.



TIP: You can use the same special characters described to format your system login announcement.

**Related
Documentation**

- *Defining Junos OS Login Classes*
- *Configuring the Junos OS to Display a System Login Message*

Configuring the Junos OS to Display a System Login Message

By default, no login message is displayed on the router or switch. To configure a system login message, include the **message** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
```

message text;

If the message text contains any spaces, enclose it in quotation marks.

You can format the message using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

The following is a sample login message configuration:

```
[edit]
system {
  login {
    message "\n\n\n\tUNAUTHORIZED USE OF THIS SYSTEM\n
\tIS STRICTLY PROHIBITED!\n\n\tPlease contact
\t'company-noc@company.com\t' to gain\naccess
\tto this equipment if you need authorization.\n\n\n";
  }
}
```

The preceding login message configuration example produces a login message similar to the following:

```
server% telnet router1
Trying 1.1.1.1...
Connected to router1.
Escape character is '^['.
```

```
UNAUTHORIZED USE OF THIS SYSTEM
IS STRICTLY PROHIBITED!
```

```
Please contact 'company-noc@company.com' to gain
access to this equipment if you need authorization.
```

```
router1 (ttyp0)
```

```
login:
```

A system login message appears before the user logs in. A system login announcement appears after the user logs in.

Related Documentation

- *Defining Junos OS Login Classes*
- [message on page 267](#)

Configuring the Junos OS to Extend the Default Port Address Range

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

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

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

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

Related Documentation

- *Configuring the Junos OS to Disable TCP RFC 1323 Extensions*
- *Configuring the Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- [source-port on page 293](#)

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

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

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

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

Related Documentation

- [icmpv4-rate-limit on page 260](#)

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

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

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

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

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

**Related
Documentation**

- [Configuring the Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 148](#)
- [default-address-selection on page 252](#)

Configuring NTP Authentication Keys

Time synchronization can be authenticated to ensure that the switch obtains its time services only from known sources. By default, network time synchronization is unauthenticated. The switch will synchronize to whatever system appears to have the most accurate time. We strongly encourage you to configure authentication of network time services.

To authenticate other time servers, include the **trusted-key** statement at the **[edit system ntp]** hierarchy level. Only time servers that transmit network time packets containing one of the specified key numbers are eligible to be synchronized. Additionally, the key needs to match the value configured for that key number. Other systems can synchronize to the local switch without being authenticated.

```
[edit system ntp]
trusted-key [ key-numbers ];
```

Each key can be any 32-bit unsigned integer except 0. Include the **key** option in the **peer**, **server**, or **broadcast** statements to transmit the specified authentication key when transmitting packets. The key is necessary if the remote system has authentication enabled so that it can synchronize to the local system.

To define the authentication keys, include the **authentication-key** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
```

number is the key number, **type** is the authentication type (only Message Digest 5 [MD5] is supported), and **password** is the password for this key. The key number, type, and password must match on all systems using that particular key for authentication.

**Related
Documentation**

- [NTP Time Server and Time Services Overview on page 52](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)
- [trusted-key on page 313](#)
- [authentication-key](#)

Configuring NTP Authentication Keys (QFabric System)

To configure the authentication keys using the CLI:

1. Configure the authentication-key number.

```
[edit system ntp]
user@switch# set authentication-key key-number
```

For example, to specify key 5:

```
user@switch# set authentication-key 5
```

2. Specify the type of authentication you want to use.

```
[edit system ntp]
user@switch# set authentication-key type type
```



NOTE: MD5 is the only authentication type supported.

For example, to specify MD5:

```
user@switch# set authentication-key type md5
```

Related Documentation

- [NTP Time Server and Time Services Overview \(QFabric System\) on page 53](#)
- [Configuring the NTP Time Server and Time Services \(QFabric System\) on page 156](#)
- [authentication-key on page 240](#)

Configuring the NTP Time Server and Time Services

When you use NTP, configure the router or switch to operate in one of the following modes:

- Client mode
- Symmetric active mode
- Broadcast mode
- Server mode

The following topics describe how to configure these modes of operation:

1. [Configuring the Router or Switch to Operate in Client Mode on page 153](#)
2. [Configuring the Router or Switch to Operate in Symmetric Active Mode on page 154](#)
3. [Configuring the Router or Switch to Operate in Broadcast Mode on page 154](#)
4. [Configuring the Router or Switch to Operate in Server Mode on page 155](#)

Configuring the Router or Switch to Operate in Client Mode

To configure the local router or switch to operate in client mode, include the **server** statement and other optional statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
```

```
server address <key key-number> <version value> <prefer>;  
authentication-key key-number type type value password;  
boot-server address;  
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement, as described in .

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in client mode:

```
[edit system ntp]  
authentication-key 1 type md5 value "$9$EgfcvX7VY4ZEcwgoHjkP5Q3CuREyv87";  
boot-server 10.1.1.1;  
server 10.1.1.1 key 1 prefer;  
trusted-key 1;
```

Configuring the Router or Switch to Operate in Symmetric Active Mode

To configure the local router or switch to operate in symmetric active mode, include the **peer** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
peer address <key key-number> <version value> <prefer>;
```

Specify the address of the remote system. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2 or 3, include the **version** option.

If you configure more than one remote system, you can mark one system preferred by including the **prefer** option:

```
peer address <key key-number> <version value> prefer;
```

Configuring the Router or Switch to Operate in Broadcast Mode

To configure the local router or switch to operate in broadcast mode, include the **broadcast** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
broadcast address <key key-number> <version value> <ttl value>;
```

Specify the broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be **224.0.1.1**.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2, or 3, include the **version** option.

Configuring the Router or Switch to Operate in Server Mode

In server mode, the router or switch acts as an NTP server for clients when the clients are configured appropriately. The only prerequisite for “server mode” is that the router or switch must be receiving time from another NTP peer or server. No other configuration is necessary on the router or switch.

To configure the local router or switch to operate as an NTP server, include the following statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
server address <key key-number> <version value> <prefer>;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, or 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in server mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$tXERuBEreWx-wtuLNdboaUjH.T3AtOESe";
server 172.17.27.46 prefer;
trusted-key 1;
```

Related Documentation

- [NTP Time Server and Time Services Overview on page 52](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)

Configuring the NTP Time Server and Time Services (QFabric System)

To configure the external time server using the CLI:

1. Configure the IP address of the external time server.

```
[edit system ntp]
user@switch# set server address
```

For example, to set an IP address of 10.1.1.1 for your external time server:

```
user@switch# set server 10.1.1.1
```

2. (Optional) Configure the key number to encrypt authentication fields in packets that are sent to the external time server.

```
[edit system ntp]
user@switch# set server address key key-number
```

For example, to set a key number of 1:

```
user@switch# set server address key 1
```

3. (Optional) Specify the external time server as a preferred host. Doing this enables the switch to synchronize with the external time server.



NOTE: The switch can synchronize with the external time server, but the external time server cannot synchronize with the switch.

```
[edit system ntp]
user@switch# set server address prefer
```

4. (Optional) Specify the NTP version number to be used in outgoing NTP packets.

```
user@switch# set server address version
```

For example, to specify version 3:

```
user@switch# set server address version 3
```

Related Documentation

- [NTP Time Server and Time Services Overview \(QFabric System\) on page 53](#)
- [ntp on page 277](#)
- [server on page 289](#)

Specifying the Physical Location of the Switch

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

- **altitude *feet***—Number of feet above sea level.
- **building *name***—Name of the building, 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").
- **country-code *code***—Two-letter country code.
- **floor *number***—Floor in the building.

- **hcoord** *horizontal-coordinate*—Bellcore Horizontal Coordinate.
- **lata** *service-area*—Long-distance service area.
- **latitude** *degrees*—Latitude in degree format.
- **longitude** *degrees*—Longitude in degree format.
- **npa-nxx** *number*—First six digits of the phone number (area code and exchange).
- **postal-code** *postal-code*—Postal code.
- **rack** *number*—Rack number.
- **vcoord** *vertical-coordinate*—Bellcore Vertical Coordinate.

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

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

**Related
Documentation**

Configuring the Root Password

The Junos OS is preinstalled on the router or switch. When the router or switch is powered on, it is ready to be configured. Initially, you log in as the user *root* with no password. The root directory of a UNIX device is the entry point to all other folders and files on that device. As a result, access to the root directory is restricted by default to a predefined user account known as *root user*. The root user (also referred to as *superuser*) has unrestricted access and full permissions within the system. The expression “log in as root” is commonly used when an action requires you to log into the device as the root user.



NOTE: If you configure a blank password using the `encrypted-password` statement at the `[edit system root-authentication]` hierarchy level for root authentication, you can commit a configuration, but you are *not* able to log in as the root user and gain root-level access to the router or switch.

After you log in, you should configure the root (superuser) password by including the **root-authentication** statement at the `[edit system]` hierarchy level and configuring one of the password options:

```
[edit system]
root-authentication {
  (encrypted-password "password" | load-key-password URL | plain-text-password);
  ssh-dsa "public-key" <from hostname>;
  ssh-ecdsa "public-key" <from hostname>;
  ssh-rsa "public-key" <from hostname>;
}
```



NOTE: ECDSA is not supported on the QFabric system.

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
- You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Valid passwords must contain at least one change of case or character class.

To enter a password that is already encrypted, use the following command to set the root password:

```
[edit system]
user@host# set root-authentication encrypted-password password
```

To view the encrypted password entries, use the configuration mode **show** command. For example:

```
[edit system]
```

```

user@host# show
root-authentication {
  "$1$ZUIES4dp$OUwWo1g7cLoV/aMWPUnC/",##SECRET-DATA
}

```

If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```

[edit system]
user@host# set root-authentication plain-text-password
New password: type password here
Retype new password: retype password here

```

To view the encrypted password entries, use the configuration mode **show** command. For example:

```

[edit system]
user@host# show
root-authentication {
  "$1$ZUIES4dp$OUwWo1g7cLoV/aMWPUnC/",##SECRET-DATA
}

```

You can use the **load-key-file** *URL filename* statement to load an SSH key file that was previously generated using **ssh-keygen**. The *URL filename* is the path to the file's location and name. When using this option, the contents of the key file are copied into the configuration immediately after entering the **load-key-file** *URL* statement. This command loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

Optionally, you can use the **ssh-dsa**, **ssh-ecdsa**, or **ssh-rsa** statements to directly configure SSH RSA, DSA, or ECDSA (not supported on the QFabric system) keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

To view the SSH keys entries, use the configuration mode **show** command. For example:

```

[edit system]
user@host# set root-authentication load-key-file
  scp://user@1.2.3.4/home/test/.ssh/id_dsa.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@host# show
root-authentication {
  ssh-dsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
SECRET-DATA
}

```

Junos-FIPS software has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the router or switch, you cannot

configure passwords unless they meet this standard. If you use the **encrypted-password** option, then a null-password (empty) is not permitted.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

**Related
Documentation**

- [Example: Configuring the Root Password on page 189](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 188](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1733](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182](#)
- [Recovering the Root Password](#)

Configuring the Router or Switch to Listen for Broadcast Messages Using NTP

When you are using NTP, you can configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet by including the **broadcast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
broadcast-client;
```

When the router or switch detects a broadcast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *broadcast client* mode, in which it listens for, and synchronizes to, succeeding broadcast messages.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

**Related
Documentation**

- [Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 160](#)
- [Configuring the NTP Time Server and Time Services on page 153](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)

Configuring the Router or Switch to Listen for Multicast Messages Using NTP

When you are using NTP, you can configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet by including the **multicast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
multicast-client <address>;
```

When the router or switch receives a multicast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *multicast client* mode, in which it listens for, and synchronizes to, succeeding multicast messages.

You can specify one or more IP addresses. (You must specify an address, not a hostname.) If you do, the router or switch joins those multicast groups. If you do not specify any addresses, the software uses **224.0.1.1**.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

- Related Documentation**
- [Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 160](#)
 - [Configuring the NTP Time Server and Time Services on page 153](#)
 - [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)

Configuring System Alarms to Appear Automatically Upon Login

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

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

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

- Related Documentation**
- [System Alarms on J Series Routers](#)
 - [show system alarms on page 914](#)

Configuring Time-Based User Access

The Junos OS enables you to configure time-based restrictions for user access to log in to a device. This is useful for restricting the time and duration of user logins for all users belonging to a login class. You can specify the days of the week when users can log in, the access start time, and the access end time.

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

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

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

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

```
    }  
  }
```

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

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

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



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

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

Related Documentation

- *Examples: Configuring Time-Based User Access*
- *Defining Junos OS Login Classes*
- *access-end*
- *access-start*
- *allowed-days*
- [access-end on page 232](#)
- [access-start on page 232](#)
- [allowed-days on page 235](#)

Configuring the Timeout Value for Idle Login Sessions

An idle login session is one in which the CLI operational mode prompt is displayed but there is no input from the keyboard. By default, a login session remains established until

a user logs out of the router or switch, even if that session is idle. To close idle sessions automatically, you must configure a time limit for each login class. If a session established by a user in that class remains idle for the configured time limit, the session automatically closes.

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

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

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

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

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

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

- Related Documentation**
- *Defining Junos OS Login Classes*
 - *idle-timeout*
 - [idle-timeout on page 261](#)

Configuring a QFX3500 Device as a Standalone Switch

If you are using the QFX3500 device as a standalone switch, you must perform the initial configuration of the QFX3500 device through the console port using the command-line interface (CLI). If you are using the QFX3500 as a Node device in a QFX3000 QFabric system, you instead perform the initial setup of a QFabric system on a QFX3100 Director device (see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#)).

Before you begin connecting and configuring a QFX3500 device, set the following parameter values on the console server or PC:

- Baud Rate—9600
- Flow Control—None
- Data—8
- Parity—None
- Stop Bits—1
- DCD State—Disregard

To connect and configure the device from the console:

1. Connect the console port to a laptop or PC using the supplied RJ-45 cable and RJ-45 to DB-9 adapter. The console (**CON**) port is located on the front panel of the device.
2. Log in as **root**. There is no password. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

Login: **root**

3. Start the CLI.

```
root@% cli
```

4. Enter configuration mode.

```
root> configure
```

5. Add a password to the root administration user account.

```
[edit]
```

```
root@# set system root-authentication plain-text-password
```

```
New password: password
```

```
Retype new password: password
```

6. (Optional) Configure the name of the device. If the name includes spaces, enclose the name in quotation marks (" ").

```
[edit]
```

```
root@# set system host-name host-name
```

7. Configure the default gateway.

```
[edit]
```

```
root@# set routing-options static route default next-hop address
```

8. Configure the IP address and prefix length for the device management interface.

```
[edit]
```

```
root@# set interfaces me0 unit 0 family inet address address/prefix-length
```



CAUTION: Configuring the two management Ethernet interfaces within the same subnet is not supported.



NOTE: The management ports are on the front panel of the QFX3500 device. They are labeled C0 and C1 on the front panel. In the CLI they are referred to as me0 and me1.

9. (Optional) Configure the static routes to remote prefixes with access to the management port.

```
[edit]
```

```
root@# set routing-options static route remote-prefix next-hop destination-ip retain no-readvertise
```

10. Enable telnet service.

```
[edit]
```

```
root@# set system services telnet
```




NOTE: When Telnet is enabled, you cannot log in to a QFX3500 device through Telnet using root credentials. Root login is allowed only for SSH access.

11. Commit the configuration to activate it on the device.

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

Related Documentation

- *Installing and Connecting a QFX3500 Device*
- *QFX3000-G QFabric System Installation Overview*
- *Understanding QFX3000-G QFabric System Hardware Configurations*

Creating an Emergency Boot Device for a QFX Series Device

If Junos OS on the QFX Series is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



NOTE: In the following procedure, we assume that you are creating the emergency boot device on a QFX3500 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device from a QFX3500 device:

1. Use FTP to copy the installation media image into the `/var/tmp` directory on the QFX3500 device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```



NOTE: The password is the root password for the QFX3500 device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

Related Documentation

- [USB Port Specifications for the QFX Series](#)
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)

Creating a Snapshot and Using It to Boot a QFX Series Switch

The system snapshot feature takes a “snapshot” of the files currently used to run the QFX Series switch—the complete contents of the `/config` and `/var` directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration—and copies all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source. You can then use these snapshots to boot the switch at the next bootup or as a backup boot option.

This topic includes the following tasks:

- [Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch on page 166](#)
- [Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch on page 167](#)

Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch

A snapshot can be created on USB flash memory after a switch is booted using files stored in internal memory.

Ensure that you have the following tools and parts available before creating a snapshot on a USB Flash drive:

- A USB flash drive that meets the QFX Series switch USB port specifications. See *USB Port Specifications for the QFX Series*.

To create a snapshot on USB flash memory and use it to boot the switch:

1. Place the snapshot into USB flash memory:

```
user@switch> request system snapshot partition
```



NOTE: This example uses the `partition` option. If you have already created a partition for the snapshot, you don't need to use the `partition` option.

2. (Optional) Perform this step if you want to boot the switch now using the snapshot stored on the external USB flash drive. If you created the snapshot as a backup, do not perform this step.

- To reboot the switch using the most recently created snapshot:

```
user@switch> request system reboot
```

- To reboot the switch using a snapshot in a specific partition on the USB flash drive:

```
user@switch> request system reboot slice 1
```

Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch

A snapshot can be created on internal memory after a switch is booted using files stored in external memory.

To create a snapshot in internal memory and use it to boot the switch:

1. Place the snapshot files in internal memory:

```
user@switch> request system snapshot partition
```



NOTE: This example uses the `partition` option. If you have already created a partition for the snapshot, you don't need to use the `partition` option.

2. (Optional) Perform this step if you want to boot the switch now using the newly created snapshot. If you created the snapshot as a backup, do not perform this step.

- To reboot the switch using the most recently created snapshot:

```
user@switch> request system reboot
```

- To reboot the switch using a snapshot in a specific partition in internal memory:

```
user@switch> request system reboot slice 1
```

Related Documentation

- *Verifying That a System Snapshot Was Created on a QFX Series Switch*
- *Understanding System Snapshot on QFX Series Switches*

Including the Year or Millisecond in Timestamps

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

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

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

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

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

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

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

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



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

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

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
 - [Examples: Configuring System Logging](#)

Mapping the Hostname of the Switch to IP Addresses

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

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

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

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

- Related Documentation**
- [Configuring Basic Router or Switch Properties](#)
 - [Configuring the Domain Name for the Router or Switch on page 146](#)

- *Example: Configuring the Name of the Router, IP Address, and System ID*
- [static-host-mapping on page 295](#)

Methods for Configuring Junos OS

You can use any of the methods shown in [Table 46 on page 169](#) to configure Junos OS:

Table 46: Methods for Configuring Junos OS

Method	Description
Command-line interface (CLI)	Create the configuration for the device using the CLI. You can enter commands from a single command line, and scroll through recently executed commands.
ASCII file	Load an ASCII file containing a configuration that you created earlier, either on this system or on another system. You can then activate and run the configuration file, or you can edit it using the CLI and then activate it.
J-Web graphical user interface (GUI)	Use the J-Web graphical user interface (GUI) to configure the device. J-Web enables you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser. The J-Web GUI is preinstalled on J Series Services Routers and is an optional software package that can be installed on M Series and T Series routers. J-Web is not available for the QFX Series.
Junos XML management protocol (API)	Use Junos XML protocol Perl client modules to develop custom applications for configuring information on devices that run Junos OS. Client applications use the Junos XML management protocol to request and change configuration information on Juniper Networks J Series, M Series, and T Series routers. The Junos XML management protocol is customized for Junos OS, and operations in the API are equivalent to those in the Junos OS CLI.
NETCONF application programming interface (API)	Use NETCONF Perl client modules to develop custom applications for configuring information on devices that run Junos OS. Client applications use the NETCONF XML management protocol to request and change configuration information on Juniper Networks J Series, M Series, and T Series routers. The NETCONF XML management protocol includes features that accommodate the configuration data models of multiple vendors.
Configuration commit scripts	Create scripts that run at commit time to enforce custom configuration rules. Commit scripts are written in Extensible Stylesheet Language Transformations (XSLT). Commit scripts are not available for the QFX Series.

The following sections contain complete descriptions of the methods you can use to configure Junos OS:

- [Junos OS Command-Line Interface \(CLI\) on page 170](#)
- [ASCII File on page 170](#)

- [J-Web Package on page 170](#)
- [Junos XML Management Protocol Software on page 171](#)
- [NETCONF XML Management Protocol Software on page 171](#)
- [Configuration Commit Scripts on page 171](#)

Junos OS Command-Line Interface (CLI)

The Junos OS CLI is a straightforward command interface. You use Emacs-style keyboard sequences to move around on a command line and scroll through a buffer that contains recently executed commands. You type commands on a single line, and the commands are executed when you press the Enter key. The CLI also provides command help and command completion. For more information about the CLI, see the *CLI User Guide* and [CLI Explorer](#).

ASCII File

You can load an ASCII file containing a configuration that you created earlier, either on this system or another system. You can then activate and run the configuration file as is, or you can edit it using the CLI and then activate it.

J-Web Package

As an alternative to entering CLI commands, the Junos OS supports the J-Web graphical user interface (GUI). The J-Web user interface enables you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser with Hypertext Transfer Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled.

The J-Web user interface is preinstalled on J Series Services Routers. It is provided as an optional, licensed software package (jweb package) on M Series and T Series routers. The jweb package is not included in jinstall and jbundle software bundles. It must be installed separately. To install the package on M Series and T Series routers, follow the procedure described in the *Junos® OS Installation and Upgrade Guide*.

J-Web supports weak (56-bit) encryption by default. This enables international customers to install J-Web and use HTTPS connections for J-Web access. Domestic customers can also install the jcrypto strong encryption package. This package automatically overrides the weak encryption. For more information about the J-Web GUI, see the *J-Web Interface User Guide*.



NOTE: Because the J-Web package is bundled separately from other packages, it is possible to have a version mismatch between J-Web and other Junos OS packages you have installed.

To check for a version mismatch, use the `show system alarms` CLI command. If the version number does not match exactly, a system alarm appears. For example, if you install the 7.4R1.2 jroute package and the 7.4R1.1 jweb package, an alarm is activated. For more information on the `show system alarms` command, see the [CLI Explorer](#).

Junos XML Management Protocol Software

The Junos XML management protocol is an Extensible Markup Language (XML) application that client applications use to request and change configuration information on Juniper Networks J Series, M Series, MX Series, and T Series routers. This API is customized for Junos OS, and operations in the API are equivalent to Junos OS CLI configuration mode commands. The Junos XML management protocol includes a set of Perl modules that enable client applications to communicate with a Junos XML protocol server on the router. The Perl modules are used to develop custom applications for configuring and monitoring Junos OS.

For a complete description of how to use Junos XML and Junos XML management protocol software, see the *Junos XML Management Protocol Developer Guide*.

NETCONF XML Management Protocol Software

The NETCONF XML management protocol is an Extensible Markup Language (XML) application that client applications can use to request and change configuration information on Juniper Networks J Series, M Series, MX Series, and T Series routers. This API is customized for Junos OS, and includes features that accommodate the configuration data models of multiple vendors. The NETCONF XML management protocol includes a set of Perl modules that enable client applications to communicate with a NETCONF server on the router. The Perl modules are used to develop custom applications for configuring and monitoring Junos OS.

For a complete description of how to use Junos XML and NETCONF XML management protocol software, see the *NETCONF XML Management Protocol Guide*.

Configuration Commit Scripts

You can create and use scripts that run at commit time to enforce custom configuration rules. If a configuration breaks the custom rules, the script can generate actions that the Junos OS performs. These actions include:

- Generating custom error messages
- Generating custom warning messages
- Generating custom system log messages
- Making changes to the configuration

Configuration commit scripts also enable you to create macros, which expand simplified custom aliases for frequently used configuration statements into standard Junos configuration statements. Commit scripts are written in Extensible Stylesheet Language Transformations (XSLT). For more information, see the *Junos OS Configuration and Operations Automation Guide*.

Related Documentation

- *Junos OS Configuration from External Devices*

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

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

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

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

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

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



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

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

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

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

Related Documentation

- [NTP Time Server and Time Services Overview on page 52](#)
- [Setting a Custom Time Zone on Routers or Switches Running Junos OS on page 175](#)

Rebooting and Halting a QFX Series Product

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

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>           Execute this command
```


at	Time at which to perform the operation
in	Number of minutes to delay before operation
media	Boot media for next boot
message	Message to display to all users
	Pipe through a command

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

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



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

```
user@switch> request system halt ?
Possible completions:
<[Enter]>      Execute this command
at             Time at which to perform the operation
in            Number of minutes to delay before operation
media         Boot media for next boot
message       Message to display to all users
|             Pipe through a command
```



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

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

Related Documentation

- [clear system reboot on page 332](#)
- [request system reboot on page 390](#)
- [request system halt on page 375](#)
- [request system power-off on page 385](#)
- [Connecting a QFX Series Device to a Management Console](#)

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1.

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

**Related
Documentation**

- [Understanding Configuration Files on page 1590](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Reverting to the Rescue Configuration on page 175](#)

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

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

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



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

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

1.

```
user@switch> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (yes)
```
2. Type **yes** to remove configuration and log files and revert to the factory default configuration.
3. Complete the initial configuration of the switch. See or [“Configuring a QFX3500 Device as a Standalone Switch” on page 163](#)

**Related
Documentation**

- [request system zeroize on page 433](#)

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1609](#)
- [Reverting to the Default Factory Configuration on page 173](#)
- [Configuration File Terms on page 48](#)

Saving Core Files Generated by Junos OS Processes

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

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

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

- To save the core files only:

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

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

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

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

Related Documentation

- [Viewing Core Files from Junos OS Processes on page 181](#)

Setting a Custom Time Zone on Routers or Switches Running Junos OS

You can update the time zone database information on routers or switches running Junos OS. This feature simplifies time zone management in devices running Junos OS by allowing

for future unforeseen time zone database adjustments. You can configure your router or switch to use a custom time zone database file that you create to meet your requirements by editing an existing time zone database file.

Tasks for setting a custom time zone are:

1. [Importing and Installing Time Zone Files on page 176](#)
2. [Configuring a Custom Time Zone on page 177](#)

Importing and Installing Time Zone Files

To import and install time zone files, follow these steps:

1. Download the time zone files archive and untar them to a temporary directory such as `/var/tmp`:

```
# mkdir -p /var/tmp/tz && cd /var/tmp/tz && rm *
# wget 'ftp://ftp.iana.org/tz/tzdata-latest.tar.gz'
# tar xvzf tzdata*.gz
africa
antarctica
asia
australasia
europe
northamerica
southamerica
pacificnew
etcetera
factory
backward
systemv
solar87
solar88
solar89
iso3166.tab
zone.tab
leapseconds
yearistype.sh
```



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

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

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

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

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

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

4. Set the time zone and commit the configuration:

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

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

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

Configuring a Custom Time Zone

To use a custom time zone, follow these steps:

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

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

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

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

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

Related Documentation

- *Modifying the Default Time Zone for a Router or Switch Running Junos OS*
- *NTP Overview*
- *NTP Time Server and Time Services Overview on page 52*
- *Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187*
- *use-imported-time-zones*

Setting the Date and Time

To set the date and time from a QFX Series product:



NOTE: The `set date` command is not operational on the QFabric system.

1. Enter operational mode in the CLI.
2. Enter the following command:

```
user@switch> set date YYYYMMDDHHMM.ss source-address
```

For example, the following command sets the date and time.

```
user@switch# set date 201102151010.55
```

3. To set the date and time from an NTP server, enter the following command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from an NTP server:

```
user@switch# set date ntp 200.40.40.1
```

4. To set the date and time from more than one NTP server, enter the same command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from more than one NTP server:

```
user@switch# set date ntp 200.40.40.1 200.40.40.2
```

Related Documentation

- *set date*

Specifying Access Privileges for Junos OS Operational Mode Commands

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

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

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

- Related Documentation**
- [Example: Configuring Access Privileges for Operational Mode Commands on page 1729](#)
 - [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1686](#)
 - *allow-commands*
 - *deny-commands*

Synchronizing and Coordinating Time Distribution Using NTP

Using NTP to synchronize and coordinate time distribution in a large network involves these tasks:

1. [Configuring NTP on page 180](#)
2. [Configuring the NTP Boot Server on page 180](#)
3. [Specifying a Source Address for an NTP Server on page 181](#)

Configuring NTP

- To configure NTP on the router or switch, include the **ntp** statement at the **[edit system]** hierarchy level:

```
[edit system]
ntp {
  authentication-key number type type value password;
  boot-server (address | hostname);
  broadcast <address> <key key-number> <version value> <ttd value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <version value> <prefer>;
  server address <key key-number> <version value> <prefer>;
  source-address source-address;
  trusted-key [ key-numbers ];
}
```

Configuring the NTP Boot Server

When you boot the router or switch, it issues an **ntpdate** request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP will not be able to synchronize to a time server if the server's time appears to be very far off of the local router's or switch's time.

- To configure the NTP boot server, include the **boot-server** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
boot-server (address | hostname);
```

Specify either the IP address or the hostname of the network server.

Specifying a Source Address for an NTP Server

For IP version 4 (IPv4), you can specify that if the NTP server configured at the **[edit system ntp]** hierarchy level is contacted on one of the loopback interface addresses, the reply always uses a specific source address. This is useful for controlling which source address NTP uses to access your network when it is either responding to or sending an NTP client request from your network.

To configure the specific source address that the reply will always use, and the source address that requests initiated by NTP server will use, include the **source-address** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.



NOTE: If a firewall filter is applied on the loopback interface, ensure that the source address specified for the NTP server at the **[edit system ntp]** hierarchy level is explicitly included as one of the match criteria in the firewall filter. This enables the Junos OS to accept traffic on the loopback interface from the specified source address.

The following example shows a firewall filter with the source address 10.0.10.100 specified in the **from** statement included at the **[edit firewall filter firewall-filter-name]** hierarchy:

```
[edit firewall filter Loopback-Interface-Firewall-Filter]
term Allow-NTP {
  from {
    source-address {
      172.17.27.46/32; // IP address of the NTP server
      10.0.10.100/32; // Source address specified for the NTP server
    }
    then accept;
  }
}
```

If no source address is configured for the NTP server, include the primary address of the loopback interface in the firewall filter.

Related Documentation

- [NTP Time Server and Time Services Overview on page 52](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)

Viewing Core Files from Junos OS Processes

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

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

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

Related Documentation

- [Saving Core Files from Junos OS Processes](#)
- [Saving Core Files Generated by Junos OS Processes on page 175](#)

Configuration Examples

- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182](#)
- [Example: Configuring the Domain Name for the Router or Switch on page 184](#)
- [Example: Configuring the Name of the Switch, IP Address, and System ID on page 184](#)
- [Example: Configuring NTP on page 185](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 187](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 188](#)
- [Example: Configuring the Root Password on page 189](#)

Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 182](#)
- [Overview on page 182](#)
- [Configuration on page 183](#)

Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**, **minimum-punctuations**, or **minimum-upper-cases**.

Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
set system login password maximum-length 22
set system login password minimum-numeric 1
set system login password minimum-upper-cases 1
set system login password minimum-lower-cases 1
set system login password minimum-punctuations 1
```

Configuring Requirements for Plain-Text Passwords

Step-by-Step Procedure This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the [system login password] hierarchy level.


```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.


```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.


```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.


```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

Related Documentation

- [Special Requirements for Junos OS Plain-Text Passwords on page 1687](#)
- [password \(Login\) on page 278](#)

Example: Configuring the Domain Name for the Router or Switch

The following example shows how to configure the router or switch domain name:

```
[edit]
user@host# set system domain-name company.net
[edit]
user@host# show
system {
    domain-name company.net;
}
```

Related Documentation

- [domain-name on page 257](#)
- [Configuring the Domain Name for the Router or Switch on page 146](#)

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

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

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

```

        alias sjl;
        sysid 1921.6800.1077;
    }
}

```

Related Documentation

- [Configuring Basic Router or Switch Properties](#)

Example: Configuring NTP

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. NTP uses a returnable-time design in which a distributed subnet of time servers operating in a self-organizing, hierarchical primary-secondary configuration synchronizes local clocks within the subnet and to national time standards by means of wire or radio. The servers also can redistribute reference time using local routing algorithms and time daemons.

This example shows how to configure NTP:

- [Requirements on page 185](#)
- [Overview on page 185](#)
- [Configuration on page 185](#)
- [Verification on page 186](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later
- A switch connected to a network on which an NTP boot server and NTP server reside

Overview

Debugging and troubleshooting are much easier when the timestamps in the log files of all switches are synchronized, because events that span a network can be correlated with synchronous entries in multiple logs. We recommend using the Network Time Protocol (NTP) to synchronize the system clocks of your switch and other network equipment.

In this example, an administrator wants to synchronize the time in a switch to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date are obtained when the router or switch boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme is used to hash the key value for authentication, which prevents the switch from synchronizing with an attacker's host that is posing as the time server.

Configuration

To configure NTP:

CLI Quick Configuration To quickly configure NTP, copy the following commands and paste them into the switch's terminal window:

```
[edit system]
set ntp boot-server 10.1.4.1
set ntp server 10.1.4.2
set ntp authentication-key 2 type md5 value "$9$aHlj8"
```

Step-by-Step Procedure To configure NTP :

1. Specify the boot server:

```
[edit system]
user@switch# set ntp boot-server 10.1.4.1
```
2. Specify the NTP server:

```
[edit system]
user@switch# set ntp server 10.1.4.2
```
3. Specify the key number, authentication type (MD5), and key for authentication:

```
[edit system]
user@switch# set ntp authentication-key 2 type md5 value "$9$aHlj8"
```

Results Check the results:

```
[edit system]
user@switch# show
ntp {
  boot-server 10.1.4.1;
  authentication-key 2 type md5 value "$9$aHlj8"; ## SECRET-DATA
  server 10.1.4.2;
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- [Checking the Time on page 186](#)
- [Displaying the NTP Peers on page 187](#)
- [Displaying the NTP Status on page 187](#)

Checking the Time

Purpose Check the time that has been set on the switch.

Action Enter the **show system uptime** operational mode command to display the time.

```
user@switch> show system uptime
fpc0:
-----
Current time: 2009-06-12 12:49:03 PDT
System booted: 2009-05-15 06:24:43 PDT (4w0d 06:24 ago)
Protocols started: 2009-05-15 06:27:08 PDT (4w0d 06:21 ago)
Last configured: 2009-05-27 14:57:03 PDT (2w1d 21:52 ago) by admin1
12:49PM up 28 days, 6:24, 1 user, load averages: 0.05, 0.06, 0.01
```

Meaning The output shows that the current date and time are June 12, 2009 and 12:49:03 PDT. The switch booted 4 weeks, 6 hours, and 24 minutes ago, and its protocols were started

approximately 3 minutes before it booted. The switch was last configured by user **admin1** on May 27, 2009, and there is currently one user logged in to the switch.

The output also shows that the load average is 0.05 seconds for the last minute, 0.06 seconds for the last 5 minutes, and 0.01 seconds for the last 15 minutes.

Displaying the NTP Peers

Purpose Verify that the time has been obtained from an NTP server.

Action Enter the **show ntp associations** operational mode command to display the NTP server from switch obtained its time.

```
user@switch> show ntp associations
      remote          refid          st t when poll reach  delay  offset  jitter
=====
*ntp5.domain1.ne .GPS.             1 u  414 1024  377   3.435   4.002   0.765
```

Meaning The asterisk (*) in front of the NTP server name, or peer, indicates that the time is synchronized and obtained from this server. The delay, offset, and jitter are displayed in milliseconds.

Displaying the NTP Status

Purpose View the configuration of the NTP server and the status of the system.

Action Enter the **show ntp status** operational mode command to view the status of the NTP.

```
user@switch> show ntp status
status=0644 leap_none, sync_ntp, 4 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Mon Apr 13 19:09:05 UTC 2009 (1)",
processor="powerpc", system="JUNOS9.5R1.8", leap=00, stratum=2,
precision=-18, rootdelay=2.805, rootdispersion=42.018, peer=48172,
refid=172.17.28.5,
reftime=cddd397a.60e6d7bf Fri, Jun 12 2009 13:30:50.378, poll=10,
clock=cddd3b1b.ec5a2bb4 Fri, Jun 12 2009 13:37:47.923, state=4,
offset=3.706, frequency=-23.018, jitter=1.818, stability=0.303
```

Meaning The output shows status information about the switch and the NTP.

Related Documentation

- [NTP Time Server and Time Services Overview on page 52](#)
- [ntp on page 277](#)
- [Configuring the NTP Time Server and Time Services on page 153](#)
- [CLI Explorer](#)
- [Junos OS Baseline Network Operations Guide](#)

Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization

Debugging and troubleshooting are much easier when the timestamps in the log files of all the routers or switches are synchronized, because events that span the network can

be correlated with synchronous entries in multiple logs. We strongly recommend using the Network Time Protocol (NTP) to synchronize the system clocks of routers, switches, and other network equipment.

By default, NTP operates in an entirely unauthenticated manner. If a malicious attempt to influence the accuracy of a router or switch's clock succeeds, it could have negative effects on system logging, make troubleshooting and intrusion detection more difficult, and impede other management functions.

The following sample configuration synchronizes all the routers or switches in the network to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date is obtained when the router boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme should be used to hash the key value for authentication, which prevents the router or switch from synchronizing with an attacker's host posing as the time server.

```
[edit]
system {
  ntp {
    authentication-key 2 type md5 value "$9$aHlj8gqQ1gijyjhgiiliii"; # SECRET-DATA
    boot-server 10.1.4.1;
    server 10.1.4.2;
  }
}
```

Related Documentation

- [NTP Overview](#)
- [NTP Time Server and Time Services Overview on page 52](#)
- [authentication-key](#)
- [boot-server on page 243](#)
- [server on page 289](#)
- [show ntp associations on page 892](#)
- [show ntp status on page 894](#)

Example: Configuring a Plain-Text Password for Root Logins

The following example shows how to set a plain-text password for root logins:

```
[edit]
user@switch# set system root-authentication plain-text-password
New password: type root password
Retype new password: retype root password
[edit]
user@switch# show
system {
  root-authentication {
    encrypted-password "$1$14c5.$sBopasddsdfs0";
  }
}
```


}

Related Documentation

- [Configuring the Root Password on page 158](#)

Example: Configuring the Root Password

Junos OS is preinstalled on the router. When the router is powered on, it is ready to be configured. Initially, you log in as the user “root” with no password. The root directory of a UNIX device is the entry point to all other folders and files on that device. As a result, access to the root directory is restricted by default to a predefined user account known as the *root user*. The root user (also referred to as *superuser*) has unrestricted access and full permissions within the system. The expression “log in as root” is commonly used when an action requires the user to log into the device as the root user.

To set the root password, you have several options: enter a clear-text password that the system will encrypt, enter a password that is already encrypted, or enter a secure shell (ssh) public key string.

To set the root password, follow these steps:

1. To enter a clear-text password that the system will encrypt, use the following command to set the root password:

```
[edit]
root@# set system root-authentication plain-text-password
New Password: type password here
Retype new password: retype password here
```

2. To enter a password that is already encrypted, use the following command to set the root password:

```
[edit]
root@# set system root-authentication encrypted-password password
```

3. To enter an ssh public string, use the following command to set the root password:

```
[edit]
root@# set system root-authentication ssh-rsa key
```

4. Commit the changes:

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

5. Exit from configuration mode:

```
[edit]
root@# exit
root@>
```

Related Documentation

- [Configuring the Root Password on page 158](#)

Configuration Statements

- [QFX Series CLI Hierarchy on page 192](#)
- [access-end on page 232](#)
- [access-start on page 232](#)
- [accounting on page 233](#)
- [accounting-port on page 234](#)
- [allow-commands on page 234](#)
- [allow-configuration on page 235](#)
- [allowed-days on page 235](#)
- [allow-transients on page 236](#)
- [announcement on page 236](#)
- [archival on page 237](#)
- [arp \(System\) on page 238](#)
- [authentication \(Login\) on page 239](#)
- [authentication-key on page 240](#)
- [authentication-order on page 241](#)
- [auxiliary on page 242](#)
- [boot-server \(NTP\) on page 243](#)
- [broadcast on page 244](#)
- [broadcast-client on page 245](#)
- [change-type on page 245](#)
- [checksum on page 246](#)
- [class \(Defining Login Classes\) on page 247](#)
- [class \(Assigning a Class to an Individual User\) on page 248](#)
- [commit on page 249](#)
- [compress-configuration-files \(System\) on page 250](#)
- [console \(Physical Port\) on page 251](#)
- [default-address-selection on page 252](#)
- [deny-commands on page 253](#)
- [deny-configuration on page 254](#)
- [destination \(Accounting\) on page 255](#)
- [destination-override on page 256](#)
- [direct-access on page 256](#)
- [domain-name on page 257](#)
- [domain-search on page 257](#)
- [explicit-priority on page 258](#)

- [events](#) on page 259
- [format](#) on page 259
- [host-name](#) on page 260
- [icmpv4-rate-limit](#) on page 260
- [idle-timeout](#) on page 261
- [internet-options](#) on page 261
- [load-key-file](#) on page 262
- [location](#) on page 263
- [login](#) on page 264
- [login-alarms](#) on page 265
- [login-tip](#) on page 265
- [max-configurations-on-flash](#) on page 266
- [maximum-length](#) on page 266
- [message](#) on page 267
- [minimum-changes](#) on page 267
- [minimum-length](#) on page 268
- [minimum-lower-cases](#) on page 269
- [minimum-numeric](#) on page 270
- [minimum-punctuations](#) on page 271
- [minimum-upper-cases](#) on page 272
- [multicast-client](#) on page 272
- [name-server](#) on page 273
- [no-multicast-echo](#) on page 274
- [no-ping-record-route](#) on page 275
- [no-ping-time-stamp](#) on page 275
- [no-redirects \(IPv4 Traffic\)](#) on page 276
- [ntp](#) on page 277
- [ntp \(QFabric\)](#) on page 277
- [optional](#) on page 278
- [password \(Login\)](#) on page 278
- [peer](#) on page 279
- [permissions](#) on page 280
- [port \(TACACS+ Server\)](#) on page 280
- [ports](#) on page 281
- [radius \(System\)](#) on page 282
- [refresh \(Commit Scripts\)](#) on page 283
- [refresh-from \(Commit Scripts\)](#) on page 283

- [retry](#) on page 284
- [retry-options](#) on page 285
- [root-authentication](#) on page 286
- [saved-core-context](#) on page 287
- [saved-core-files](#) on page 287
- [secret](#) on page 288
- [server \(NTP\)](#) on page 289
- [server \(RADIUS Accounting\)](#) on page 290
- [server \(TACACS+ Accounting\)](#) on page 290
- [single-connection](#) on page 291
- [source \(Commit Scripts\)](#) on page 291
- [source-address \(NTP, RADIUS, System Logging, or TACACS+\)](#) on page 292
- [source-port \(Port Addresses\)](#) on page 293
- [ssh-dsa](#) on page 293
- [ssh-rsa](#) on page 294
- [static-host-mapping](#) on page 295
- [structured-data](#) on page 296
- [syslog \(System\)](#) on page 297
- [system](#) on page 299
- [tacplus](#) on page 304
- [tacplus-server](#) on page 305
- [time-format](#) on page 306
- [timeout](#) on page 307
- [time-zone](#) on page 308
- [traceoptions \(Commit Scripts\)](#) on page 310
- [tracing](#) on page 312
- [trusted-key](#) on page 313
- [uid](#) on page 313
- [use-imported-time-zones](#) on page 314
- [user \(Access\)](#) on page 314

QFX Series CLI Hierarchy

This topic contains the full command-line interface (CLI) statement hierarchy for the QFX Series.

- [\[edit access\] Hierarchy](#) on page 193
- [\[edit accounting-options\] Hierarchy](#) on page 193
- [\[edit chassis\] Hierarchy](#) on page 195
- [\[edit class-of-service\] Hierarchy](#) on page 196

- [\[edit ethernet-switching-options\] Hierarchy on page 199](#)
- [\[edit fabric\] Hierarchy on page 200](#)
- [\[edit fc-fabrics\] Hierarchy on page 201](#)
- [\[edit fc-options\] Hierarchy on page 202](#)
- [\[edit firewall\] Hierarchy on page 202](#)
- [\[edit groups\] Hierarchy on page 203](#)
- [\[edit interfaces\] Hierarchy on page 203](#)
- [\[edit policy-options\] Hierarchy on page 209](#)
- [\[edit protocols\] Hierarchy on page 210](#)
- [\[edit security\] Hierarchy on page 222](#)
- [\[edit snmp\] Hierarchy on page 222](#)
- [\[edit system\] Hierarchy on page 226](#)
- [\[edit vlans\] Hierarchy on page 231](#)

[\[edit access\] Hierarchy](#)

```
access {
  address-assignment
  pool pool-name
  address-pool pool-name
  profile profile-name {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      (authentication-order (ldap radius | none);
      order (radius | none);
    }
    radius {
      accounting-server [server-addresses];
      authentication-server [server-addresses];
    }
  }
}
```

[\[edit accounting-options\] Hierarchy](#)

```
accounting-options {
  class-usage-profile profile-name {
    destination-classes {
      destination-class-name;
    }
    file filename;
    interval minutes;
    source-classes {
      source-class-name;
    }
  }
  file filename {
    archive-sites {
      site-name;
    }
  }
}
```

```
    }
    files number;
    nonpersistent;
    size bytes;
    start-time time;
    transfer-interval minutes;
  }
  filter-profile profile-name {
    counters {
      counter-name;
    }
    file filename;
    interval minutes;
  }
  interface-profile profile-name {
    fields {
      input-bytes;
      input-errors;
      input-multicast;
      input-packets;
      input-unicast;
      output-bytes;
      output-errors;
      output-multicast;
      output-packets;
      output-unicast;
      rpf-check-bytes;
      rpf-check-packets;
      rpf-check6-bytes;
      rpf-check6-packets;
      unsupported-protocol;
    }
    file filename;
    interval minutes;
  }
  mib-profile profile-name {
    file filename;
    interval minutes;
    object-names {
      mib-object-name;
    }
    operation (get | get-next | walk);
  }
  policy-decision-statistics-profile profile-name {
    application-aware-access-list-fields {
      address;
      application;
      application-group;
      input-bytes;
      input-interface;
      input-packets;
      mask;
      output-bytes;
      output-packets;
      subscriber-name;
      timestamp;
    }
  }
}
```

```

        vrf-name;
    }
    file filename;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}

```

[edit chassis] Hierarchy

```

interconnect-device {
    alarm {
        interface-type {
            link-down (red | yellow | ignore);
        }
    }
    container-devices {
        device-count number;
    }
    craft-lockout {
        alarm {
            interface-type {
                link-down (red | yellow | ignore);
            }
        }
        container-devices {
            device-count number;
        }
        fpc slot {
            power (on | off);
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
    fpc slot {
        power (on | off);
    }
    routing-engine {
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
    }
}
chassis {
    routing-engine {
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
    }
}

```

```
aggregated-devices {
  ethernet {
    device-count number;
  }
  alarm {
    interface-type {
      alarm-name (red | yellow | ignore);
    }
  }
}
fpc slot {
  pic pic-number {
    fibre-channel {
      port-range {
        port-range-low port-range-high;
      }
    }
  }
  xe {
    (port port-number | port-range port-range-low port-range-high);
  }
  xle {
    (port port-number | port-range port-range-low port-range-high);
  }
}
}
```

[\[edit class-of-service\] Hierarchy](#)

```
class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority level {
          code-points [ aliases ] [ bit-patterns ];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | ieee-802.1) {
      alias-name bits;
    }
  }
  congestion-notification-profile profile-name {
    input {
      ieee-802.1 {
        code-point [ code-point-bits ] {
          pfc {
            mru mru-value;
          }
        }
      }
    }
    cable-length cable-length-value;
  }
}
```



```

output {
  ieee-802.1 {
    code-point [code-point-bits] {
      flow-control-queue [queue | list-of-queues];
    }
  }
}
drop-profiles {
  profile-name {
    interpolate {
      fill-level low-value fill-level high-value drop-probability 0 drop-probability
        high-value;
    }
  }
}
forwarding-class class-name {
  loss-priority level {
    code-points [aliases] [bit-patterns];
  }
}
forwarding-class class-name {
  scheduler scheduler-name;
}
forwarding-class-sets forwarding-class-set-name {
  class class-name;
}
forwarding-classes {
  class {
    class-name {
      queue-num queue-number <no-loss>;
    }
  }
}
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point code-point;
}
interfaces {
  interface-name {
    congestion-notification-profile profile-name {
    }
    forwarding-class lossless-forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
      output-traffic-control-profile profile-name;
    }
    rewrite-value {
      input {
        ieee-802.1 {
          code-point code-point-bits;
        }
      }
    }
  }
  unit logical-unit-number {
    classifiers {
      (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
    }
  }
}

```

```

    }
    forwarding-class class-name;
    rewrite-rules {
        (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
    }
}
}
multi-destination {
    classifiers {
        (dscp | ieee-802.1) classifier-name;
    }
}
rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority priority code-point (alias | bits);
        }
    }
}
scheduler-map-forwarding-class-sets {
    fabric-scheduler-map-name {
        forwarding-class-set fabric-forwarding-class-set-name scheduler scheduler-name;
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-priority (low | medium-high | high) protocol protocol
            drop-profile drop-profile-name;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (percent percentage);
    }
}
shared-buffer {
    egress {
        percent percent;
        buffer-partition (lossless | lossy | multicast) {
            percent percent
        }
    }
    ingress {
        percent percent;
        buffer-partition (lossless | lossless-headroom | lossy) {
            percent percent
        }
    }
}
traffic-control-profiles profile-name {

```

```

    guaranteed-rate(rate| percent percentage);
    scheduler-map map-name;
    shaping-rate (rate| percent percentage);
  }
}

```

[edit ethernet-switching-options] Hierarchy

```

ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
      output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds {
  }
  port-error-disable {
    disable-timeout timeout;
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      fcoe-trusted;
      mac-limit limit action action;
      no-allowed-mac-log;
    }
  }
}

```

```
vlan (all | vlan-name) {
  (arp-inspection | no-arp-inspection) [
    forwarding-class (for DHCP Snooping or DAI Packets) class-name;
  ]
  dhcp-option82 {
    circuit-id {
      prefix (Circuit ID for Option 82) hostname;
      use-interface-description;
      use-vlan-id;
    }
    remote-id {
      prefix (Remote ID for Option 82) hostname | mac | none;
      use-interface-description;
      use-string string;
    }
    vendor-id <string>;
  }
  (examine-dhcp | no-examine-dhcp) {
    forwarding-class (for DHCP Snooping or DAI Packets) class-name;
  }
  examine-fip {
    examine-vn2vn {
      beacon-period milliseconds;
    }
    fc-map fc-map-value;
  }
  mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
```

[edit fabric] Hierarchy

```
fabric
  aliases {
    director-device director-device-name {
      assigned-director-device-name;
    }
  }
```

```

interconnect-device interconnect-device-name {
    assigned-interconnect-device-name;
}
node-device node-device-name {
    assigned-node-device-name;
}
}
resources {
    node-group node-group-name {
        node-device node-device-name;
        network-domain;
    }
}
routing-options {
    multicast
        fabric-optimized-distribution;
        no-make-before-break;
}
}
}

```

[edit fc-fabrics] Hierarchy

```

fc-fabrics {
    fc-fabric-name {
        description
        fabric-id fc-fabric-id;
        fabric-type proxy;
        interface {
            interface-name {
                max-login-sessions max-login-sessions;
            }
            interface-name {
                max-login-sessions max-login-sessions;
            }
            <...>;
            max-login-sessions max-login-sessions;
        }
        vlan.interface-name;
    }
    fc2 {
        traceoptions {
            file filename <replace> <size size> <files number> <no-stamp>;
            <world-readable | no-world-readable>;
            flag flag <flag-modifier>;
        }
    }
    max-login-sessions max-login-sessions;
    protocols {
        fip {
            fcoe-trusted;
            fc-map fc-map-value;
            fka-adv-period milliseconds;
            interface {
                interface-name {
                    fka-adv-period milliseconds;
                }
            }
        }
    }
}

```

```
        priority priority;
    }
}
max-sessions-per-enode max-sessions-per-enode;
priority priority;
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>;
    <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}
proxy {
    auto-load-rebalance
    load-balance-algorithm (simple | enode-based | flogi-based);
    no-fabric-wwn-verify;
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>;
        <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
}
}
```

[edit fc-options] Hierarchy

```
fc-options
max-login-sessions-per-node max-login-sessions-per-node;
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>;
    <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

[edit firewall] Hierarchy

```
firewall {
    family family-name {
        filter filter-name {
            interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
    }
}
policer policer-name {
    filter-specific;
    if-exceeding {
        bandwidth-limit bps;
    }
}
```

```

        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
three-color-policer policer-name {
    action {
        loss-priority high then discard;
    }
    single-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}
}
}

```

[edit groups] Hierarchy

```

groups {
    group-name {
        configuration-data;
    }
    global {
        configuration-data
    }
    if-config {
        configuration-data
    }
    rel {
        configuration-data
    }
}

```

[edit interfaces] Hierarchy

```

interfaces {
    aex {
        disable;
        aggregated-ether-options {
            configured-flow-control {
                rx-buffers (on | off);
                tx-buffers (on | off);
            }
            (flow-control | no-flow-control);
            lacp mode {
                admin-key key;
            }
        }
    }
}

```

```
    periodic interval;  
    system-id mac-address;  
  }  
  link-speed speed;  
  loopback;  
  no-loopback;  
  minimum-links number;  
}  
mc-ae {  
  chassis-id chassis-id;  
  mc-ae-id mc-ae-id;  
  mode (active-active);  
  status-control (active | standby);  
}  
description text;  
gratuitous-arp-reply | no-gratuitous-arp-reply)  
hold-time down milliseconds up milliseconds;  
mtu bytes;  
no-gratuitous-arp-request;  
traceoptions;  
(traps | no traps);  
unit logical-unit-number {  
  disable;  
  description text;  
  family {  
    ethernet-switching {  
      filter input filter-name;  
      filter output filter-name;  
      native-vlan-id vlan-id;  
      port-mode mode;  
      reflective-relay;  
      vlan {  
        members [ (all | names | vlan-ids) ];  
      }  
    }  
  }  
  inet {  
    address address {  
      primary;  
    }  
    filter input filter-name;  
    filter output filter-name;  
    primary;  
    targeted-broadcast;  
  }  
  (traps | no traps);  
  vlan-id vlan-id-number;  
}  
vlan-tagging;  
}  
interface-range interface-range-name {  
  disable;  
  description text;  
  ether-options {  
    802.3ad aex {  
      lacp {  
        force-up;  
      }  
    }  
  }  
}
```



```

    }
  }
  (auto-negotiation| no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
  disable;
  description text;
  family family-name {...}
  (traps | no traps);
  vlan-id vlan-id-number;
}
}
}
lo0 {
  disable;
  description text;
  hold-time milliseconds down milliseconds;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    disable;
    description text;
    family {
      inet {
        address address {
          primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
      }
    }
    (traps | no traps);
  }
}
}
mex {
  disable;
  description text;
  hold-time milliseconds down milliseconds;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  no-gratuitous-arp-request;
  traceoptions;
  traps;
  unit logical-unit-number {
    disable;
    description text;

```

```
family {
  ethernet-switching {
    filter input filter-name;
    filter output filter-name;
    native-vlan-id vlan-id;
    port-mode mode;
    reflective-relay;
    vlan {
      members [ (all | names | vlan-ids) ];
    }
  }
  inet {
    address address {
      primary;
      filter input filter-name;
      filter output filter-name;
      primary;
      targeted-broadcast;
    }
  }
  traps;
  vlan-id vlan-id-number;
}
vlan-tagging;
vlan {
  disable;
  description text;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    description text;
    disable;
    family {
      inet {
        address address {
          primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
      }
      (traps | no traps);
    }
  }
}
fc-0/0/port {
  fibrechannel-options {
    bb-sc-n;
    (loopback | no-loopback);
    speed (auto-negotiation | 2g | 4g | 8g);
  }
  unit logical-unit-number {
```

```

    disable;
    description text;
    family {
        fibre-channel {
            port-mode np-port;
        }
        (traps | no traps);
    }
ge-0/0/port {
    disable;
    description text;
    ether-options {
        802.3ad aex {
            lacp {
                force-up;
                primary;
            }
        }
    }
    (auto-negotiation | no-auto-negotiation);
    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    (flow-control | no-flow-control);
    link-mode mode;
    loopback;
    no-loopback;
    speed (auto-negotiation | speed);
}
gratuitous-arp-reply| no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    description text;
    disable;
    family {
        ethernet-switching {
            filter input filter-name;
            filter output filter-name;
            native-vlan-id vlan-id;
            port-mode mode;
            reflective-relay;
            vlan {
                members [ (all | names | vlan-ids) ];
            }
        }
    }
    inet {
        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
    }
}

```

```

        targeted-broadcast;
    }
    (traps | no traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
vrp-group group-id {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route prefix/prefix-length routing-instance instance-name priority-cost priority;
    }
}
virtual-address [ addresses ];
}
xe-0/0/port {
    disable;
    description text;
    ether-options {
        802.3ad aex {
            lacp {
                force-up;
                (primary | backup);
            }
        }
    }
    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    (flow-control | no-flow-control);
    loopback;
    no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply)
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    disable;
    description text;
    family {

```

```

ethernet-switching {
  filter input filter-name;
  filter output filter-name;
  native-vlan-id vlan-id;
  port-mode mode;
  reflective-relay;
  vlan {
    members [ (all | names | vlan-ids) ];
  }
}
fibre-channel {
  port-mode (f-port | np-port);
}
inet {
  address address {
    primary;
  }
  filter input filter-name;
  filter output filter-name;
  primary;
  targeted-broadcast;
}
(traps | no traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

[edit policy-options] Hierarchy

```

policy-options
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
}

```

[edit protocols] Hierarchy

```
protocols {
  bgp {
    disable;
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain key-chain;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      hold-down-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | single-hop);
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family family-name {
      ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
    }
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
    group group-name {
      ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
    }
    hold-time seconds;
    import [ policy-names ];
    include-mp-next-hop;
    keep (all | none);
    local-address address;
    local-as autonomous-system <loops number> <alias> <private>;
```

```

local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}

```

```
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
        version (Liveness Detection) (1 | automatic);
    }
    local-ip-addr ipv4-address;
    session-establishment-hold-time seconds;
}
session-establishment-hold-time seconds;
traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable {
        interface interface-name;
    }
}
interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
        group ip-address;
    }
}
```



```

        robust-count number;
    }
}
isis {
    disable;
    export [ policy-names ];
    ignore-attached-bit;
    interface interface-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
    }
    checksum;
    csnp-interval (seconds | disable);
    disable;
    hello-padding (adaptive | loose | strict);
    level (1 | 2) {
        disable;
        hello-authentication-key key;
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric number;
        metric metric;
        passive;
        priority number;
    }
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-ipv4-multicast;
    no-unicast-topology;
    passive;
    point-to-point;
}
level (1 | 2) {
    disable;
    authentication-key key;
    authentication-type authentication;
    external-preference preference;
    no-csnp-authentication;
}

```

```
no-hello-authentication;
no-psnp-authentication;
preference preference;
prefix-export-limit number;
wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
topologies {
    ipv4-multicast;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag (detail | disable | receive | send);
    }
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
    forward-delay seconds;
```

```

hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix </prefix-length> <exact> <override-metric metric > <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}

```

```

    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication {

```

```

        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
    }
    dead-interval seconds;
    demand-circuit;
    flood-reduction;
    hello-interval seconds;
    ipsec-sa sa-name;
    no-neighbor-down-notification;
    retransmit-interval seconds;
    topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
    }
    transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
}

```

```
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
  bootstrap-import [ policy-names ];
```

```

bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
prefix-list prefix-list-addresses {
    source source-address {
        next-hop next-hop-address;
    }
    wildcard-source {
        next-hop next-hop-address;
    }
}
}

```

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}
rip {
  authentication-key password;
  authentication-type type;
  (check-zero | no-check-zero);
  group group-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  export [ policy-names ];
  import [ policy-names ];
  metric-out metric;
  neighbor neighbor-name {
    any-sender;
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
      ... same statements as at the [edit protocols rip group group-name
        bfd-liveness-detection] hierarchy level ...
    }
    (check-zero | no-check-zero);
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive (both | none | version-1 | version-2);
    route-timeout seconds;
    send (broadcast | multicast | none | version-1);
    update-interval seconds;
  }
  preference preference;
  route-timeout seconds;
  update-interval seconds;
}
holddown seconds;
```



```

import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number > <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
}

```

```
    }
    max-age seconds;
  }
  traceoptions {
    file filename <files number> <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
  uplink-failure-detection {
    group group-name {
      link-to-monitor interface-name;
      link-to-disable interface-name;
    }
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number> <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
```

[edit security] Hierarchy

```
security {
  certificates
  pki
  ssh-known-hosts
  traceoptions
}
```

[edit snmp] Hierarchy

```
snmp {
```

```

client-list client-list-name {
    ip-addresses;
}
community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
        address restrict;
    }
    logical-system logical-system-name {
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
    }
    routing-instance routing-instance-name {
        clients {
            addresses;
        }
    }
    view view-name;
}
contact contact;
description description;
filter-duplicates;
filter-interfaces;
health-monitor {
    falling-threshold integer;
    interval seconds;
    rising-threshold integer;
}
interface [ interface-names ];
location location;
name name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type;
        rising-event-index index;
        rising-threshold integer;
        sample-type (absolute-value | delta-value);
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag syslog-subtag;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
    }
}

```

```
    type type;
  }
  history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
```

```

message-processing-model (v1 | v2c | V3);
security-level (authentication | none | privacy);
security-model (usm | v1 | v2c);
security-name security-name;
}
}
usm {
  local-engine {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
      privacy-aes128 {
        privacy-password privacy-password;
      }
      privacy-des {
        privacy-password privacy-password;
      }
      privacy-3des {
        privacy-password privacy-password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
      privacy-aes128 {
        privacy-password privacy-password;
      }
      privacy-des {
        privacy-password privacy-password;
      }
      privacy-3des {
        privacy-password privacy-password;
      }
      privacy-none {
        privacy-password privacy-password;
      }
    }
  }
}
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {

```

```
    security-model (any | usm | v1 | v2c) {
      security-level (authentication | none | privacy) {
        notify-view view-name;
        read-view view-name;
        write-view view-name;
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
view view-name {
  oid object-identifier (include | exclude);
}
}
```

[edit system] Hierarchy

```
system {
  accounting {
    events [ login change-log interactive-commands ];
    destination {
      radius {
        server {
          server-address {
            accounting-port port-number;
            retry number;
            secret password;
            source-address address;
            timeout seconds;
          }
        }
      }
    }
    tacplus {
      server {
        server-address {
          port port-number;
          secret password;
          single-connection;
          timeout seconds;
        }
      }
    }
  }
  archival {
    configuration {
      archive-sites {

```

```

        ftp://<username>:<password>@<host>:<port>/<url-path>;
        ftp://<username>:<password>@<host>:<port>/<url-path>;
    }
    transfer-interval interval;
    transfer-on-commit;
}
}
arp {
    aging-timer minutes;
    interfaces;
}
authentication-order [ authentication-methods ];
(compress-configuration-files | no-compress-configuration-files);
default-address-selection;
domain-name domain-name;
domain-search [ domain-list ];
host-name hostname;
internet-options {
    icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
    source-port upper-limit <upper-limit>;
}
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
login {
    announcement text;
    class class-name {
        access-end;
        access-start;
        allow-configuration "regular-expression";
        allowed-days "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
}

```

```
retry-options {
  backoff-factor seconds;
  backoff-threshold number;
  minimum-time seconds;
  tries-before-disconnect number;
}
user username {
  authentication {
    (encrypted-password "password" | plain-text-password);
    load-key-file URL;
    remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
  }
  uid uid-value;
  class class-name;
  full-name complete-name;
}
}
name-server {
  address;
}
no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
  authentication-key number type type value password;
  serveraddress <key key-number> <version value> <prefer>;
}
ports {
  auxiliary {
    disable;
    insecure;
    type terminal-type;
  }
  console {
    disable;
    insecure;
    log-out-on-disconnect;
    type terminal-type;
  }
}
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
}
radius-options {
  password-protocol mschap-v2;
}
attributes {
```



```

        nas-ip-address ip-address;
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
    ftp {
        connection-limit limit;
        rate-limit limit;
    }
    service-deployment {
        servers server-address {
            port port-number;
        }
        source-address source-address;
    }
    ssh {
        root-login (allow | deny | deny-password);
        protocol-version [v1 v2];
        connection-limit limit;
        rate-limit limit;
    }
    telnet {
        connection-limit limit;
        rate-limit limit;
    }
    web-management {
        http {
            interfaces [ interface-names ];
            port port;
        }
        https {
            interfaces [ interface-names ];
            local-certificate name;
            port port;
        }
        session {
            idle-timeout [ minutes ];
            session-limit [ session-limit ];
        }
    }
    xnm-clear-text {
        connection-limit limit;
    }
}

```

```
    rate-limit limit;
  }
  xnm-ssl {
    connection-limit limit;
    local-certificate name;
    rate-limit limit;
  }
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
  console {
    facility severity;
  }
  file filename {
    archive {
      files number;
      size maximum-file-size;
      start-time "YYYY-MM-DD.hh:mm";
      transfer-interval minutes;
      (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
      brief;
    }
  }
}
host (hostname | other-routing-engine | scc-master) {
  explicit-priority;
  facility-override facility;
  facility severity;
  log-prefix string;
  match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
  facility severity;
  match "regular-expression";
}
}
tacplus-options {
```

```

    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

[edit vlans] Hierarchy

```

vlans {
    vlan-name {
        description text-description;
        dot1q-tunneling {
            customer-vlans (id | range);
        }
        filter input filter-name;
        filter output filter-name;
        interface interface-name {
            isolated;
            mapping (policy | tag push | native push);
            promiscuous;
        }
        isolation-vlan-id;
        l3-interface vlan.logical-interface-number;
        mac-limit number;
        mac-table-aging-time seconds;
        no-local-switching;
        no-mac-learning;
        primary-vlan vlan-name;
        pvlan extend-secondary-vlan-id vlan-id;
        vlan-id number;
        vlan-range vlan-id-low-vlan-id-high;
    }
}

```

access-end

Syntax	access-end <i>HH:MM</i> ;
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the end time for login access.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Time-Based User Access on page 161

access-start

Syntax	access-start <i>HH:MM</i> ;
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the start time for login access.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Time-Based User Access on page 161

accounting

```
Syntax  accounting {
          destination {
            radius {
              server {
                server-address {
                  accounting-port port-number;
                  secret password;
                  source-address address;
                  retry number;
                  timeout seconds;
                }
              }
            }
          }
          tacplus {
            server {
              server-address {
                port port-number;
                secret password;
                single-connection;
                timeout seconds;
              }
            }
          }
        }
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring RADIUS Accounting](#)
- [Configuring TACACS+ System Accounting on page 1714](#)

accounting-port

Syntax	<code>accounting-port <i>port-number</i>;</code>
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system radius server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the accounting port number on which to contact the RADIUS server.
Options	<i>number</i> —Port number on which to contact the RADIUS server. Default: 1813
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Accounting

allow-commands

Syntax	<code>allow-commands "<i>regular-expression</i>";</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the operational mode commands that members of a login class can use.
Default	If you omit this statement and the deny-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.
Options	<i>regular-expression</i> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying Access Privileges for Junos OS Operational Mode Commands on page 178• deny-commands on page 253• user on page 314

allow-configuration

Syntax	<code>allow-configuration "regular-expression";</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Explicitly allow configuration access to the specified levels in the hierarchy even if the permissions set with the permissions statement do not grant such access by default.
Default	If you omit this statement and the deny-configuration statement, users can edit only those commands for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1719 • Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1685 • deny-configuration on page 254 • user on page 314

allowed-days

Syntax	<code>allowed-days [<i>days-of-the-week</i>];</code>
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the days of the week when users can log in.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Time-Based User Access on page 161



allow-transients

Syntax	allow-transients;
Hierarchy Level	[edit systems scripts commit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, enable transient configuration changes to be committed.
Default	Transient changes are disabled by default. If you do not include the allow-transients statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Generating a Persistent or Transient Change</i>• <i>Creating a Macro to Read the Custom Syntax and Generate Related Configuration Statements</i>

announcement

Syntax	announcement <i>text</i> ;
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a system login announcement. This announcement appears after a user logs in.
Options	text —Text of the announcement. If the text contains any spaces, enclose it in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Junos OS to Display a System Login Announcement</i>• Configuring the Junos OS to Display a System Login Message on page 149• message on page 267

archival

Syntax	<pre> archival { configuration { archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } transfer-interval interval; transfer-on-commit; } } </pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.
<div>  NOTE: The <code>edit system archival</code> hierarchy is not available on QFabric systems. </div>	
Options	The remaining statements are explained separately.
<div>  NOTE: The <code>[edit system archival]</code> hierarchy is not available on QFabric systems. </div>	
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611

arp (System)

Syntax

```
arp {  
    aging-timer minutes;  
    gratuitous-arp-delay seconds;  
    gratuitous-arp-on-ifup;  
    interfaces {  
        interface-name {  
            aging-timer minutes;  
        }  
    }  
    passive-learning;  
    purging;  
}
```

For EX-Series switches:

```
arp {  
    aging-timer minutes;  
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

For EX-Series switches, set only the time interval between ARP updates.

Options **aging-timer**—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

passive-learning (QFX-Series only)—Configure switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests.

Default: 20 minutes

Range: 1 to 240 minutes

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Junos® OS Network Interfaces*

- For more information about ARP updates, see the [Junos OS System Basics Configuration Guide](#).

authentication (Login)

Syntax	<pre>authentication { encrypted-password <i>password</i>; load-key-file <i>URL</i>; plain-text-password <i>password</i>; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "<i>public-key</i>"; ssh-rsa "<i>public-key</i>"; }</pre>
Hierarchy Level	[edit system login user <i>username</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Authentication methods that a user can use to log in to the switch. You can assign multiple authentication methods to a single user.
Options	<p>encrypted-password "<i>password</i>"—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.</p> <p>You cannot configure a blank password for encrypted-password using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p>load-key-file—Load RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.</p> <p>plain-text-password—Plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it.</p> <p>remote-debug-permission (QFabric systems only)—QFabric component authentication. Specifies permission levels for users to access individual components in a QFabric system.</p> <p>ssh-dsa "<i>public-key</i>"—SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p> <p>ssh-rsa "<i>public-key</i>"—SSH version 1 and SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS User Accounts on page 1692 • root-authentication on page 286

authentication-key

Syntax	<code>authentication-key <i>key-number</i> type <i>type</i> value <i>password</i>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure Network Time Protocol (NTP) authentication keys so that the router or switch can send authenticated packets. If you configure the router or switch to operate in authenticated mode, you must configure a key.</p> <p>Both the keys and the authentication scheme (MD5) must be identical between a set of peers sharing the same key number.</p>
Options	<p><i>key-number</i>—An integer in the range of 1 to 65533.</p> <p><i>type type</i>—Authentication type. It can only be md5.</p> <p><i>value password</i>—Key itself, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring NTP Authentication Keys (QFabric System) on page 153• NTP Time Server and Time Services Overview (QFabric System) on page 53

authentication-order

Syntax	<code>authentication-order [<i>authentication-methods</i>];</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.
Default	If you do not include the authentication-order statement, users are verified based on their configured passwords.
Options	<p><i>authentication-methods</i>—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:</p> <ul style="list-style-type: none"> • password—Use the password configured for the user with the authentication statement at the [edit system login user] hierarchy level. • radius—Use RADIUS authentication services. • tacplus—Use TACACS+ authentication services.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1694 • authentication on page 239

auxiliary

Syntax	<pre>auxiliary { disable; insecure; type <i>terminal-type</i>; }</pre>
Hierarchy Level	[edit system ports]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the characteristics of the auxiliary port.
Default	The auxiliary port is disabled.
Options	<p>disable—Disable the port.</p> <p>insecure—Disable superuser access or root logins to establish a terminal connection.</p> <p>type <i>terminal-type</i>—Type of terminal that is connected to the port.</p> <p>Range: ansi, vt100, small-xterm, xterm</p> <p>Default: The terminal type is unknown, and the user is prompted for the terminal type.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Console and Auxiliary Port Properties on page 6177

boot-server (NTP)

Syntax	<code>boot-server (address hostname);</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the server that NTP queries when the router or switch boots to determine the local date and time.</p> <p>When you boot the router or switch, it issues an ntpdate request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP cannot synchronize to a time server if the server time significantly differs from the local router's or switch's time. You can configure either an IP address or a hostname for the boot server. If you configure a hostname instead of an IP address, the ntpdate request resolves the hostname to an IP address when the router or switch boots up.</p>
Options	<ul style="list-style-type: none">• address—IP address of an NTP boot server.• hostname—Hostname of an NTP boot server.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Synchronizing and Coordinating Time Distribution Using NTP on page 180

broadcast

Syntax	<code>broadcast address <key key-number> <version value> <tll value>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the local router or switch to operate in broadcast mode with the remote system at the specified address to send periodic broadcast messages to a client population. Normally, you include this statement only when the local router or switch is operating as a transmitter.
Options	<p>address—Broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be 224.0.1.1.</p> <p>key key-number—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer).</p> <p>tll value—(Optional) Time-to-live (TTL) value to use. Range: 1 through 255 Default: 1</p> <p>version value—(Optional) Specify the version number to be used in outgoing NTP packets. Range: 1 through 4 Default: 4</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the NTP Time Server and Time Services on page 153

broadcast-client

Syntax	<code>broadcast-client;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 160

change-type

Syntax	<code>change-type (character-sets set-transitions);</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set requirements for using character sets in plain-text passwords. When you combine this statement with the minimum-changes statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.
Options	Specify one of the following: <ul style="list-style-type: none"> • character-sets—Number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters. • set-transitions—Number of transitions between character sets.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1687 • minimum-changes on page 267

checksum

Syntax	<code>checksum (md5 sha-256 sha1) <i>hash</i>;</code>
Hierarchy Level	[edit event-options event-script file <i>filename</i>], [edit system scripts commit file <i>filename</i>],
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event or commit script, the Junos OS verifies the authenticity of the script by using the configured checksum hash.
Options	md5 <i>hash</i> —MD5 checksum of this script. sha-256 <i>hash</i> —SHA-256 checksum of this script. sha1 <i>hash</i> —SHA-1 checksum of this script.
Required Privilege Level	maintenance —To view this statement in the configuration. maintenance-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i>• <i>Configuring Checksum Hashes for an Event Script</i>• <i>Configuring Checksum Hashes for an Op Script</i>• file checksum md5 on page 339• file checksum sha-256 on page 341• file checksum sha1 on page 340

class (Defining Login Classes)

Syntax	<pre> class <i>class-name</i> { access-end; access-start; allow-commands "<i>regular-expression</i>"; allow-configuration "<i>regular-expression</i>"; deny-commands "<i>regular-expression</i>"; deny-configuration "<i>regular-expression</i>"; idle-timeout <i>minutes</i>; login-tip; permissions [<i>permissions</i>]; } </pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define a login class.
Options	<p><i>class-name</i>—A name you choose for the login class.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Defining Junos OS Login Classes on page 1716 • user on page 314


class (Assigning a Class to an Individual User)

Syntax	<pre>class <i>class-name</i> { operator; read-only; super-user; unauthorized; }</pre>
Hierarchy Level	[edit system login user <i>username</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a user's login class. You must configure one class for each user.
Options	<i>class-name</i> —One of the classes defined at the [edit system login class] hierarchy level.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 1692

commit

Syntax	<pre> commit { allow-transients; direct-access; file <i>filename</i> { checksum (md5 sha-256 sha1) <i>hash</i>; optional; refresh; refresh-from <i>url</i>; sourceurl; } refresh; refresh-from <i>url</i>; traceoptions { file <<i>filename</i>> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } } </pre>
Hierarchy Level	[edit system scripts]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, configure the commit-time scripting mechanism.
Options	The statements are explained separately.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.

compress-configuration-files (System)

Syntax	(compress-configuration-files no-compress-configuration-files);
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Compress the current operational configuration file. The file is stored in the file juniper.conf, in the /config file system, along with the last three committed versions of the configuration. However, with large networks, the current configuration file might exceed the available space in the /config file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. The current configuration file is compressed on the second commit of the configuration after the first commit is made to include the compress-configuration-files statement.</p>
<div> NOTE: We recommend that you enable compression of the configuration files to minimize the amount of disk space that they require.</div>	
Default	The current operational configuration file is uncompressed.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Compressing the Current Configuration File on page 1595

console (Physical Port)

Syntax	<pre>console { disable; insecure; log-out-on-disconnect; type <i>terminal-type</i>; }</pre>
Hierarchy Level	[edit system ports]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the characteristics of the console port.
Default	The console port is enabled and its speed is 9600 baud.
Options	<p>disable—Disable console login connections.</p> <p>insecure—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.</p> <p>log-out-on-disconnect—Log out the session when the data carrier on the console port is lost.</p> <p>type <i>terminal-type</i>—Type of terminal that is connected to the port: ansi, vt100, small-xterm, or xterm.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Console and Auxiliary Port Properties on page 6177

default-address-selection

Syntax	default-address-selection;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Use the loopback interface, lo0, as the source address for all locally generated IP packets when the packet is sent through a routed interface, but not when the packet is sent through a local interface such as fxp0. The lo0 interface is the interface to the switch's Routing Engine.</p>
Default	<p>The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:</p> <ul style="list-style-type: none">• The primary address on the loopback interface lo0 that is <i>not</i> 127.0.0.1 is used.• The primary address for the primary interface or the preferred address (if configured) for the primary interface is used. <p>By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.</p> <p>An interface's <i>primary address</i> is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. An interface's <i>preferred address</i> is the default local address used for packets sourced by the local switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.</p> <p>To configure a different primary address or preferred address, include the primary or preferred statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] or [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] hierarchy levels.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 151• <i>Junos® OS Network Interfaces</i>

deny-commands

Syntax	<code>deny-commands "regular-expression";</code>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the operational mode commands that the user is denied permission to issue even though the permissions set with the permissions statement would allow their use.
Default	If you omit this statement and the allow-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying Access Privileges for Junos OS Operational Mode Commands on page 178• allow-commands on page 234• user on page 314

deny-configuration

Syntax	<code>deny-configuration "regular-expression";</code>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Explicitly deny configuration access to the specified levels in the hierarchy even if the permissions set with the permissions statement grant such access by default.
Default	If you omit this statement and the allow-configuration statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the permissions statement.
Options	regular-expression —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Specifying Access Privileges Using allow/deny-configuration Statements</i>• allow-configuration on page 235• user on page 314

destination (Accounting)

```
Syntax destination {
    radius {
        server {
            server-address {
                accounting-port port-number;
                secret password;
                source-address address;
                retry number;
                timeout seconds;
            }
        }
    }
    tacplus {
        server {
            server-address {
                port port-number;
                secret password;
                single-connection;
                timeout seconds;
            }
        }
    }
}
```

Hierarchy Level [edit system accounting]

Release Information Statement introduced before Junos OS Release 7.4.
radius statement added in Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the authentication server.

Options The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring RADIUS System Accounting on page 1697](#)
- [Configuring TACACS+ System Accounting](#)

destination-override

Syntax	<code>destination-override { syslog host <i>ip-address</i>; }</code>
Hierarchy Level	[edit system tracing]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Override the system-wide configuration of the switch at the [edit system tracing] hierarchy level. This statement has no effect if system tracing is not configured.
Options	syslog —System process log files to send to the remote tracing host. <ul style="list-style-type: none">• syslog—System process log files to send to the remote tracing host.• host <i>ip-address</i>—IP address to which to send tracing information.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Tracing and Logging Operations on page 6081• tracing on page 312

direct-access

Syntax	<code>direct-access;</code>
Hierarchy Level	[edit system scripts commit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify that commit scripts read input configurations directly from the database when inspecting these scripts for errors.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	

domain-name

Syntax	<code>domain-name <i>domain-name</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the name of the domain in which the switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.
Options	<i>domain-name</i> —Name of the domain.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Domain Name for the Router or Switch on page 146 • Example: Configuring the Domain Name for the Router or Switch on page 184

domain-search

Syntax	<code>domain-search <i>domain-list</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a list of domains to be searched.
Options	<i>domain-list</i> —List of domain names to search. The list can contain up to 6 domain names, with a total of up to 256 characters.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Domains to Search When a Router or Switch Is Included in Multiple Domains on page 147

explicit-priority

Syntax	<pre>explicit-priority { archive <files <i>number</i>> <size <i>size</i> <start-time<i>time</i>> <transfer-interval <i>interval</i>><world-readable no-world-readable>; archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } structured-data { brief; } }</pre>
Hierarchy Level	[edit system syslog file <i>filename</i>], [edit system syslog host]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.</p> <p>When the structured-data statement is also included at the [edit system syslog file <i>filename</i>] hierarchy level, this statement is ignored for the file.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Including Priority Information in System Log Messages on page 6213

events

Syntax	<code>events [<i>events</i>];</code>
Hierarchy Level	[edit system accounting]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the types of events to track and log.
Options	<p>events—Event types; can be one or more of the following:</p> <ul style="list-style-type: none"> • change-log—Audit configuration changes. • interactive-commands—Audit interactive commands (any command-line input). • login—Audit logins.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring TACACS+ System Accounting on page 1714

format

Syntax	<code>format (des md5 sha1);</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the authentication algorithm for plain-text passwords.
Default	For Junos OS, the default encryption format is md5 . For Junos OS-FIPS software, the default encryption format is sha1 .
Options	<p>The hash algorithm that authenticates the password can be one of three algorithms:</p> <ul style="list-style-type: none"> • des—Has a block size of 8 bytes; its key size is 48 bits long. • md5—Produces a 128-bit digest. • sha1—Produces a 160-bit digest.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1687

host-name

Syntax	<code>host-name <i>hostname</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the hostname of the switch.
Options	<i>hostname</i> —Name of the switch.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Hostname of the Router or Switch on page 147

icmpv4-rate-limit

Syntax	<code>icmpv4-rate-limit { bucket-size <i>seconds</i>; packet-rate <i>pps</i>; }</code>
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure rate-limiting parameters for ICMPv4 messages sent.
Options	bucket-size <i>seconds</i> —Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds Default: 5 packet-rate <i>pps</i> —Rate-limiting packets earned per second. Range: 0 through 4294967295 pps Default: 1000
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>ping</i>• Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages• Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages


idle-timeout

Syntax	<code>idle-timeout <i>minutes</i>;</code>
Hierarchy Level	<code>[edit system login class <i>class-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For a login class, configure the maximum time that a session can be idle before the user is logged off the switch. The session times out after remaining at the CLI operational mode prompt for the specified time.
Default	If you omit this statement, a user is never forced off the system after extended idle times.
Options	<i>minutes</i> —Maximum idle time. Range: 0 through 4294967295 minutes
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Timeout Value for Idle Login Sessions on page 162

internet-options

Syntax	<pre>internet-options { icmpv4-rate-limit bucket-size <i>bucket-size</i> packet-rate <i>packet-rate</i>; source-port upper-limit <i>upper-limit</i>; }</pre>
Hierarchy Level	<code>[edit system]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure system IP options to protect against certain types of denial-of-service (DoS) attacks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 151 • Configuring the Junos OS to Extend the Default Port Address Range on page 151

load-key-file

Syntax	<code>load-key-file URL filename;</code>
Hierarchy Level	[edit system root-authentication], [edit system login user <i>username</i> authentication]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<div> NOTE: ECDSA is not supported on the QFabric system.</div> <p>Load RSA (SSH version 1 and SSH version 2) and DSA or ECDSA (SSH version 2) public keys from a previously-generated named file at a specified URL location or local path. The file contains one or more SSH keys that are copied into the configuration when the command is issued.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Password on page 158• Configuring the Root Password on page 1702• Configuring Junos OS User Accounts• Configuring Junos OS User Accounts on page 1692

location

Syntax	<pre>location { altitude <i>feet</i>; building <i>name</i>; country-code <i>code</i>; floor <i>number</i>; hcoord <i>horizontal-coordinate</i>; lata <i>service-area</i>; latitude <i>degrees</i>; longitude <i>degrees</i>; npa-nxx <i>number</i>; postal-code <i>postal-code</i>; rack <i>number</i>; vcoord <i>vertical-coordinate</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the system location.
Options	<p>altitude <i>feet</i>—Number of feet above sea level.</p> <p>building <i>name</i>—Name of the building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").</p> <p>country-code <i>code</i>—Two-letter country code.</p> <p>floor <i>number</i>—Floor in the building.</p> <p>hcoord <i>horizontal-coordinate</i>—Bellcore Horizontal Coordinate.</p> <p>lata <i>service-area</i>—Long-distance service area.</p> <p>latitude <i>degrees</i>—Latitude in degree format.</p> <p>longitude <i>degrees</i>—Longitude in degree format.</p> <p>npa-nxx <i>number</i>—First six digits of the phone number (area code and exchange).</p> <p>postal-code <i>postal-code</i>—Postal code.</p> <p>rack <i>number</i>—Rack number.</p> <p>vcoord <i>vertical-coordinate</i>—Bellcore Vertical Coordinate.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Specifying the Physical Location of the Switch on page 156

login

```
Syntax login {
    announcement text;
    class class-name {
        access-end "regular-expression";
        access-start "regular-expression";
        allow-commands "regular-expression";
        allow-configuration "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication authentication;
        (encrypted-password "password" | plain-text-password);
        load-key-file URL;
        remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
        ssh-dsa "public-key";
        ssh-rsa "public-key";
    }
    class class-name;
    full-name complete-name;
    uid uid-value;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure user access to the switch.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

- Related Documentation**
- [Defining Junos OS Login Classes on page 1716](#)

login-alarms

Syntax	login-alarms;
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Show system alarms automatically when an admin user logs in to the router or switch.
Options	<i>class-name</i> —Login class name.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring System Alarms to Appear Automatically Upon Login on page 161

login-tip

Syntax	login-tip;
Hierarchy Level	[edit system login class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Enable CLI tips at login.
Default	Disabled.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring CLI Tips on page 72

max-configurations-on-flash

Syntax	max-configurations-on-flash <i>number</i> ;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the number of configurations stored on the internal fixed media storage (for example, USB device).
Options	<i>number</i> —The number of configurations stored on the CompactFlash card. Range: 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	

maximum-length

Syntax	maximum-length <i>length</i> ;
Hierarchy Level	[edit system login passwords]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.
Default	For Junos OS-FIPS software, the maximum number of characters for plain-text passwords is 20. For Junos OS, no maximum is set.
Options	<i>length</i> —Maximum number of characters the password can include. Range: 1 to 64 characters
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1687• minimum-length on page 268

message

Syntax	<code>message text;</code>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a system login message. This message appears before a user logs in.
Options	<i>text</i> —Text of the message.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS to Display a System Login Message on page 149 • announcement on page 236

minimum-changes

Syntax	<code>minimum-changes number;</code>
Hierarchy Level	[edit system login passwords]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement is used in combination with the change-type statement. If the change type is character-sets, then the number of character sets included in the password is checked against the specified minimum. If the change type is set-transitions, then the number of character set changes in the password is checked against the specified minimum.</p>
Default	For Junos OS, the minimum number of changes is 1. For Junos-FIPS software, the minimum number of changes is 3.
Options	<i>number</i> —Minimum number of character sets (or character set changes) required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1687 • change-type on page 245

minimum-length

Syntax	minimum-length <i>length</i> ;
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.
Default	For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.
Options	length —Minimum number of characters the password must include. Range: 6 to 20 characters
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1687• maximum-length on page 266

minimum-lower-cases

Syntax	minimum-lower-cases <i>number</i> ;
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	<p>Specify the minimum number of lower-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-upper-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of lower-case letters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1687 • Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182 • password (Login) on page 278

minimum-numeric

Syntax	minimum-numeric <i>number</i> ;
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	<p>Specify the minimum number of numeric class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of numeric class characters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1687• Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182• password (Login) on page 278

minimum-punctuations

Syntax	<code>minimum-punctuations <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	<p>Specify the minimum number of punctuation class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-upper-cases, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of punctuation class characters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Special Requirements for Junos OS Plain-Text Passwords on page 1687 • Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182 • password (Login) on page 278

minimum-upper-cases

Syntax	<code>minimum-upper-cases <i>number</i>;</code>
Hierarchy Level	[edit system login password]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	<p>Specify the minimum number of upper-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as minimum-length, minimum-punctuations, minimum-lower-cases, and so on.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p>
Options	<i>number</i> —The minimum number of upper-case letters required for the password.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1687• Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182• password (Login) on page 278

multicast-client

Syntax	<code>multicast-client <<i>address</i>>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Network Time Protocol (NTP), configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet.
Options	<i>address</i> —(Optional) One or more IP addresses. If you specify addresses, the router or switch joins those multicast groups. Default: 224.0.1.1.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 160

name-server

Syntax	<code>name-server { <code>address</code>; }</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure one or more Domain Name System (DNS) name servers.
Options	<code>address</code> —Address of the name server. To configure multiple name servers, include multiple <code>address</code> options.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 146

no-multicast-echo

```
Syntax  no-multicast-echo {
        arp {
            aging-timer minutes;
            gratuitous-arp-delay seconds;
            gratuitous-arp-on-ifup;
            interfaces {
                interface-name {
                    aging-timer minutes;
                }
            }
            passive-learning;
            purging;
        }
        host-name hostname;
        location {
            altitude feet;
            building name;
            country-code code;
            floor number;
            hcoord horizontal-coordinate;
            lata service-area;
            latitude degrees;
            longitude degrees;
            npa-nxx number;
            postal-code postal-code;
            rack number;
            vcoord vertical-coordinate;
        }
        license {
            autoupdate URL;
        }
        renew before-expiration (number | interval number)
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

Default The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring the Junos OS to Disable the Routing Engine Response to Multicast Ping Packets</i> |
|------------------------------|--|

no-ping-record-route

Syntax	no-ping-record-route;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Junos OS to disable the reporting of the IP address in ping responses.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 148

no-ping-time-stamp

Syntax	no-ping-time-stamp;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Junos OS to disable the recording of timestamps in ping responses.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 148

no-redirects (IPv4 Traffic)

Syntax	no-redirects;
Hierarchy Level	[edit system], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	<p>Stop protocol redirect messages for IPv4 traffic from being sent on the entire switch or on an interface on the router or switch.</p> <p>To disable the sending of protocol redirect messages for the entire router or switch, include the no-redirects statement at the [edit system] hierarchy level.</p> <p>To disable the sending of protocol redirect messages on a specific interface, include the no-redirects statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] hierarchy level.</p>
Default	The router or switch sends redirect messages.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 148• <i>Understanding the Protocol Redirect Mechanism on EX Series Switches</i>• <i>Configuring Junos OS to Disable Sending Protocol Redirect Messages on EX Series Switches (CLI Procedure)</i>• <i>Junos® OS Network Interfaces</i>

ntp

Syntax	<pre>ntp { authentication-key <i>number</i> type <i>type</i> value <i>password</i>; boot-server <i>address</i>; broadcast <<i>address</i>> <<i>key key-number</i>> <<i>version value</i>> <<i>ttl value</i>>; broadcast-client; multicast-client <<i>address</i>>; peer <i>address</i> <<i>key key-number</i>> <<i>version value</i>> <<i>prefer</i>>; server <i>address</i> <<i>key key-number</i>> <<i>version value</i>> <<i>prefer</i>>; source-address <i>source-address</i>; trusted-key [<i>key-numbers</i>]; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure Network Time Protocol (NTP) on the switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Synchronizing and Coordinating Time Distribution Using NTP on page 180

ntp (QFabric)

Syntax	<pre>ntp { authentication-key <i>number</i> type <i>type</i> value <i>password</i>; server <i>address</i> <<i>key key-number</i>> <<i>version value</i>> <<i>prefer</i>>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure Network Time Protocol (NTP) on the switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring NTP Authentication Keys (QFabric System) on page 153 • Configuring the NTP Time Server and Time Services (QFabric System) on page 156 • authentication-key on page 240 • server on page 289

optional

Syntax	optional;
Hierarchy Level	[edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the file statement is missing from the /var/db/scripts/commit directory on the router.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Controlling Execution of Commit Scripts During Commit Operations

password (Login)

Syntax	<pre>password { change-type (set-transitions character-set); format (md5 sha1); maximum-length <i>length</i>; minimum-changes <i>number</i>; minimum-length <i>length</i>; minimum-lower-cases <i>number</i>; minimum-numeric <i>number</i>; minimum-punctuations <i>number</i>; minimum-upper-cases <i>number</i>; }</pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements.</p> <p>Using several password minimum requirement options will cause the minimum-length to be reset if the total sum of the required minimums exceeds the minimum-length setting.</p> <p>The individual statements are explained separately.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Special Requirements for Junos OS Plain-Text Passwords on page 1687• Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 182

peer

Syntax	<code>peer address <key key-number> <version value> <prefer>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For NTP, configure the local router or switch to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router or switch and the remote system can synchronize with each other. This configuration is useful in a network in which either the local router or switch or the remote system might be a better source of time.
Options	<p>address—Address of the remote system. You must specify an address, not a hostname.</p> <p>key key-number—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.</p> <p>Range: Any unsigned 32-bit integer</p> <p>prefer—(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p>version value—(Optional) Specify the NTP version number to be used in outgoing NTP packets.</p> <p>Range: 1 through 4</p> <p>Default: 4</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the NTP Time Server and Time Services on page 153

permissions

Syntax	<pre>permissions { storage; storage-control; }</pre>
Hierarchy Level	[edit system login class]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the login access privileges to be provided on the switch.
Options	<i>permissions</i> —Privilege type.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Access Privilege Levels on page 1692• Table 164 on page 1672• user on page 314

port (TACACS+ Server)

Syntax	<pre>port <i>port-number</i>;</pre>
Hierarchy Level	[edit system accounting destination tacplus server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the port number on which to contact the TACACS+ server.
Options	<i>number</i> —Port number on which to contact the TACACS+ server. Default: 49
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring TACACS+ System Accounting on page 1714

ports

Syntax	<pre> ports { auxiliary { disable; insecure; type <i>terminal-type</i>; } console { disable; insecure; log-out-on-disconnect; type <i>terminal-type</i>; } }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the properties of the console and auxiliary ports. The ports are located on the craft interface.</p> <p>See the switch hardware documentation for port locations.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Console and Auxiliary Port Properties on page 6177

radius (System)

Syntax	<pre>radius { server { server-address { accounting-port <i>port-number</i>; secret <i>password</i>; source-address <i>address</i>; retry <i>number</i>; timeout <i>seconds</i>; } } }</pre>
Hierarchy Level	[edit system accounting destination]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the RADIUS accounting server.
Options	<i>server-address</i> —Address of the RADIUS accounting server. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS System Accounting on page 1697

refresh (Commit Scripts)

Syntax	<code>refresh;</code>
Hierarchy Level	[edit system scripts commit], [edit system scripts file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at the source URL, as specified in the source statement at the same hierarchy level.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • refresh-from on page 283 • source on page 291

refresh-from (Commit Scripts)

Syntax	<code>refresh-from url;</code>
Hierarchy Level	[edit system scripts commit], [edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at a URL other than the URL specified in the source statement.
Options	url —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • refresh on page 283 • source on page 291

retry

Syntax	<code>retry number;</code>
Hierarchy Level	[edit system radius server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.
Options	<i>number</i> —Number of retries allowed for contacting a RADIUS server. Range: 1 through 10 Default: 3



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Authentication on page 1699• Configuring RADIUS Accounting• timeout on page 307

retry-options

Syntax	<pre> retry-options { backoff-threshold <i>number</i>; backoff-factor <i>seconds</i>; maximum-time <i>seconds</i>; minimum-time <i>seconds</i>; tries-before-disconnect <i>number</i>; } </pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Maximum number of times a user can attempt to enter a password while logging in through SSH or Telnet before being disconnected.
Options	<p>backoff-threshold <i>number</i>—Threshold for the number of failed login attempts before the user experiences a delay when attempting to reenter a password. Use the backoff-factor option to specify the length of delay, in seconds.</p> <p>Range: 1 through 3</p> <p>Default: 2</p> <p>backoff-factor <i>seconds</i>—Length of delay after each failed login attempt. The length of delay increases by this value for each subsequent login attempt after the value specified in the backoff-threshold option.</p> <p>Range: 5 through 10</p> <p>Default: 5</p> <p>maximum-time <i>seconds</i>—Maximum length of time that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the configured maximum-time, the connection is closed.</p> <p>Range: 20 through 300</p> <p>Default: 120</p> <p>minimum-time <i>seconds</i>—Minimum length of time that the connection remains open while the user is attempting to enter a password to log in.</p> <p>Range: 20 through 60</p> <p>Default: 20</p> <p>tries-before-disconnect <i>number</i>—Maximum number of times a user is allowed to attempt to enter a password to log in through SSH or Telnet.</p> <p>Range: 1 through 10</p> <p>Default: 10</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1716](#)

root-authentication

Syntax	<pre>root-authentication { (encrypted-password "<i>password</i>" load-key-password <i>URL</i> plain-text-password); ssh-dsa "<i>public-key</i>"; ssh-rsa "<i>public-key</i>"; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the authentication methods for the root-level user, whose username is root .
Options	<p>encrypted-password "<i>password</i>"— Specify the MD5 or other encrypted authentication password. You can specify only one encrypted password.</p> <p>You cannot configure a blank password for the encrypted-password option using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p>plain-text-password—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.</p> <p>ssh-dsa "<i>public-key</i>"—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.</p> <p>ssh-rsa "<i>public-key</i>"—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Password on page 1702• authentication on page 239

saved-core-context

Syntax	(saved-core-context no-saved-core-context);
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure whether the switch saves core files generated by internal Junos OS processes, along with contextual information (system log files and a copy of the current configuration):</p> <ul style="list-style-type: none"> • saved-core-context—The switch saves each core file and its associated context in a compressed tar file named <code>/var/tmp/process-name.core.core-number.tgz</code>. • no-saved-core-context—The switch does not save core files and their associated context.
Default	The switch saves core files.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Saving Core Files from Junos OS Processes</i> • saved-core-files on page 287

saved-core-files

Syntax	saved-core-files <i>number</i> ;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Save core files generated by internal Junos OS processes, but not the associated contextual information (configuration and system log files).
Options	<p><i>number</i>—Maximum number of core files to save.</p> <p>Range: 1 through 10</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Saving Core Files from Junos OS Processes</i> • saved-core-context on page 287

secret

Syntax	<code>secret <i>password</i>;</code>
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>], [edit system radius-server <i>server-address</i>], [edit system tacplus-server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local switch must match that used by the server.
Options	<i>password</i> —Password to use; can include spaces included in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Accounting• Configuring RADIUS Authentication on page 1699• Configuring TACACS+ Authentication on page 1711• Configuring TACACS+ System Accounting on page 1714

server (NTP)

Syntax	<code>server address <key key-number> <version value> <prefer>;</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For NTP, configure the switch to operate in client mode with the remote system at the specified server address. In this mode, the local switch can be synchronized with the remote system, but the remote system can never be synchronized with the local switch.
Options	<p>address—Address of the remote system. You must specify an address, not a hostname.</p> <p>key key-number—(Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.</p> <p>Range: Any unsigned 32-bit integer</p> <p>prefer—(Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p>version value—(Optional) Specify the version number to be used in outgoing NTP packets.</p> <p>Range: 1 through 4</p> <p>Default: 4</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • ntp on page 277

server (RADIUS Accounting)

Syntax	<pre>server { server-address { accounting-port port-number; retry number secret password; source-address address; timeout seconds; } }</pre>
Hierarchy Level	[edit system accounting destination radius]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure RADIUS logging. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS System Accounting on page 1697

server (TACACS+ Accounting)

Syntax	<pre>server { server-address { port port-number; secret password; single-connection; timeout seconds; } }</pre>
Hierarchy Level	[edit system accounting destination tacplus]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure TACACS+ logging. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring TACACS+ System Accounting on page 1714


single-connection

Syntax	single-connection;
Hierarchy Level	[edit system accounting destination tacplus server <i>server-address</i>], [edit system tacplus <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring TACACS+ Authentication on page 1711 • Configuring TACACS+ System Accounting on page 1714

source (Commit Scripts)

Syntax	source <i>url</i> ;
Hierarchy Level	[edit system scripts commit file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series .
Description	For Junos OS commit scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/commit</code> directory. When you include the refresh statement at the same hierarchy level and commit the configuration, the local copy is overwritten by the version stored at the specified URL.
Options	<i>url</i> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.

source-address (NTP, RADIUS, System Logging, or TACACS+)

Syntax	source-address <i>source-address</i> ;
Hierarchy Level	[edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>], [edit system ntp], [edit system radius-server <i>server-address</i>], [edit system syslog], [edit system tacplus-server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<div> NOTE: The source address is not supported at the [edit system radius-server] hierarchy on the QFabric.</div> <p>Specify a source address for each configured TACACS+ server, RADIUS server, NTP server, or the source address to record in system log messages that are directed to a remote machine.</p>
Options	source-address —Valid IP address configured on one of the switch interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all host <i>hostname</i> statements at the [edit system syslog] hierarchy level.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Authentication on page 1699• Synchronizing and Coordinating Time Distribution Using NTP on page 180• Specifying an Alternative Source Address for System Log Messages

source-port (Port Addresses)

Syntax	<code>source-port upper-limit <upper-limit>;</code>
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the range of port addresses.
Options	upper-limit <i>upper-limit</i> —(Optional) The range of port addresses and can be a value from 5000 through 65,355.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS to Extend the Default Port Address Range on page 151

ssh-dsa

Syntax	<code>ssh-dsa "public-key";</code>
Hierarchy Level	[edit system root-authentication] [edit system login user authentication]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the DSA (SSH version 2) public key. You can specify one or more public keys.
Options	ssh-dsa "public-key" —SSH version 2 authentication.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Root Password on page 158 • authentication on page 239 • <i>root-authentication</i>

ssh-rsa

Syntax	ssh-dsa " <i>public-key</i> ";
Hierarchy Level	[edit system root-authentication] [edit system login user authentication]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the RSA (SSH version 1) public key. You can specify one or more public keys.
Options	ssh-rsa " <i>public-key</i> "—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Password on page 1702• authentication on page 239• <i>root-authentication</i>

static-host-mapping

Syntax	<pre>static-host-mapping { hostname { alias [<i>alias</i>]; inet [<i>address</i>]; sysid <i>system-identifier</i>; } }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).
Options	<p>alias <i>alias</i>—Alias for the hostname.</p> <p>hostname—Fully qualified hostname.</p> <p>inet <i>address</i>—IP address. You can specify one or more IP addresses for the host.</p> <p>sysid <i>system-identifier</i>—ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Hostname of the Router or Switch on page 147

structured-data

Syntax structured-data {
 brief;
 }

Hierarchy Level [edit system syslog file *filename*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, *The syslog Protocol* (<http://tools.ietf.org/html/draft-ietf-syslog-protocol-23>).



NOTE: When this statement is included, other statements that specify the format for messages written to the file are ignored (the `explicit-priority` statement at the [edit system syslog file *filename*] hierarchy level and the `time-format` statement at the [edit system syslog] hierarchy level).

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- *Logging Messages in Structured-Data Format*
- [explicit-priority on page 258](#)
- [time-format on page 306](#)

syslog (System)

```

Syntax  syslog {
        archive {
            (binary-data| no-binary-data);
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        console {
            facility severity;
        }
        file filename {
            facility severity;
            explicit-priority;
            match "regular-expression";
            archive {
                (binary-data| no-binary-data);
                files number;
                size maximum-file-size;
                start-time "YYYY-MM-DD.hh:mm";
                transfer-interval minutes;
                (world-readable | no-world-readable);
            }
            structured-data {
                brief;
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
            match "regular-expression";
            source-address source-address;
            structured-data {
                brief;
            }
            port port number;
        }
        log-rotate-frequency frequency;
        server server name;
        source-address source-address;
        time-format (millisecond | year | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }

```

Hierarchy Level [edit logical-systems *logical-system-name* system],
[edit system]

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Support at the [edit logical-systems logical-system-name system] hierarchy level introduced in Junos OS Release 11.4.
Description	Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console. The remaining statements are explained separately.
Options	archive —Define parameters for archiving log messages. console —Send log messages of a specified class and severity to the console. file —Send log messages to a named file. host —Remote location to be notified of specific log messages. log-rotate-frequency —Configure the interval for checking logfile size and archiving messages. server —Name of the system log server in the inet.0 routing instance. source-address —Include a specified address as the source address for log messages. time-format —Additional information to include in the system log time stamp. user —Notify a specific user of the log event.
Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS System Log Configuration Overview</i>• <i>Junos OS System Log Messages Reference</i>• Overview of Single-Chassis System Logging Configuration on page 6155

system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```

```
altitude feet;  
building name;  
country-code code;  
floor number;  
hcoord horizontal-coordinate;  
lata service-area;  
latitude degrees;  
longitude degrees;  
npa-nxx number;  
postal-code postal-code;  
rack number;  
vcoord vertical-coordinate;  
}  
login {  
  announcement text;  
  class class-name {  
    access-end;  
    access-start;  
    allow-configuration "regular-expression";  
    allowed-days "regular-expression";  
    deny-commands "regular-expression";  
    deny-configuration "regular-expression";  
    idle-timeout minutes;  
    login-tip;  
    permissions [ permissions ];  
  }  
  message text;  
  password {  
    change-type (set-transitions | character-set);  
    format (md5 | sha1 | des);  
    maximum-length length;  
    minimum-changes number;  
    minimum-length length;  
  }  
  retry-options {  
    backoff-factor seconds;  
    backoff-threshold number;  
    minimum-time seconds;  
    tries-before-disconnect number;  
  }  
  user username {  
    authentication {  
      (encrypted-password "password" | plain-text-password);  
      load-key-file URL;  
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);  
      ssh-rsa "public-key";  
      ssh-dsa "public-key";  
    }  
    uid uid-value;  
    class class-name;  
    full-name complete-name;  
  }  
}  
name-server {  
  address;  
}
```



```

no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}

```

```
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}
```

```

}
console {
    facility severity;
}
file filename {
    archive {
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine | scc-master) {
    explicit-priority;
    facility-override facility;
    facility severity;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}
tacplus-options {
    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure system management properties.



NOTE: The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

Required Privilege system—To view this statement in the configuration.

Level system-control—To add this statement to the configuration.

tacplus

Syntax

```
tacplus {  
  server {  
    server-address {  
      port port-number;  
      secret password;  
      single-connection;  
      timeout seconds;  
    }  
  }  
}
```

Hierarchy Level [edit system accounting destination]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure TACACS+.

Options *server-address*—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

Required Privilege system—To view this statement in the configuration.

Level system-control—To add this statement to the configuration.


Related Documentation

- [Configuring TACACS+ System Accounting on page 1714](#)

tacplus-server

Syntax	<pre>tacplus-server server-address { port secret password; single-connection; source-address source-address; timeout seconds; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the TACACS+ server.
Options	<p>server-address—Address of the TACACS+ authentication server.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring TACACS+ Authentication on page 1711

time-format

Syntax	time-format (year millisecond year millisecond);
Hierarchy Level	[edit system syslog]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a file , console , or user statement at the [edit system syslog] hierarchy level, but not to destinations configured by a host statement.
Default	The timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, Aug 21 12:36:30 .
<div> NOTE: When the structured-data statement is included at the [edit system syslog file <i>filename</i>] hierarchy level, this statement is ignored for the file.</div>	
Options	millisecond —Include the millisecond in the timestamp. year —Include the year in the timestamp.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Including the Year or Millisecond in Timestamps on page 167• structured-data on page 296

timeout

Syntax	<code>timeout seconds;</code>
Hierarchy Level	[edit system radius-server <i>server-address</i>], [edit system tacplus-server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>], [edit system accounting destination tacplus server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the length of time that the local router or switch waits to receive a response from a RADIUS or TACACS+ server.
Options	seconds —Length of time to wait. Range: 1 through 90 seconds Default: 3 seconds
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS Accounting • Configuring TACACS+ System Accounting on page 1714 • retry on page 284

time-zone

Syntax	<code>time-zone (GMT <i>hour-offset</i> <i>time-zone</i>);</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the local time zone. To have the time zone change take effect for all processes running on the switch, you must reboot the switch.
Default	UTC
Options	<p>GMT <i>hour-offset</i>—Set the time zone relative to UTC time.</p> <p>Range: -14 through +12</p> <p>Default: 0</p> <p><i>time-zone</i>—Specify the time zone as UTC, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:</p> <p>Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek</p> <p>America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/Porto_Acre, America/Puerto_Rico, America/Rainy_River, America/Rankin_Inlet, America/Regina, America/Rosario, America/Santiago,</p>

America/Santo_Domingo, America/Sao_Paulo, America/Scoresbysund,
 America/Shiprock, America/St_Johns, America/St_Kitts, America/St_Lucia,
 America/St_Thomas, America/St_Vincent, America/Swift_Current, America/Tegucigalpa,
 America/Thule, America/Thunder_Bay, America/Tijuana, America/Tortola,
 America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat,
 America/Yellowknife
 Antarctica/Casey, Antarctica/DumontDURville, Antarctica/Mawson, Antarctica/McMurdo,
 Antarctica/Palmer, Antarctica/South_Pole
 Arctic/Longyearbyen
 Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqtau, Asia/Aqtobe,
 Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut,
 Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca,
 Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong_Kong,
 Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul,
 Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk,
 Asia/Kuala_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila,
 Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom_Penh,
 Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul,
 Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran,
 Asia/Thimbu, Asia/Tokyo, Asia/Ujung_Pandang, Asia/Ulan_Bator, Asia/Urumqi,
 Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan
 Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape_Verde, Atlantic/Faeroe,
 Atlantic/Jan_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South_Georgia,
 Atlantic/St_Helena, Atlantic/Stanley
 Australia/Adelaide, Australia/Brisbane, Australia/Broken_Hill, Australia/Darwin,
 Australia/Hobart, Australia/Lindeman, Australia/Lord_Howe, Australia/Melbourne,
 Australia/Perth, Australia/Sydney
 Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade,
 Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest,
 Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki,
 Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana,
 Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk,
 Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague,
 Europe/Riga, Europe/Rome, Europe/Samara, Europe/San_Marino, Europe/Sarajevo,
 Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn,
 Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius,
 Europe/Warsaw, Europe/Zagreb, Europe/Zurich
 Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro,
 Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte,
 Indian/Reunion
 Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate,
 Pacific/Enderbury, Pacific/Fakaofo, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos,
 Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston,
 Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas,
 Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea,
 Pacific/Pago_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port_Moresby,
 Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu,
 Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation • [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 172](#)

traceoptions (Commit Scripts)

Syntax	<pre>traceoptions { file <filename> <files number> <size size> <world-readable no-world-readable>; flag flag; no-remote-trace; }</pre>
Hierarchy Level	[edit system scripts commit],
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define tracing operations for commit or op scripts.
Default	If you do not include this statement, no script-specific tracing operations are performed.
Options	<p>filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory <code>/var/log</code>. By default, commit script process tracing output is placed in the file <code>cscript.log</code> and op script process tracing is placed in the file <code>op-script.log</code>. If you include the file statement, you must specify a filename. To retain the default, you can specify <code>cscript.log</code> or <code>op-script.log</code> as the filename.</p> <p>files number—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed and compressed to <i>trace-file.0.gz</i>. When <i>trace-file</i> again reaches its maximum size, <i>trace-file.0.gz</i> is renamed <i>trace-file.1.gz</i> and <i>trace-file</i> is renamed and compressed to <i>trace-file.0.gz</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.</p> <p>Range: 2 through 1000</p> <p>Default: 10 files</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none">• all—Log all operations• events—Log important events• input—Log script input data• offline—Generate data for offline development• output—Log script output data• rpc—Log script RPCs• xslt—Log the XSLT library <p>no-world-readable—Restrict file access to owner. This is the default.</p>

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB


Range: 10 KB through 1 GB

Default: 128 KB

world-readable—Enable unrestricted file access.

Required Privilege	maintenance—To view this statement in the configuration.
Level	maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Tracing Commit Script Processing</i>

tracing

Syntax	<pre>tracing { destination-override syslog host <i>ip-address</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the switch to enable remote tracing to a specified host IP address.
<hr/>	
<div> NOTE: The tracing statement is not supported on the QFX3000 QFabric system.</div> <hr/>	
<p>The following processes are supported:</p> <ul style="list-style-type: none">• chassisd—Chassis-control process• eventd—Event-processing process• cosd—Class-of-service process <p>If you enabled remote tracing but wish to disable it for specific processes on the switch, use the no-remote-trace statement at the [edit system <i>process-name</i> traceoptions] hierarchy level.</p>	
Default	Remote tracing is disabled by default.
Options	destination-override syslog host <i>ip-address</i> —Overrides the global configuration for system tracing and has no effect if the tracing statement is not configured.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Tracing and Logging Operations on page 6081• destination-override on page 256

trusted-key

Syntax	<code>trusted-key [<i>key-numbers</i>];</code>
Hierarchy Level	[edit system ntp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For NTP, configure the keys to use when you configure the switch to synchronize its time with other systems on the network.
Options	<i>key-numbers</i> —One or more key numbers. Each key can be any 32-bit unsigned integer except 0.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring NTP Authentication Keys on page 152 • <i>authentication-key</i> • server on page 289

uid

Syntax	<code>uid <i>uid-value</i>;</code>
Hierarchy Level	[edit system login user]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a user identifier for a login account.
Options	<i>uid-value</i> —Number associated with the login account. This value must be unique on the router or switch. Range: 100 through 64000
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS User Accounts on page 1692

use-imported-time-zones

Syntax	use-imported-time-zones;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a custom time zone from a locally generated time zone database.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Setting a Custom Time Zone on Routers or Switches Running Junos OS on page 175

user (Access)

Syntax	<pre>user username { authentication { (encrypted-password "password" plain-text-password); load-key-file URL; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "public-key" <from hostname>; ssh-rsa "public-key" <from hostname>; } class class-name; full-name "complete-name"; uid uid-value; }</pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure access permission for individual users.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Junos OS User Accounts on page 1692• class on page 247

CHAPTER 7

Administration

- [Routine Monitoring on page 315](#)
- [Operational Commands on page 318](#)

Routine Monitoring

- [Monitoring System Process Information on page 315](#)
- [Monitoring System Properties on page 316](#)
- [Monitoring Interface Status and Traffic on page 317](#)
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 318](#)

Monitoring System Process Information

Purpose View the processes running on the QFX Series.

Action To view the software processes running on the QFX Series:

[edit system]

user@switch> [show system processes](#)

Meaning [Table 47 on page 315](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 47: Summary of System Process Information Output Fields

Field	Values
PID	Identifier of the process.
Name	Owner of the process.
State	Current state of the process.
CPU Load	Percentage of the CPU that is being used by the process.
Memory Utilization	Amount of memory that is being used by the process.

Table 47: Summary of System Process Information Output Fields (*continued*)

Field	Values
Start Time	Time of day when the process started.

- Related Documentation**
- [Monitoring System Properties on page 316](#)
 - [show system uptime on page 1067](#)

Monitoring System Properties

Purpose View system properties such as the name and IP address of a QFX Series product and resource usage.

Action To monitor system properties in the CLI, enter the following commands:

- [show system uptime](#)
- [show system users](#)
- [show system storage](#)

Meaning [Table 48 on page 316](#) summarizes key output fields in the system properties display.

Table 48: Summary of Key System Properties Output Fields

Field	Values	Additional Information
General Information		
Serial Number	Serial number for a QFX Series product.	
Junos OS Version	Version of Junos OS active on the switch, including whether the software is for domestic or export use.	Export software is for use outside the USA and Canada.
Hostname	Name of the QFX Series product.	
IP Address	IP address of the QFX Series product.	
Loopback Address	Loopback address.	
Domain Name Server	Address of the domain name server.	
Time Zone	Time zone on the QFX Series product.	
Time		
Current Time	Current system time, in Coordinated Universal Time (UTC).	

Table 48: Summary of Key System Properties Output Fields (*continued*)

Field	Values	Additional Information
System Booted Time	Date and time when the QFX Series product was last booted and how long it has been running.	
Protocol Started Time	Date and time when the protocols were last started and how long they have been running.	
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last commit command.	
Load Average	CPU load average for 1, 5, and 15 minutes.	
Storage Media		
Internal Flash Memory	Usage details of internal flash memory.	
External Flash Memory	Usage details of external USB flash memory.	
Logged in Users Details		
User	Username of any user logged in to the switch.	
Terminal	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the user@switch field in show system users command output.
Idle Time	How long the user has been idle.	

- Related Documentation**
- [Monitoring System Process Information on page 315](#)
 - [show system processes on page 983](#)

Monitoring Interface Status and Traffic

Purpose View interface status to monitor interface bandwidth utilization and traffic statistics on the QFX Series product.

- Action**
- To view interface status for all the interfaces, enter **show interfaces xe**.
 - To view status and statistics for a specific interface, enter **show interfaces xe interface-name**.

- To view status and traffic statistics for all interfaces, enter either [show interfaces xe detail](#) or [show interfaces xe extensive](#).

Meaning For details about output from the CLI commands, see [show interfaces xe](#).

Other Tools to Configure and Monitor Devices Running Junos OS

Apart from the command-line interface, Junos OS also supports the following applications, scripts, and utilities that enable you to configure and monitor devices running Junos OS:

- Junos XML Management Protocol Application Programming Interface (API)—Application programmers can use the Junos XML Management Protocol API to monitor and configure Juniper Networks QFX Series products. Juniper Networks provides a Perl module with the API to help you more quickly and easily develop custom Perl scripts for configuring and monitoring the QFX Series.
- NETCONF Application Programming Interface (API)—Application programmers can also use the NETCONF API to monitor and configure Juniper Networks QFX Series products.
- Junos OS commit scripts—You can define scripts to enforce custom configuration tasks, enforce consistency, prevent common mistakes, and more. Every time you commit a new candidate configuration, the active commit scripts are called to inspect the new candidate configuration. If a configuration violates your custom rules, the script can instruct the Junos OS to perform various actions, including making changes to the configuration and generating custom, warning, and system log messages.
- Junos OS Op scripts—You can add your own commands to the operation-mode CLI. You can use these scripts to automate troubleshooting of known network problems and correct them.
- Junos OS event scripts—You can use event scripts to diagnose and fix issues, monitor the overall status of the system, and examine errors periodically. Event scripts are similar to op scripts except that certain events on the switch will trigger these scripts.
- Junos Space—The Junos Space application design allows multiple users concurrent access to its user interface. It also includes applications for network infrastructure automation.

Related Documentation

- [CLI User Interface Overview on page 70](#)
- [QFX Series Software Features Overview on page 15](#)
- [NETCONF XML Management Protocol Guide](#)
- [Understanding Device and Network Management Features on page 6077](#)

Operational Commands

- [commit](#)
- [clear log](#)

- clear chassis display message
- clear system commit
- clear system reboot
- file
- file archive
- file checksum md5
- file checksum sha1
- file checksum sha-256
- file compare
- file delete
- file list
- file rename
- file show
- load
- ping
- request chassis beacon
- request chassis cb
- request chassis fpc
- request chassis routing-engine master
- request message
- request system configuration rescue delete
- request system configuration rescue save
- request system halt
- request system license add
- request system license delete
- request system license save
- request system logout
- request system power-off
- request system reboot
- request system snapshot
- request system software add
- request system software configuration-backup
- request system software configuration-restore
- request system software delete
- request system software download
- request system software nonstop-upgrade
- request system software rollback

- [request system software validate](#)
- [request system storage cleanup](#)
- [request system zeroize](#)
- [restart](#)
- [rollback](#)
- [save](#)
- [show chassis alarms](#)
- [show chassis beacon](#)
- [show chassis ethernet-switch interconnect-device fpc](#)
- [show chassis ethernet-switch interconnect-device cb](#)
- [show chassis environment](#)
- [show chassis environment cb](#)
- [show chassis environment fpc](#)
- [show chassis environment pem](#)
- [show chassis environment routing-engine](#)
- [show chassis fan](#)
- [show chassis firmware](#)
- [show chassis fpc](#)
- [show chassis hardware](#)
- [show chassis lcd](#)
- [show chassis led](#)
- [show chassis location](#)
- [show chassis mac-addresses](#)
- [show chassis nonstop-upgrade node-group](#)
- [show chassis pic](#)
- [show chassis routing-engine](#)
- [show chassis temperature-thresholds](#)
- [show chassis zones](#)
- [show cli](#)
- [show cli authorization](#)
- [show cli directory](#)
- [show cli history](#)
- [show host](#)
- [show interfaces diagnostics optics](#)
- [show log](#)
- [show ntp associations](#)
- [show ntp status](#)

- `show subscribers`
- `show system alarms`
- `show system audit`
- `show system boot-messages`
- `show system buffers`
- `show system certificate`
- `show system commit`
- `show system configuration archival`
- `show system configuration rescue`
- `show system connections`
- `show system core-dumps`
- `show system directory-usage`
- `show system license`
- `show system processes`
- `show system reboot`
- `show system resource-cleanup processes`
- `show system rollback`
- `show system services service-deployment`
- `show system software`
- `show system statistics`
- `show system storage`
- `show system uptime`
- `show system users`
- `show system virtual-memory`
- `show version`
- `start shell`
- `test configuration`
- `traceroute`
- `traceroute monitor`

commit

Syntax `commit <<at <"string">> <and-quit> <check> <comment <"comment-string">>
<confirmed> <display detail> <fast-synchronize> <minutes> <synchronize><force>>`

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Option **fast-synchronize** added in Junos OS Release 12.2.

Description Commit the set of changes to the database and cause the changes to take operational effect.



NOTE: Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

Options `at <"string">`—(Optional) Save software configuration changes and activate the configuration at a future time, or upon reboot.

string is **reboot** or the future time to activate the configuration changes. Enclose the **string** value (including **reboot**) in quotation marks (" "). You can specify time in two formats:

- A time value in the form **hh:mm[:ss]** (hours, minutes, and optionally seconds)—Commit the configuration at the specified time, which must be in the future but before 11:59:59 PM on the day the **commit at** configuration command is issued. Use 24-hour time for the **hh** value; for example, **04:30:00** is 4:30:00 AM, and **20:00** is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the router.
- A date and time value in the form **yyyy-mm-dd hh:mm[:ss]** (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the **commit at** command is issued. Use 24-hour time for the **hh** value. For example, **2003-08-21 12:30:00** is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

For example, **commit at "18:00:00"**. For date and time, include both values in the same set of quotation marks. For example, **commit at "2005-03-10 14:00:00"**.

A *commit check* is performed when you issue the **commit at** configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.



NOTE: If Junos OS fails before the configuration changes become active, all configuration changes are lost.

You cannot enter the **commit** at configuration command when there is a pending reboot.

You cannot enter the **request system reboot** command once you schedule a commit operation for a specific time in the future.

You cannot commit a configuration when a scheduled commit is pending. For information about how to use the **clear** command to cancel a scheduled configuration, see the [CLI Explorer](#).

and-quit—(Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

check—(Optional) Verify the syntax of the configuration, but do not activate it.

comment <"*comment-string*">—(Optional) Add a comment that describes the committed configuration. The comment can be as long as 512 bytes and must be typed on a single line. You cannot include a comment with the **commit check** command. Enclose *comment-string* in quotation marks (" "). For example, **commit comment "Includes changes recommended by SW Lab"**.

confirmed <*minutes*>—(Optional) Require that the commit be confirmed within the specified amount of time. To confirm a commit, enter either a **commit** or **commit check** command. If the commit is not confirmed within the time limit, the configuration rolls back automatically to the precommit configuration and a broadcast message is sent to all logged-in users. To show when a rollback is scheduled, enter the **show system commit** command. The allowed range is 1 through 65,535 minutes, and the default is 10 minutes.

In Junos OS Release 11.4 and later, you can also use the **commit confirmed** command in the **[edit private]** configuration mode.

display detail—(Optional) Monitors the commit process.



NOTE: In Junos OS Release 10.4 and later, if the number of commit details or messages exceeds a page when used with the **| display detail** pipe option, the **more** pagination option on the screen is no longer available. Instead, the messages roll up on the screen by default, just like using the **commit** command with the **| no more** pipe option.

fast-synchronize—(Optional) Configure the commits to run in parallel on both the master and backup Routing Engines to reduce the time taken for commit synchronization.

synchronize <*force*>—(Optional) If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the

commit synchronize command. The Routing Engine on which you execute this command (request Routing Engine) copies and loads its candidate configuration to the other (responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. The **commit synchronize** command does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines by using the **force** option. When you issue the **commit synchronize** command with the **force** option from one Routing Engine, the configuration sessions on the other Routing Engine is terminated and its configuration synchronized with that on the Routing Engine from which you issued the command.



NOTE: When you issue the **commit synchronize** command, you must use the **apply-groups re0** and **re1** commands. For information about how to use groups, see *Disabling Inheritance of a Junos OS Configuration Group*.

The responding Routing Engine must use Junos OS Release 5.0 or later.

Required Privilege Level

configure—To enter configuration mode.



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

load merge
load replace
load override
load update

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*

Related Documentation

- *Verifying a Junos Configuration, Committing a Junos OS Configuration*
- *Scheduling a Junos Commit Operation*
- *Deactivating and Reactivating Statements and Identifiers in a Junos Configuration*
- *Monitoring the Junos Commit Process*
- *Adding a Comment to Describe the Committed Configuration*

Sample Output

commit | display detail (QFX Series)

```
user@host> commit | display detail
```



```

-----
2011-08-24 01:08:08.00691 PDT: begin creating snapshots
2011-08-24 01:08:09.00210 PDT: end creating snapshots
2011-08-24 01:08:09.00211 PDT: begin preparing metadata
2011-08-24 01:08:09.00228 PDT: end preparing metadata
2011-08-24 01:08:09.00229 PDT: begin computing dcf root changes
2011-08-24 01:08:09.00236 PDT: end computing dcf root changes
2011-08-24 01:08:09.00244 PDT: begin computing additions
2011-08-24 01:08:09.00251 PDT: end computing additions
2011-08-24 01:08:09.00251 PDT: begin local object validation
2011-08-24 01:08:09.00251 PDT: end local object validation
2011-08-24 01:08:09.00252 PDT: begin update instances
2011-08-24 01:08:09.00252 PDT: end update instances
2011-08-24 01:08:09.00252 PDT: begin adjust metadata
2011-08-24 01:08:09.00252 PDT: end adjust metadata
2011-08-24 01:08:09.00253 PDT: begin validate metadata
2011-08-24 01:08:09.00253 PDT: end validate metadata
2011-08-24 01:08:09.00253 PDT: begin adjust allocations
2011-08-24 01:08:09.00254 PDT: end adjust allocations
2011-08-24 01:08:09.00254 PDT: begin adjust dependencies
2011-08-24 01:08:09.00254 PDT: end adjust dependencies
2011-08-24 01:08:09.00255 PDT: begin instance validation
2011-08-24 01:08:09.00255 PDT: end instance validation
2011-08-24 01:08:09.00255 PDT: begin opening all sessions eagerly
2011-08-24 01:08:09.00277 PDT: begin request #1 [login]
2011-08-24 01:08:09.00278 PDT: end request #1 [login]
2011-08-24 01:08:09.00325 PDT: begin processing globals
2011-08-24 01:08:09.00330 PDT: begin waiting for stamp check
(qfabric-default---node0)
2011-08-24 01:08:09.00334 PDT: end reply #1 [login]
2011-08-24 01:08:09.00351 PDT: end reply #1 [login]
2011-08-24 01:08:09.00451 PDT: begin request #2 [open]
2011-08-24 01:08:09.00451 PDT: end request #2 [open]
2011-08-24 01:08:09.00451 PDT: begin request #3 [get commit history]
2011-08-24 01:08:09.00452 PDT: end request #3 [get commit history]
2011-08-24 01:08:09.00452 PDT: begin request #4 [load]
2011-08-24 01:08:09.00453 PDT: end request #4 [load]
2011-08-24 01:08:09.00453 PDT: begin request #5 [load]
2011-08-24 01:08:09.00454 PDT: begin reply #2 [open]
2011-08-24 01:08:09.00456 PDT: end reply #2 [open]
2011-08-24 01:08:09.00457 PDT: begin reply #3 [get commit history]
2011-08-24 01:08:09.00475 PDT: end reply #3 [get commit history]
2011-08-24 01:08:09.00476 PDT: begin reply #4 [load]
2011-08-24 01:08:09.00499 PDT: begin reply #5 [load]
2011-08-24 01:08:09.00501 PDT: end waiting for stamp check
(qfabric-default---node0)
2011-08-24 01:08:09.00501 PDT: begin waiting for open (qfabric-default---node0)
2011-08-24 01:08:09.00502 PDT: end waiting for open (qfabric-default---node0)
2011-08-24 01:08:09.00504 PDT: end processing globals
2011-08-24 01:08:09.00617 PDT: end request #5 [load]
2011-08-24 01:08:09.00617 PDT: begin request #6 [check]
2011-08-24 01:08:09.00617 PDT: end request #6 [check]
2011-08-24 01:08:09.00619 PDT: end reply #5 [load]
2011-08-24 01:08:09.00619 PDT: begin reply #6 [check]
2011-08-24 01:08:09.00730 PDT: end session
2011-08-24 01:08:09.00752 PDT: end request #5 [load]
2011-08-24 01:08:09.00754 PDT: begin request #6 [check]
2011-08-24 01:08:09.00755 PDT: end request #6 [check]
2011-08-24 01:08:09.00881 PDT: end request #5 [load]
2011-08-24 01:08:09.00961 PDT: begin commit to devices
2011-08-24 01:08:10.00668 PDT: begin request #8 [get commit history]

```

```
2011-08-24 01:08:10.00669 PDT: end request #8 [get commit history]
2011-08-24 01:08:10.00721 PDT: end session
2011-08-24 01:08:10.00727 PDT: end commit to devices
2011-08-24 01:08:10.00733 PDT: begin committing metadata
2011-08-24 01:08:10.00772 PDT: end committing metadata
2011-08-24 01:08:10.00772 PDT: begin calling commit callbacks
2011-08-24 01:08:10.00773 PDT: end calling commit callbacks
commit complete
```

clear log

Syntax	<code>clear log <i>filename</i></code> <code><all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Remove contents of a log file.
Options	<i>filename</i> —Name of the specific log file to delete. <code>all</code> —(Optional) Delete the specified log file and all archived versions of it.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show log on page 889
List of Sample Output	clear log on page 327
Output Fields	See file list for an explanation of output fields.

Sample Output

clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```

clear chassis display message

List of Syntax	Syntax on page 328 Syntax (TX Matrix Router) on page 328 Syntax (TX Matrix Plus Router) on page 328 Syntax (QFabric Systems) on page 328
Syntax	clear chassis display message
Syntax (TX Matrix Router)	clear chassis display message <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	clear chassis display message <lcc <i>number</i> sfc <i>number</i> >
Syntax (QFabric Systems)	clear chassis display message <node-device <i>name</i> interconnect-device <i>name</i> >
Release Information	Command introduced in Junos OS Release 7.5. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option for the TX Matrix Plus routers introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	(M40e, M160, M320, T Series routers, EX Series, and QFabric systems only) Clear or stop a text message on the craft interface display, which is on the front of the router or switch or on the LCD panel display on the router or switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.
Options	none —Clear or stop a text message on the craft interface display. interconnect-device <i>name</i> —(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Interconnect device. lcc <i>number</i> —(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none">• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

node-device *name*—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Node device in a Node group.

scc—(TX Matrix routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Required Privilege Level clear

Related Documentation

- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- [set chassis display message on page 1469](#)
- *show chassis craft-interface*

List of Sample Output [clear chassis display message on page 329](#)

Output Fields See *show chassis craft-interface* for an explanation of output fields.

Sample Output

clear chassis display message

The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
      +-----+
      |NOC contact Dusty|
      |(888) 526-1234   |
      +-----+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
```

```
Host OK LED:  On
Host fail LED: Off
FPCs      0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|host    |
|Up: 0+17:05:47|
|        |
|Temperature OK|
+-----+
```

clear system commit

Syntax	clear system commit
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear any pending commit operation.
Options	This command has no options.
Required Privilege Level	maintenance (or the actual user who scheduled the commit)
Related Documentation	<ul style="list-style-type: none"> • show system commit on page 939
List of Sample Output	clear system commit on page 331 clear system commit (None Pending) on page 331 clear system commit (User Does Not Have Required Privilege Level) on page 331
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

clear system reboot

List of Syntax	Syntax on page 332 Syntax (EX Series Switches) on page 332 Syntax (TX Matrix Router) on page 332 Syntax (TX Matrix Plus Router) on page 332 Syntax (QFX Series) on page 332
Syntax	clear system reboot <both-routing-engines>
Syntax (EX Series Switches)	clear system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	clear system reboot <both-routing-engines> <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	clear system reboot <both-routing-engines> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (QFX Series)	clear system reboot <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear any pending system software reboots or halts. When issued on a TX Matrix router without any options, the default behavior clears all pending system software reboots or halts on all T640 routers connected to the TX Matrix router. When issued on a TX Matrix Plus router without any options, the default behavior clears all pending system software reboots or halts on all T1600 or T4000 routers connected to the TX Matrix Plus router.
Options	none —Clear all pending system software reboots or halts. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for all the Routing Engines in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests on the l connected T1600 or T4000 LCCs.

all-members—(EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

infrastructure *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the fabric control Routing Engines or fabric manager Routing Engines.

interconnect-device *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

node-group *name*—(QFabric systems) (Optional) Clear all halt or reboot requests on the Node group.

scc—(TX Matrix routers only) (Optional) Clear all halt or reboot requests for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level maintenance

Related Documentation	<ul style="list-style-type: none">• <i>request system reboot</i>• request system reboot on page 390• Rebooting and Halting a QFX Series Product on page 172• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	clear system reboot on page 335 clear system reboot (TX Matrix Router) on page 335 clear system reboot (QFX Series) on page 335
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

clear system reboot (TX Matrix Router)

```
user@host> clear system reboot
scc-re0:
-----
No shutdown/reboot scheduled.
lcc0-re0:
-----
No shutdown/reboot scheduled.
lcc2-re0:
-----
No shutdown/reboot scheduled.
```

clear system reboot (QFX Series)

```
user@switch> clear system reboot node-group node1
No shutdown/reboot scheduled.
```

file

Syntax	<code>file <archive checksum compare copy delete list rename show source address archive></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Archive files from the device, copy files to and from the router or switch, calculate the file checksum, compare files, delete a file from the device, list files on the device, rename a file, show file contents, or show the local address to initiate a connection.
Options	<p>archive (Optional) —Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.</p> <p>checksum (Optional) —Calculate the Message Digest 5 (MD5) checksum of a file.</p> <p>compare (Optional) —Compare two local files and describe the differences between them in default, context, or unified output styles.</p> <p>copy (Optional) —Copy files from one place to another on the local switch or between the local switch and a remote system.</p> <p>delete (Optional) —Delete a file on the local switch.</p> <p>list (Optional) —Display a list of files on the local switch.</p> <p>rename (Optional) —Rename a file on the local switch.</p> <p>show (Optional) —Display the contents of a file.</p> <p>source address (Optional) —Specify the source address of the local file.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Viewing Files and Directories on a Device Running Junos OS</i>• CLI Explorer

file archive

Syntax	<code>file archive destination <i>destination</i> source <i>source</i> <compress></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
Options	<p>destination <i>destination</i>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none"> For archived files—The suffix .tar For archived and compressed files—The suffix .tgz <p>source <i>source</i>—Source of the original file or files. Specify the source as a URL or filename.</p> <p>compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix .tgz.</p>
Required Privilege Level	maintenance
List of Sample Output	file archive (Multiple Files) on page 337 file archive (Single File) on page 337 file archive (with Compression) on page 338
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination  
/var/log/messages-archive.tgz  
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

file checksum md5

Syntax	<code>file checksum md5 <pathname> filename</code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Calculate the Message Digest 5 (MD5) checksum of a file.
Options	<p>pathname—(Optional) Path to a filename.</p> <p>filename—Name of a local file for which to calculate the MD5 checksum.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • file checksum sha-256 on page 341 • file checksum sha1 on page 340 • <i>op</i>
List of Sample Output	file checksum md5 on page 339
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

file checksum sha1

Syntax	<code>file checksum sha1 <pathname> filename</code>
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the SHA-1 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum md5 on page 339• file checksum sha-256 on page 341• <i>op</i>
List of Sample Output	file checksum sha1 on page 340
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```


file checksum sha-256

Syntax	<code>file checksum sha-256 <pathname> filename</code>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
Options	<p>pathname—(Optional) Path to a filename.</p> <p>filename—Name of a local file for which to calculate the SHA-256 checksum.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • file checksum md5 on page 339 • file checksum sha1 on page 340 • <i>op</i>
List of Sample Output	file checksum sha-256 on page 341
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha-256

```

user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71

```

file compare

Syntax	<code>file compare (files <i>filename filename</i>)</code> <code><context unified></code> <code><ignore-white-space></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none">• Default—In the first line of output, c means lines were changed between the two files, d means lines were deleted between the two files, and a means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (<) in front of output lines refers to the first file. A right angle bracket (>) in front of output lines refers to the second file.• Context—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-).• Unified—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.
Options	<p>files <i>filename</i>—Names of two local files to compare.</p> <p>context—(Optional) Display output in context format.</p> <p>ignore-white-space—(Optional) Ignore changes in the amount of white space.</p> <p>unified—(Optional) Display output in unified format.</p>
Required Privilege Level	none
List of Sample Output	file compare files on page 343 file compare files context on page 343 file compare files unified on page 343 file compare files unified ignore-white-space on page 343
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

file delete

Syntax	<code>file delete <i>filename</i></code> <code><purge></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete a file on the local router or switch.
Options	<i>filename</i> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued. <i>purge</i> —(Optional) Overwrite regular files before deleting them.
Required Privilege Level	maintenance
List of Sample Output	file delete on page 345 file delete (Routing Matrix) on page 345
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax	<code>file list</code> <code><detail recursive></code> <code><filename></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of files on the local router or switch.
Options	none —Display a list of all files for the current directory. detail recursive —(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively. filename —(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.
Additional Information	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the file list command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the file list command.
Required Privilege Level	maintenance
List of Sample Output	file list on page 346 file list (Routing Matrix) on page 346
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

file rename

Syntax	<code>file rename <i>source destination</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Rename a file on the local router or switch.
Options	<i>destination</i> —New name for the file. <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
Required Privilege Level	maintenance
List of Sample Output	file rename on page 348 file rename (Routing Matrix) on page 348
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file rename

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file rename (Routing Matrix)

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```



```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
```

```
-----
/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```

file show

Syntax	<code>file show filename</code> <encoding (base64 raw)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a file.
Options	filename —Name of a file. For a routing matrix, the filename must include the chassis information. encoding (base64 raw) —(Optional) Encode file contents with base64 encoding or show raw text.
Required Privilege Level	maintenance
List of Sample Output	file show on page 350 file show (Routing Matrix) on page 350
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

load

Syntax	<code>load (factory-default merge override patch replace set update) load (<i>filename</i> terminal) <relative></code>
QFX Series	<code>load (dhcp-snooping <i>filename</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Load a configuration from an ASCII configuration file, from terminal input, or from the factory default. Your current location in the configuration hierarchy is ignored when the load operation occurs.
Options	<p>dhcp-snooping—(QFX Series switches) Loads DHCP snooping entries.</p> <p>factory-default—Loads the factory configuration. The factory configuration contains the manufacturer's suggested configuration settings. The factory configuration is the router or switch's first configuration and is loaded when the router or switch is first installed and powered on.</p> <p>On J Series Services Routers, pressing and holding down the Config button on the router for 15 seconds causes the factory configuration to be loaded and committed. However, this operation deletes all other configurations on the router; using the load factory-default command does not.</p> <p>filename—Name of the file to load. For information about specifying the filename, see <i>Viewing Files and Directories on a Device Running Junos OS</i>.</p> <p>merge—Combine the configuration that is currently shown in the CLI with the configuration.</p> <p>override—Discard the entire configuration that is currently shown in the CLI and load the entire configuration. Marks every object as changed.</p> <p>patch—Change part of the configuration and mark only those parts as changed.</p> <p>replace—Look for a replace tag in filename, delete the existing statement of the same name, and replace it with the configuration.</p> <p>set—Merge a set of commands with an existing configuration. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as set, edit, exit, and top.</p> <p>relative—(Optional) Execute set of commands relative to current edit point.</p> <p>terminal—Use the text you type at the terminal as input to the configuration. Type Ctrl+d to end terminal input.</p> <p>update—Discard the entire configuration that is currently shown in the CLI, and load the entire configuration. Marks changed objects only.</p>



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

Required Privilege Level configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Related Documentation

- [Loading a Configuration from a File on page 1597](#)


ping

List of Syntax [Syntax on page 354](#)
 [Syntax \(QFX Series\) on page 354](#)

Syntax `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet | inet6>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict >`
 `<strict-source value.>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`
 `<vpls instance-name>`
 `<wait seconds>`

Syntax (QFX Series) `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict>`
 `< strict-source value>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`

<wait *seconds*>

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Check host reachability and network connectivity. The ping command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>
Options	<p>host—IP address or hostname of the remote system to ping.</p> <p>bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p>count requests—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p>detail—(Optional) Include in the output the interface on which the ping reply was received.</p> <p>do-not-fragment—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: In Junos OS Release 11.1 and later, when issuing the ping command for an IPv6 route with the do-not-fragment option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p>inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p>inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p>interface source-interface—(Optional) Interface to use to send the ping requests.</p> <p>interval seconds—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p>logical-system logical-system-name—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the set cli logical-system logical-system-name command and then run the ping command. To return to the main router, enter the clear cli logical-system command.</p>

loose-source value—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address mac-address—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern string—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

record-route—(Optional) Record and report the packet's path (IPv4).

routing-instance routing-instance-name—(Optional) Name of the routing instance for the ping attempt.

size bytes—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

strict—(Optional) Use the strict source route option (IPv4).

strict-source value—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos type-of-service—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

ttl value—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

verbose—(Optional) Display detailed output.

vpls instance-name—(Optional) Ping the instance to which this VPLS belongs.

wait seconds—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level network

Related Documentation • *Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages*

List of Sample Output [ping hostname on page 357](#)
[ping hostname rapid on page 357](#)
[ping hostname size count on page 357](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

Sample Output

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

request chassis beacon

Syntax (QFX Series)	<code>request chassis beacon</code> <code><all (off on)></code> <code><fpc slot-number (off on)></code> <code><interconnect-device name (cb slot-number fpc slot-number (off on)></code> <code><node-device name (off on)></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	(QFX Series only) Enable or disable the beacon LED on a QFX Series device.
Options	<p>all—Turn the beacon LED either on or off on all QFabric system Interconnect and Node devices.</p> <p>cb slot-number—Turn the beacon LED either on or off on the Control Board of the QFX3008-I Interconnect device.</p> <p>fpc slot-number—Turn the beacon LED either on or off on the Flexible PIC Concentrator on the standalone QFX3500 switch or the Interconnect device.</p> <p>interconnect-device name—Turn the beacon LED either on or off on the Interconnect device.</p> <p>node-device name—Turn the beacon LED either on or off on the Node device.</p> <p>off—Turn the beacon LED off.</p> <p>on—Turn the beacon LED on.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• show chassis beacon on page 464
List of Sample Output	request chassis beacon fpc 0 on (QFX Series) on page 358 request chassis beacon node-device (QFabric System) on page 358 request chassis beacon on interconnect-device fpc (QFabric System) on page 359
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis beacon fpc 0 on (QFX Series)

```
user@switch> request chassis beacon fpc 0 on

Beacon set to ON
```

request chassis beacon node-device (QFabric System)

```
user@switch> request chassis beacon node-device node1 on
```

node1 ON

request chassis beacon on interconnect-device fpc (QFabric System)

user@switch> request chassis beacon on interconnect-device fpc 2

FPC 2 ON

request chassis cb

List of Syntax	Syntax on page 360 Syntax (TX Matrix Router) on page 360 Syntax (TX Matrix Plus Router) on page 360 Syntax (QFabric System) on page 360
Syntax	<code>request chassis cb (offline online) slot <i>slot-number</i></code>
Syntax (TX Matrix Router)	<code>request chassis cb (offline online) <slot <i>slot-number</i> lcc <i>number</i> slot <i>cb-slot-number</i> scc <i>number</i> slot <i>cb-slot-number</i>></code>
Syntax (TX Matrix Plus Router)	<code>request chassis cb (offline online) <slot <i>slot-number</i> lcc <i>number</i> slot <i>cb-slot-number</i> sfc <i>number</i> slot <i>cb-slot-number</i>></code>
Syntax (QFabric System)	<code>request chassis cb (offline online) interconnect-device <i>name</i> slot <i>slot-number</i></code> <code><interconnect-device <i>name</i> slot <i>slot-number</i> (offline online)></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS 9.4 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS 11.3 for QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	(M120, M320, and MX Series routers and T Series routers, QFabric systems, and EX8200 switches only) Control the operation of the Control Board (CB). For information about the meaning of “CBs” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> .
Options	offline —Take the Control Board offline.



NOTE: On a QFabric system, to bring the backup Control Board on a QFX3008-I Interconnect device offline, issue the `request chassis cb slot backup-slot-number offline` command.



NOTE: Only backup Control Board can be turned offline or online. To turn a Control Board offline or to bring it back online, the Routing Engine should be turned offline first.

online—Bring the Control Board online.

interconnect-device *name*—(QFabric systems only) (Optional) Bring the QFX3008-I Interconnect device Control Board either offline or online:

slot slot-number—Control Board slot number:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **cb-slot-number** with a value from 0 through 1.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **cb-slot-number** with a value from 0 through 1.

- M320 router—Replace **slot-number** with a value from 0 through 1.
- MX480/MX240 routers—Replace **slot-number** with a value from 0 through 1.
- MX960 router—Replace **slot-number** with a value from 0 through 2.
- MX2020 and MX2010 routers—Replace **slot-number** with 0 or 1.
- EX8208 switch—Replace **slot-number** with a value from 0 through 2.
- EX8216 switch—Replace **slot-number** with a value from 0 through 1.
- QFabric System—Replace **slot-number** with a value from 0 through 1.

lcc number—(TX Matrix, TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace **number** with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc number—(TX Matrix Plus routers only) (Optional) Change the CB status for the TX Matrix Plus router (switch-fabric chassis). Replace **number** with 0.

Required Privilege Level maintenance

Related Documentation

- [show chassis environment cb on page 561](#)
- *Switching Control Board Redundancy*
- *Routing Engine and Switching Control Board Redundancy Configuration Statements*

List of Sample Output

- [request chassis cb on page 362](#)
- [request chassis cb interconnect-device \(QFabric System\) on page 362](#)
- [request chassis cb \(MX2020 Router\) on page 362](#)
- [request chassis cb \(MX2010 Router\) on page 362](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cb

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb interconnect-device (QFabric System)

```
user@switch> request chassis cb interconnect-device interconnect1 offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2020 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2010 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis fpc

List of Syntax	Syntax on page 363 Syntax (TX Matrix and TX Matrix Plus Routers) on page 363 Syntax (MX Series Routers) on page 363 Syntax (MX2020 3D Universal Edge Routers) on page 363 Syntax (MX2010 3D Universal Edge Routers) on page 363 Syntax (QFabric System) on page 363 Syntax (PTX Series Packet Transport Switches) on page 363
Syntax	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (TX Matrix and TX Matrix Plus Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i> <lcc <i>number</i>></code>
Syntax (MX Series Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i> <all-members></code> <code><local></code> <code><member <i>member-id</i>></code>
Syntax (MX2020 3D Universal Edge Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (MX2010 3D Universal Edge Routers)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Syntax (QFabric System)	<code>request chassis fpc</code> <code><interconnect-device <i>name</i> slot <i>slot-number</i> (offline online)></code> <code><(offline online) interconnect-device <i>name</i> slot <i>slot-number</i>></code> <code><slot <i>slot-number</i> interconnect-device <i>name</i> (offline online)></code>
Syntax (PTX Series Packet Transport Switches)	<code>request chassis fpc (offline online restart) slot <i>slot-number</i></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS 11.3 for QFX Series.</p> <p>Command introduced in Junos OS 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	<p>(M20, M40, M40e, M120, M160, M320, MX Series, and T Series routers, QFabric systems, EX Series switches, and PTX Series Packet Transport Switches only) Control the operation of the Flexible PIC Concentrator (FPC). For information about the meaning of “FPCs” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i>.</p>



NOTE: Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

Options **offline**—Take the FPC offline.

online—Bring the FPC online.

interconnect-device *name*—(QFabric systems only) Bring the Flexible Port Concentrator (FPC) on the QFX3008-I Interconnect device either offline or online:

- (QFabric System) On a QFabric system, specify the name of the QFX3008-I Interconnect device containing the Flexible Port Concentrator (FPC) you want to bring either offline or online.

restart—Restart the FPC.

slot *slot-number*—FPC slot number:

- M20 router—0 through 3.
- M120 router—0 through 5.
- MX240 router—0 through 2. On the MX240 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX480 router—0 through 5. On the MX480 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX960 router—0 through 11. On the MX960 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX2020 router—0 through 19.
- MX2010 router—0 through 9.
- TX Matrix and TX Matrix Plus routers only—On the TX Matrix router, if you specify the number of the T640 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31. In case of TX Matrix Plus router with 3D SIBs, replace

slot-number with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

- Other routers—0 through 7.
- QFabric System—Replace *slot-number* with a value from 0 through 2.
- EX Series switches:
 - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9.
 - EX6210 switches—Replace *slot-number* with a value from 0 through 9.



NOTE: These commands are not supported for slots 4 and 5 when a Switch Fabric and Routing Engine (SRE) module is installed in those slots. These commands are supported for slots 4 and 5 only if a line card is installed in them.

- EX8208 switches—Replace *slot-number* with a value from 0 through 7.
- EX8216 switches—Replace *slot-number* with a value from 0 through 15.
- PTX5000 Packet Transport Switch—Replace *slot-number* with a value from 0 through 7.

all-members—(MX Series routers only) (Optional) Change FPC status of all members of the Virtual Chassis configuration.

local—(MX Series routers only) (Optional) Change FPC status of the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Change FPC status of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level maintenance

Related Documentation

- [show chassis fpc on page 641](#)
- *show chassis fpc-feb-connectivity*
- *show chassis fabric fpcs*
- *Configuring the Junos OS to Make a Flexible PIC Concentrator Stay Offline*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*

List of Sample Output

- [request chassis fpc on page 366](#)
- [request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\) on page 366](#)
- [request chassis fpc \(MX2020 Router\) on page 366](#)
- [request chassis fpc \(MX2010 Router\) on page 366](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request chassis fpc](#)

```
user@host> request chassis fpc online slot 0
FPC 0 already online
```

[request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\)](#)

```
user@host> request chassis fpc slot 0
Possible completions:
offline           Take FPC offline
online            Bring FPC online
restart           Restart FPC
```

[request chassis fpc \(MX2020 Router\)](#)

```
user@host >request chassis fpc online slot 2
FPC 2 already online
```

[request chassis fpc \(MX2010 Router\)](#)

```
user@host >request chassis fpc offline slot 5
Offline initiated, use "show chassis fpc" to verify
```

request chassis routing-engine master

List of Syntax	Syntax on page 367 Syntax (TX Matrix Routers) on page 367 Syntax (TX Matrix Plus Routers) on page 367 Syntax (MX Series Routers) on page 367 Syntax (MX2010 3D Universal Edge Routers) on page 367 Syntax (MX2020 3D Universal Edge Routers) on page 367 Syntax (QFabric Systems) on page 367
Syntax	request chassis routing-engine master (acquire release switch) <force> <no-confirm>
Syntax (TX Matrix Routers)	request chassis routing-engine master (acquire release switch) (lcc <i>number</i> scc all-chassis) <force> <no-confirm>
Syntax (TX Matrix Plus Routers)	request chassis routing-engine master (acquire release switch) (lcc <i>number</i> sfc all-chassis all-lcc) <force> <no-confirm>
Syntax (MX Series Routers)	request chassis routing-engine master (acquire release switch) <all-members> <force> <local> <member <i>member-id</i> > <no-confirm>
Syntax (MX2010 3D Universal Edge Routers)	request chassis routing-engine master (acquire release switch <check>) <no-confirm>
Syntax (MX2020 3D Universal Edge Routers)	request chassis routing-engine master (acquire release switch <check>) <no-confirm>
Syntax (QFabric Systems)	request chassis routing-engine master (release switch) <check> <interconnect-device <i>name</i> > <node-group <i>name</i> > <no-confirm>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>all-chassis option added in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>

Description For routers or switches with multiple Routing Engines, control which Routing Engine is the master.



CAUTION: (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS Release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix, one or all routers (in a routing matrix based on the TX Matrix router or in a routing matrix based on a TX Matrix Plus router) might become logically disconnected from the TX Matrix router and cause data loss. For more information, see the [TX Matrix Router Hardware Guide](#) or the [Junos OS High Availability Configuration Guide](#).



NOTE: Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

You will receive an error message stating “Command aborted. Not ready for mastership switch, try after n seconds” when this command is re-entered before 240 seconds have elapsed on EX Series switches.



NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the routing engine to the backup routing engine, and then reboot.

Options

- acquire**—Attempt to become the master Routing Engine.
- release**—Request that the other Routing Engine become the master.
- switch**—Toggle mastership between Routing Engines.

The **acquire**, **release**, and **switch** options have the following suboptions:

all-chassis—(TX Matrix and TX Matrix Plus routers only) On a routing matrix composed of a TX Matrix router and the attached T640 routers, switch mastership on all the Routing Engines in the routing matrix. Likewise, on a routing matrix composed of a TX Matrix Plus router and the attached T1600 or T4000 routers, switch mastership on all the Routing Engines in the routing matrix.

all-lcc—(TX Matrix Plus routers only) Request to acquire mastership for all line-card chassis (LCC).

all-members—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in all member routers of the Virtual Chassis configuration.

check—(QFabric systems, MX240, MX480, MX960, MX2010, and MX2020 Routers only) (Optional) Available only with the **switch** option. Check graceful switchover status of the standby Routing Engine before toggling mastership between Routing Engines.

interconnect-device *name*—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on an Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines of the specified member in the Virtual Chassis Configuration. Replace *member-id* with a value of 0 or 1.

no-confirm—(Optional) Do not request confirmation for the switch.

node-group *name*—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on a Node group.

scc—(TX Matrix routers only) TX Matrix (switch-card chassis).

sfc—(TX Matrix Plus routers only) TX Matrix Plus router (or switch-fabric chassis).

force—(Optional) Available only with the **acquire** option. Force the change to a new master Routing Engine.

Additional Information Because both Routing Engines are always running, the transition from one to the other as the master Routing Engine is immediate. However, the changeover interrupts communication to the System and Switch Board (SSB). The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. Interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

By default, the Routing Engine in slot 0 (RE0) is the master and the Routing Engine in slot 1 (RE1) is the backup. To change the default master Routing Engine, include the **routing-engine** statement at the **[edit chassis redundancy]** hierarchy level in the configuration. For more information, see the *Junos OS System Basics Configuration Guide*.

To have the backup Routing Engine become the master Routing Engine, use the **request chassis routing-engine master switch** command. If you use this command to change the master and then restart the chassis software for any reason, the master reverts to the default setting.



NOTE: Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

Required Privilege Level maintenance

Related Documentation

- [show chassis routing-engine on page 831](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

List of Sample Output [request chassis routing-engine master acquire on page 370](#)
[request chassis routing-engine master switch on page 371](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request chassis routing-engine master acquire](#)

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

request chassis routing-engine master switch

```
user@host> request chassis routing-engine master switch
```

```
warning: Traffic will be interrupted while the PFE is re-initialized  
Toggle mastership between Routing Engines ? [yes,no] (no) yes
```

```
Resolving mastership...
```

```
Complete. The other Routing Engine becomes the master.
```

Switch mastership back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch
```

```
warning: Traffic will be interrupted while the PFE is re-initialized  
Toggle mastership between routing engines ? [yes,no] (no) yes
```

```
Resolving mastership...
```

```
Complete. The local routing engine becomes the master.
```

request message


Syntax	<code>request message all message "text"</code> <code>request message message "text" (terminal <i>terminal-name</i> user <i>user-name</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a message on the screens of all users who are logged in to the router or switch or on specific screens.
Options	all —Display a message on the terminal of all users who are currently logged in. message "text" —Message to display. terminal <i>terminal-name</i> —Name of the terminal on which to display the message. user <i>user-name</i> —Name of the user to whom to direct the message.
Required Privilege Level	maintenance
List of Sample Output	request message message on page 372
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on tty0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```


request system configuration rescue delete


Syntax	request system configuration rescue delete
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete an existing rescue configuration.
<div>  NOTE: The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • request system configuration rescue save on page 374 • request system software rollback on page 416 • show system commit on page 939
List of Sample Output	request system configuration rescue delete on page 373
Output Fields	This command produces no output.

Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax	request system configuration rescue save
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the rollback command.
<div> NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.</div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• request system software delete on page 404• request system software rollback on page 416• show system commit on page 939
List of Sample Output	request system configuration rescue save on page 374
Output Fields	This command produces no output.


Sample Output

request system configuration rescue save

```
user@host> request system configuration rescue save
```

request system halt

List of Syntax	Syntax on page 375 Syntax (EX Series Switches) on page 375 Syntax (PTX Series) on page 375 Syntax (TX Matrix Router) on page 375 Syntax (TX Matrix Plus Router) on page 375 Syntax (MX Series Router) on page 376 Syntax (QFX Series) on page 376
Syntax	<pre>request system halt <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk removable-compact-flash usb)> <message "<i>text</i>"></pre>
Syntax (EX Series Switches)	<pre>request system halt <all-members> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> <slice <i>slice</i>></pre>
Syntax (PTX Series)	<pre>request system halt <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Router)	<pre>request system halt <all-lcc lcc <i>number</i> scc> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Plus Router)	<pre>request system halt <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></pre>

	<pre><at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "text"></pre>
Syntax (MX Series Router)	<pre>request system halt <all-members> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "text"> <other-routing-engine></pre>
Syntax (QFX Series)	<pre>request system halt <all-members> <at <i>time</i>> <director-device <i>director-device-id</i>> <in <i>minutes</i>> <local> <media > <member <i>member-id</i>> <message "text"> <slice <i>slice</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>other-routing-engine option introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>director-device option introduced for QFabric systems in Junos OS Release 12.2.</p> <p>backup-routing-engine option introduced in Junos OS Release 13.1.</p>
Description	<p>Stop the router or switch software.</p>
	<div> NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, use the member option from the Node group CLI. You cannot issue this command from the QFabric CLI.</div>
Options	<p>none—Stop the router or switch software immediately.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Halt all chassis.</p>

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 or T4000 routers connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Halt all members of the Virtual Chassis configuration.

at time —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

backup-routing-engine—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

both-routing-engines—(Optional) Halt both Routing Engines at the same time.

director-device *director-device-id*—(QFabric systems only) Halt a specific Director device.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Halt the local Virtual Chassis member.

in minutes—(Optional) Number of minutes from now to stop the software. This option is an alias for the **at +minutes** option.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to J Series routers only.)

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Halt the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before stopping the software.

other-routing-engine—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

slice *slice*—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.



NOTE: If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the **request system halt** command. You can also halt both Routing Engines at the same time by issuing the **request system halt both-routing-engines** command.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• clear system reboot on page 332• request system power-off on page 385• Rebooting and Halting a QFX Series Product on page 172• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	request system halt on page 380 request system halt (In 2 Hours) on page 380 request system halt (Immediately) on page 380 request system halt (At 1:20 AM) on page 380
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host> request system halt at now
```

request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```


request system license add

Syntax	<code>request system license add (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a license key.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Adding New Licenses (CLI Procedure) on page 93
List of Sample Output	request system license add on page 381
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license add

```
user@host> request system license add terminal
```

request system license delete

Syntax	<code>request system license delete (<i>license-id</i> license-identifier-list [<i>licenseid001 licenseid002 licenseid003</i>] all)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option license-identifier-list introduced in Junos OS Release 13.1.
Description	Delete a license key. You can choose to delete one license at a time, all licenses at once, or a list of license identifiers enclosed in brackets.
Options	license-identifier —Text string that uniquely identifies a license key. license-identifier-list [<i>licenseid001 licenseid002 licenseid003....</i>] — Delete multiple license identifiers as a list enclosed in brackets. all —Delete all licenses on the device.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Deleting a License (CLI Procedure) on page 94

request system license save

Syntax	<code>request system license save (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Save installed license keys to a file or URL.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Saving License Keys on page 95
List of Sample Output	request system license save on page 383
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license save

```
user@host> request system license save ftp://user@host/license.conf
```

request system logout

Syntax	<code>request system logout (pid <i>pid</i> terminal <i>terminal</i> user <i>username</i>) <all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Log out users from the router or switch and the configuration database. If a user held the configure exclusive lock, this command clears the exclusive lock.
Options	all —(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.) pid <i>pid</i> —Log out the user session using the specified management process identifier (PID). The PID type must be management process. terminal <i>terminal</i> —Log out the user for the specified terminal session. user <i>username</i> —Log out the specified user.
Required Privilege Level	configure
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS System Basics Configuration Guide</i>
List of Sample Output	request system logout on page 384
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system logout

```
user@host> request system logout user tammy all
Connection closed by foreign host.
```

request system power-off

List of Syntax	Syntax on page 385 Syntax (EX Series Switches) on page 385 Syntax (TX Matrix Router) on page 385 Syntax (TX Matrix Plus Router) on page 385 Syntax (MX Series Router) on page 385 Syntax (QFX Series) on page 386
Syntax	<pre>request system power-off <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk removable-compact-flash usb)> <message "<i>text</i>"></pre>
Syntax (EX Series Switches)	<pre>request system power-off <all-members> <at <i>time</i>> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> <slice <i>slice</i>></pre>
Syntax (TX Matrix Router)	<pre>request system power-off <all-chassis all-lcc lcc <i>number</i> scc> <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Plus Router)	<pre>request system power-off <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>> <both-routing-engines> <other-routing-engine> <at <i>time</i>> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (MX Series Router)	<pre>request system power-off <all-members> <at <i>time</i>> <both-routing-engines> <in <i>minutes</i>> <local></pre>

	<code><media (external internal)></code> <code><member <i>member-id</i>></code> <code><message "text"></code> <code><other-routing-engine></code>
Syntax (QFX Series)	<code>request system power-off</code> <code><at <i>time</i>></code> <code><in <i>minutes</i>></code> <code><media (external internal)></code> <code><message "text"></code> <code><slice <i>slice</i>></code>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Power off the software.



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

Options	none —Power off the router or switch software immediately.
	all-chassis —(Optional) (TX Matrix and TX Matrix Plus router only) Power off all Routing Engines in the chassis.
	all-lcc —(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, power off all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.
	all-members —(EX4200 switches and MX Series routers only) (Optional) Power off all members of the Virtual Chassis configuration.
	at <i>time</i> —(Optional) Time at which to power off the software, specified in one of the following ways: <ul style="list-style-type: none">• now—Power off the software immediately. This is the default.• +<i>minutes</i>—Number of minutes from now to power off the software.• <i>yymmddhhmm</i>—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.• <i>hh:mm</i>—Absolute time on the current day at which to power off the software.
	both-routing-engines —(Optional) Power off both Routing Engines at the same time.

in *minutes*—(Optional) Number of minutes from now to power off the software. This option is an alias for the **at +*minutes*** option.

lcc *number*—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off a T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, power off a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Power off the local Virtual Chassis member.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Power off the boot media:

- **external**—Power off the external mass storage device.
- **internal**—Power off the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Power off the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before powering off the software.

other-routing-engine—(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(Optional) (TX Matrix router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix router (or switch-card chassis). If you issue the command from the master Routing Engine, the master SCC is powered off. If you issue the command from the backup Routing Engine, the backup SCC is powered off.

sfc *number*—(Optional) (TX Matrix Plus router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix Plus router (or switch-fabric chassis). If you issue the command from the master Routing Engine, the master SFC is powered off. If you issue the command from the backup Routing Engine, the backup SFC is powered off. Replace *number* with zero.

slice *slice*—(EX Series and QFX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information On a routing matrix composed of a TX Matrix router and T640 routers, if you issue the **request system power-off** command on the TX Matrix master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

Likewise, on a routing matrix composed of a TX Matrix Plus router and T1600 routers, if you issue the **request system power-off** command on the TX Matrix Plus master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

If you issue the **request system power-off both-routing-engines** command on the TX Matrix or TX Matrix Plus router, all the Routing Engines on the routing matrix are powered off.

Required Privilege Level maintenance

List of Sample Output [request system power-off on page 388](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`request system power-off`

```
user@host> request system power-off message "This router will be powered off in 30 minutes.
Please save your data and log out immediately."
warning: This command will not halt the other routing-engine.
If planning to switch off power, use the both-routing-engines option.
Power Off the system ? [yes,no] (no) yes
```

```
*** FINAL System shutdown message from remote@nutmeg ***
System going down IMMEDIATELY
```

```
This router will be powered off in 30 minutes. Please save your data and log out
immediately.
```

```
Shutdown NOW!
[pid 5177]
```


request system reboot

Syntax	<code>request system reboot</code> <code><at time></code> <code><in minutes></code> <code><media (external internal)></code> <code><message "text"></code> <code><routing-engine></code>
Syntax (QFX Series)	<code>request system reboot</code> <code><all <graceful>></code> <code><director-device <i>name</i>></code> <code><director-group <graceful>></code> <code><fabric <graceful>></code> <code><node-group <i>name</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Reboot the Junos OS.



NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Reboot requests are recorded in the system log files, which you can view with the **show log messages** command. You can view the process names with the **show system processes** command.

- Options**
- none**—Reboots the software immediately.
- all**—(QFabric systems only) (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.
- at time**—(Optional) Time at which to reboot the software, specified in one of the following ways:
- **+minutes**—Number of minutes from now to reboot the software.
 - **hh:mm**—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
 - **now**—Stop or reboot the software immediately. This is the default.
 - **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- director-device *name***—(QFabric systems only) (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group—(QFabric systems only) (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric—(QFabric systems only) (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

graceful—(QFabric systems only) (Optional) Allows the QFabric component to reboot with minimal impact to network traffic. This option is only available for the **all**, **fabric**, and **director-group** options.

in minutes—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at +minutes** option.

media (external | internal)—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

message "text"—(Optional) Message to display to all system users before rebooting the software.

node-group name—(QFabric systems only) (Optional) Reboots the software on a server Node group or a network Node group.

routing-engine—(Optional) Reboot the Routing Engine.

Required Privilege Level maintenance

Related Documentation

- [clear system reboot on page 332](#)
- [Rebooting and Halting a QFX Series Product on page 172](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

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

request system reboot (At 2300)

```
user@switch> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request the system to reboot in 2 hours:

```
user@switch> request system reboot at +120
user@switch> request system reboot in 120
user@switch> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@switch> request system reboot at now
```

request system reboot (At 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@switch> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot director-device

```
user@switch> request system reboot director-device Node1
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group

```
user@switch> request system reboot director-group
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group graceful

```
user@switch> request system reboot director-group graceful
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system snapshot

Syntax	<code>request system snapshot</code> <code><partition></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Copy the currently running Junos OS and configuration to alternate media. This command takes a snapshot of the contents of the <code>/</code> (root), and <code>/var</code> partitions on the media used to boot the switch and then copies the snapshot to alternate media. If the switch was booted from internal flash memory, the snapshot is copied to an external USB flash drive. If the switch was booted from an external USB flash drive, the snapshot is copied to internal flash memory.
Options	<p>none—Create a snapshot on the alternate media—that is, the external media if you booted the switch using software stored on internal media or internal media if you booted the switch using software stored on external media.</p> <p>partition—Partition the destination media before copying over the snapshot.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>show system snapshot</i> • Creating a Snapshot and Using It to Boot a QFX Series Switch on page 166 • <i>Verifying That a System Snapshot Was Created on a QFX Series Switch</i>
List of Sample Output	request system snapshot partition on page 393

Sample Output

request system snapshot partition

```

user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (da1) ...
Verifying compatibility of destination media partitions...
Running newfs (334MB) on external media / partition ...
Running newfs (404MB) on external media /config partition ...
Running newfs (222MB) on external media /var partition ...
Copying '/dev/da0s2a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s2f' to '/dev/da1s1f' .. (this may take a few minutes)
The following filesystems were archived: / /config /var

```

request system software add

List of Syntax	Syntax on page 394 Syntax (EX Series Switches) on page 394 Syntax (TX Matrix Router) on page 394 Syntax (TX Matrix Plus Router) on page 394 Syntax (MX Series Router) on page 395 Syntax (QFX Series) on page 395
Syntax	<code>request system software add <i>package-name</i></code> <code><best-effort-load></code> <code><delay-restart></code> <code><force></code> <code><no-copy></code> <code><no-validate></code> <code><re0 re1></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code> <code><unlink></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code> <code><validate></code>
Syntax (EX Series Switches)	<code>request system software add <i>package-name</i></code> <code><best-effort-load></code> <code><delay-restart></code> <code><force></code> <code><no-copy></code> <code><no-validate></code> <code><re0 re1></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code> <code><validate></code>
Syntax (TX Matrix Router)	<code>request system software add <i>package-name</i></code> <code><best-effort-load></code> <code><delay-restart></code> <code><force></code> <code><lcc <i>number</i> scc></code> <code><no-copy></code> <code><no-validate></code> <code><re0 re1></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code> <code><unlink></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code> <code><validate></code>
Syntax (TX Matrix Plus Router)	<code>request system software add <i>package-name</i></code> <code><best-effort-load></code>

```

<delay-restart>
<force>
<lcc number | sfc number>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>

```

Syntax (MX Series Router) request system software add *package-name*

```

<best-effort-load>
<delay-restart>
<force>
<member member-id>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>

```

Syntax (QFX Series) request system software add *package-name*

```

<best-effort-load>
<component all>
<delay-restart>
<force>
<no-copy>
<no-validate>
<partition>
<reboot>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>

```

Release Information Command introduced before Junos OS Release 7.4.
best-effort-load and **unlink** options added in Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
set [*package-name package-name*] option added in Junos OS Release 11.1 for EX Series switches.
set [*package-name package-name*] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.



NOTE: On EX Series switches, the set `[package-name package-name]` option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis, whereas, on M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways, the set `[package-name package-name]` option allows you to install multiple software packages and software add-on packages at the same time.

`upgrade-with-config` and `upgrade-with-config-format` *format* options added in Junos OS Release 12.3 for M Series routers, MX Series routers, T Series routers, EX Series Ethernet switches, and QFX Series devices.

Description

NOTE: We recommend that you always download the software image to `/var/tmp` only. On EX Series and QFX Series switches, you must use the `/var/tmp` directory. Other directories are not supported.

Install a software package or bundle on the router or switch.

Options *package-name*—Location from which the software package or bundle is to be installed.

For example:

- `/var/tmp/package-name`—For a software package or bundle that is being installed from a local directory on the router or switch.
- `protocol://hostname/pathname/package-name`—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
 - **ftp**—File Transfer Protocol.
Use `ftp://hostname/pathname/package-name`. To specify authentication credentials, use `ftp://<username>:<password>@hostname/pathname/package-name`. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
 - **http**—Hypertext Transfer Protocol.
Use `http://hostname/pathname/package-name`. To specify authentication credentials, use `http://<username>:<password>@hostname/pathname/package-name`. If a password is required and you omit it, you are prompted for it.
 - **scp**—Secure copy (available only for Canada and U.S. version).
Use `scp://hostname/pathname/package-name`. To specify authentication credentials, use `scp://<username>:<password>@hostname/pathname/package-name`.

**NOTE:**

- The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
- Do not use the `scp` protocol in the `request system software add` command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the MGD process which does not support `scp`.
Use the `file copy` command to copy the software package or bundle from the remote location to the `/var/tmp` directory on the hard disk:
`file copy scp://source/package-name /var/tmp`
Then install the software package or bundle using the `request system software add` command:
`request system software add /var/tmp/package-name`
- On a J Series Services Router, when you install the software from a remote location, the package is removed at the earliest opportunity in order to make room for the installation to be completed. If you copy the software to a local directory on the router and then install the new package, use the `unlink` option to achieve the same effect and allow the installation to be completed.

best-effort-load—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

component all—(QFabric systems only) (Optional) Install software package on all of the QFabric components.

delay-restart—(Optional) Install a software package or bundle, but do not restart software processes.

force—(Optional) Force the addition of the software package or bundle (ignore warnings).

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) In a routing matrix based on the TX Matrix router, install a software package or bundle on a T640 router (or line-card chassis) that is connected to the TX Matrix router. In a routing matrix based on the TX Matrix Plus router, install a software package or bundle on a T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a value from 0 through 3.

member member-id—(MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *member-id* with a value of 0 or 1.

partition—(QFX3500 switches only) (Optional) Format and repartition the media before installation.

scc—(TX Matrix routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

no-copy—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-validate—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

re0 | re1—(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (**re0**) or the Routing Engine in slot 1 (**re1**).

reboot—(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

set [package-name package-name]—(Mixed EX4200 and EX4500 Virtual Chassis only) (Optional) Install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.

set [package-name package-name]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages and software add-on packages at the same time.

unlink—(Optional) On J Series Services Routers, this option ensures that the software package is removed at the earliest opportunity in order to make room for the installation to be completed. On M Series, T Series, and MX Series routers, use the **unlink** option to remove the software package from this directory after a successful upgrade is completed.

upgrade-with-config—(Optional) Install one or more configuration files.

upgrade-with-config-format format—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

validate—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.

Additional Information



NOTE: The **request system snapshot** command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, **jkernl**, last. Add the operating system package, **jkernl**, first and the routing software package, **jroute**, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernl
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

By default, when you issue the **request system software add package-name** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, when you issue the **request system software add package-name** command on a TX Matrix Plus master Routing Engine, all the T1600 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix Plus backup Routing Engine, all the T1600 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level maintenance

Related Documentation	request system software delete on page 404
	request system software rollback on page 416
	request system storage cleanup on page 423
	Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132
	Upgrading Software on a QFabric System on page 134
List of Sample Output	request system software add validate on page 400
	request system software add (Mixed EX4200 and EX4500 Virtual Chassis) on page 401
	request system software add component all (QFabric Systems) on page 401
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software add validate

```
user@host> request system software add validate /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING: This package will load JUNOS 7.2R1.7 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
```

WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...

Sample Output

request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set  
[/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-signed.tgz  
/var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]  
...
```

request system software add component all (QFabric Systems)

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm  
component all  
...
```

request system software configuration-backup

Syntax	request system software configuration-backup (<i>path</i>)
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Save the currently active configuration and any installation-specific parameters such as a configuration that you have entered outside of the CLI, Director group IP addresses, and the default partition IP address.
Options	path —(QFabric System) Provide the path to the location of the backup configuration files. You can save the backup configuration files to either a URL, local directory, remote server, or removable drive.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none">• request system software configuration-restore on page 403• Performing a QFabric System Recovery Installation on the Director Group on page 112
List of Sample Output	request system software configuration-backup on page 402
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software configuration-backup

```
user@switch request system software configuration-backup ftp://ftp.test.net/test
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                       Dload  Upload  Total   Spent    Left
Speed
100      4035    0    0   100 4035    0    138k  --:--:-- --:--:-- --:--:--
0
```

request system software configuration-restore

Syntax	request system software configuration-restore (<i>path</i>)
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Restore a previously saved configuration and any installation-specific parameters, such as a configuration that you have entered outside of the CLI, Director group IP addresses, and the default partition IP address.
Options	path —(QFabric System) Provide the path to the location of the backup configuration files. The path can be to a local file, a file on an external flash drive, or an SCP or FTP destination.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none"> • request system software configuration-backup on page 402 • Performing a QFabric System Recovery Installation on the Director Group on page 112
List of Sample Output	request system software configuration-restore on page 403
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software configuration-restore

```

user@switch request system software configuration-restore ftp://ftp.test.net/test
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   Dload  Upload  Total    Dload  Upload    Total   Spent    Left    Speed
100 4035 100 4035    0     0  153k      0  --:--:-- --:--:-- --:--:-- 3803k

```

request system software delete

List of Syntax	Syntax on page 404 Syntax (TX Matrix Router) on page 404 Syntax (TX Matrix Plus Router) on page 404
Syntax	<code>request system software delete <i>software-package</i></code> <code><force></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code>
Syntax (TX Matrix Router)	<code>request system software delete <i>software-package</i></code> <code><force></code> <code><lcc <i>number</i> scc></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code>
Syntax (TX Matrix Plus Router)	<code>request system software delete <i>software-package</i></code> <code><force></code> <code><lcc <i>number</i> sfc <i>number</i>></code> <code><reboot></code> <code><set [<i>package-name package-name</i>]></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Option set [<i>package-name package-name</i>] added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Services Gateways. Option reboot introduced in Junos OS Release 12.3.
Description	Remove a software package or bundle from the router or switch.



CAUTION: Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.

- | | |
|----------------|---|
| Options | <p><i>software-package</i>—Software package or bundle name. You can delete any or all of the following software bundles or packages:</p> <ul style="list-style-type: none">• jbase—(Optional) Junos base software suite• jcrypto—(Optional, in domestic version only) Junos security software• jdocs—(Optional) Junos online documentation file• jkernel—(Optional) Junos kernel software suite• jpfe—(Optional) Junos Packet Forwarding Engine support |
|----------------|---|

- **jroute**—(Optional) Junos routing software suite
- **junos**—(Optional) Junos base software



NOTE: On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

force—(Optional) Ignore warnings and force removal of the software.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, remove an extension or upgrade package from a specific T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, remove an extension or upgrade package from a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software delete** command.

scc—(TX Matrix routers only) (Optional) Remove an extension or upgrade package from the TX Matrix router (or switch-card chassis).

set [package-name package-name]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc number—(TX Matrix Plus routers only) (Optional) Remove an extension or upgrade package from the TX Matrix Plus router. Replace *number* with 0.

Additional Information Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the `/altroot` and `/altconfig` file systems (on routers) or the `/`, `/altroot`, `/config`, `/var`, and `/var/tmp` file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the `/altroot` and `/altconfig` file systems (on routers) or the `/`, `/altroot`,

/config, /var, and /var/tmp file systems (on switches). After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Required Privilege Level maintenance

Related Documentation

- [request system software add on page 394](#)
- [request system software rollback on page 416](#)
- [request system software validate on page 420](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [request system software delete jdocs on page 406](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software delete jdocs](#)

The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:
```

```
Comment:
JUNOS Base OS Software Suite [7.2R1.7]
```

```
Information for jcrypto:
```

```
Comment:
JUNOS Crypto Software Suite [7.2R1.7]
```

```
Information for jdocs:
```

```
Comment:
JUNOS Online Documentation [7.2R1.7]
```

```
Information for jkernel:
```

```
Comment:
JUNOS Kernel Software Suite [7.2R1.7]
```

```
...
```

```
user@host> request system software delete jdocs
Removing package 'jdocs' ...
```

```
user@host> show system software
```

Information for jbase:

Comment:

JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [7.2R1.7]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [7.2R1.7]

...

request system software download

Syntax (QFabric System)	<code>request system software download <i>path package-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Download a software package from a location on the Director device, mounted external USB flash drive, remote FTP or SCP location, or other location.
Options	<p><i>path</i>—Location where the software package is located. For example:</p> <ul style="list-style-type: none">• <i>/pbdata/packages/package-name</i>—For a software package that is being installed from a local directory on the switch.• <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following:<ul style="list-style-type: none">• <i>ftp</i>—File Transfer Protocol. Use <i>ftp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>ftp://<username>:<password>@hostname/pathname/package-name</i>. To have the system prompt you for the password, specify <i>prompt</i> in place of the password. If a password is required, and you do not specify the password or <i>prompt</i>, an error message is displayed.• <i>scp</i>—Secure copy (available only for Canada and U.S. version). Use <i>scp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>scp://<username>:<password>@hostname/pathname/package-name</i>.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• request system software add on page 394• request system software delete on page 404• request system software rollback on page 416• request system storage cleanup on page 423• Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132• Upgrading Software on a QFabric System on page 134
List of Sample Output	request system software download on page 409
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software download

```
user@switch> request system software download
ftp://ftp.install-directory/jinstall-qfabric-11.3X30.6.rpm
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total     Spent    Left     Speed
100 186M 100 186M    0     0 18.4M      0  0:00:10  0:00:10 --:--:-- 18.6M
```

request system software nonstop-upgrade

Syntax	<code>request system software nonstop-upgrade <i>package-name</i></code> <code><fabric ></code> <code><director-group></code> <code><node-group <i>name</i>></code>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.
Options	<p><i>package-name</i>—Location from which the software is to be installed. For example:</p> <ul style="list-style-type: none">• <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following:<ul style="list-style-type: none">• ftp—File Transfer Protocol. Use <i>ftp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>ftp://<username>:<password>@hostname/pathname/package-name</i>. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed.• http—Hypertext Transfer Protocol. Use <i>http://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>http://<username>:<password>@hostname/pathname/package-name</i>. If a password is required and you omit it, you are prompted for it.• scp—Secure copy (available only for Canada and U.S. version). Use <i>scp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>scp://<username>:<password>@hostname/pathname/package-name</i>.

**NOTE:**

- The ***pathname*** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
-

director-group—Install software package on the Director group and Fabric managers.

fabric—Install software package on the Interconnect devices and Fabric controls.

node-group *name* —Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 105 • Verifying Nonstop Software Upgrade for QFabric Systems on page 1410 • show chassis nonstop-upgrade node-group on page 817
List of Sample Output	request system software nonstop-upgrade director-group on page 411 request system software nonstop-upgrade fabric on page 413 request system software nonstop-upgrade node-group (Redundant Server Node Group) on page 413 request system software nonstop-upgrade node-group (Server Node Group) on page 415
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software nonstop-upgrade director-group

```

user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade

```

```
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
Setting peer DG node as the master SFC
```

```
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0 2012-06-05 16:12:43]: Retrieving package
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
```



```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0 2012-06-05 16:48:53]: Retrieving package
[FC-1 2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0 2012-06-05 16:49:51]: ----- re0: -----
[FC-1 2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition
[Status 2012-06-05 16:58:34]: Load completed with 0 errors...
[Status 2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status 2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status 2012-06-05 16:58:39]: Rebooting FC-0
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status 2012-06-05 16:58:44]: Rebooting FC-1
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): RSNG

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
```

```
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last group
:1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed
```

```
[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): P1507-C

[P1507-C 2012-06-26 14:02:44]: Retrieving package
[P1507-C 2012-06-26 14:03:21]: ----- P1507-C: -----
[P1507-C 2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C 2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C 2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C 2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C 2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C 2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C 2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C 2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
[P1507-C 2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C 2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C 2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C 2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C 2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C 2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C 2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C 2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C 2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C 2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in
installation
[P1507-C 2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[P1507-C 2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status 2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status 2012-06-26 14:10:32]: Rebooting P1507-C
[Status 2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```

request system software rollback

List of Syntax	Syntax on page 416 Syntax (EX Series Switches) on page 416 Syntax (TX Matrix Router) on page 416 Syntax (TX Matrix Plus Router) on page 416 Syntax (MX Series Router) on page 416
Syntax	request system software rollback
Syntax (EX Series Switches)	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
Syntax (TX Matrix Router)	request system software rollback <lcc <i>number</i> scc> <reboot>
Syntax (TX Matrix Plus Router)	request system software rollback <lcc <i>number</i> sfc <i>number</i> > <reboot>
Syntax (MX Series Router)	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command behavior changed in Junos OS Release 12.1. Option reboot introduced in Junos OS Release 12.3.
Description	<p>For all versions of Junos OS up to and including Junos OS 11.4, revert to the software that was loaded at the last successful request system software add command.</p> <p>As of Junos OS 12.1 and greater, revert to the last known good state before the most recent request system software (add delete) command. For example, using rollback in Junos OS 12.1 after using request system software add restores the system to a known good state prior to using the add command. Similarly, using rollback in Junos OS 12.1 after using request system software delete restores the system to a known good state prior to using the delete command.</p> <p>A software rollback fails if any required package (or a bundle package containing the required package) cannot be found in /var/sw/pkg.</p> <p><i>Additional Information</i></p>

- On M Series and T Series routers, if **request system software add <jinstall> reboot** was used for the previous installation, then **request system software rollback** has no effect. In this case, use **jinstall** to reinstall the required package.
- On M Series and T Series routers, if **request system software add <sdk1>** was used for the previous installation, then **request system software rollback** removes the last installed SDK package (**sdk1** in this example).
- On SRX Series devices with dual root systems, when **request system software rollback** is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Rollback takes each root back to the previously installed image.

Options **all-members**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, attempt to roll back to the previous set of packages on a T640 router connected to the TX Matrix router. On a TX Matrix Plus router, attempt to roll back to the previous set of packages on a connected router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

member member-id—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

none—For all versions of Junos OS up to and including Junos OS 11.4, revert to the set of software as of the last successful **request system software add**. As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command.

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software rollback** command.

scc—(TX Matrix routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level maintenance

Related Documentation

- *request system software abort*
- [request system software add on page 394](#)
- [request system software delete on page 404](#)
- [request system software validate on page 420](#)
- [request system configuration rescue delete on page 373](#)
- [request system configuration rescue save on page 374](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [request system software rollback on page 419](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software rollback

```

user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host

```

request system software validate

List of Syntax	Syntax on page 420 Syntax (TX Matrix Router) on page 420 Syntax (TX Matrix Plus Router) on page 420 Syntax (MX Series Router) on page 420
Syntax	<code>request system software validate <i>package-name</i></code> <code><set [<i>package-name package-name</i>]></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code>
Syntax (TX Matrix Router)	<code>request system software validate <i>package-name</i></code> <code><lcc <i>number</i> scc></code> <code><set [<i>package-name package-name</i>]></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code>
Syntax (TX Matrix Plus Router)	<code>request system software validate <i>package-name</i></code> <code><lcc <i>number</i> sfc <i>number</i>></code> <code><set [<i>package-name package-name</i>]></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code>
Syntax (MX Series Router)	<code>request system software validate <i>package-name</i></code> <code><member <i>member-id</i>></code> <code><set [<i>package-name package-name</i>]></code> <code><upgrade-with-config></code> <code><upgrade-with-config-format <i>format</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. set [<i>package-name package-name</i>] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways. upgrade-with-config and upgrade-with-config-format <i>format</i> options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.
Description	Validate candidate software against the current configuration of the router.
Options	lcc <i>number</i> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package for a specific router that is connected to the TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none">• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

member *member-id*—(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

package-name—Name of the software bundle or package to test.

scc—(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).

set [*package-name package-name*]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc *number*—(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router.

upgrade-with-config—(Optional) Install one or more configuration files.

upgrade-with-config-format *format*—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

Additional Information By default, when you issue the **request system software validate** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the **request system software validate** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>request system software abort</i>• request system software add on page 394• request system software delete on page 404• request system software rollback on page 416• <i>Routing Matrix with a TX Matrix Plus Router Solutions Page</i>
List of Sample Output	request system software validate (Successful Case) on page 422 request system software validate (Failure Case) on page 422
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software validate (Successful Case)

```
user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

request system software validate (Failure Case)

```
user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0

user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```

request system storage cleanup

List of Syntax	Syntax on page 423 Syntax (EX Series Switches) on page 423 Syntax (MX Series Router) on page 423 Syntax (QFX Series) on page 423
Syntax	request system storage cleanup <dry-run>
Syntax (EX Series Switches)	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
Syntax (MX Series Router)	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
Syntax (QFX Series)	request system storage cleanup <component (<i>serial number</i> <i>UUID</i> all)> <director-group <i>name</i> > <dry-run> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <name-tag <i>name-tag</i> > <node-group <i>name</i> > <prune> <qfabric (component <i>name</i>) dry-run name-tag repository> <repository (core log)>
Release Information	Command introduced in Junos OS Release 7.4. dry-run option introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Free storage space on the router or switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion. On a QFabric system, you can delete debug files located on individual devices or on the entire QFabric system.
Options	<p>all-members—(EX4200 switches and MX Series routers only) (Optional) Delete files on all members of the Virtual Chassis configuration.</p> <p>component (<i>UUID</i> <i>serial number</i> all)—(QFabric systems only) (Optional) Delete files located on individual QFabric system devices or on the entire QFabric system.</p> <p>director-group <i>name</i>—(QFabric systems only) (Optional) Delete files on the Director group.</p> <p>dry-run—(Optional) List files proposed for deletion (without deleting them).</p>

infrastructure *name*—(QFabric systems only) (Optional) Delete files on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Delete files on the Interconnect device.

local—(EX4200 switches and MX Series routers only) (Optional) Delete files on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

name-tag *name-tag*—(QFabric systems only) (Optional) Delete debug files that match a specific regular expression.

node-group *name*—(QFabric systems only) (Optional) Delete files on the Node group.

prune—(QFabric systems only) (Optional) Delete debug files located in either the core or log debug repositories of a QFabric system device.

qfabric component *name*—(QFabric systems only) (Optional) Delete debug files located in the debug repositories of a QFabric system device.

repository (*core* | *log*)—(QFabric systems only) (Optional) Specify the repository on the QFabric system device for which you want to delete debug files.

Additional Information If logging is configured and being used, the **dry-run** option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

Required Privilege Level maintenance

List of Sample Output [request system storage cleanup dry-run on page 425](#)
[request system storage cleanup on page 426](#)
[request system storage cleanup director-group \(QFabric Systems\) on page 426](#)
[request system storage cleanup infrastructure device-name \(QFabric Systems\) on page 428](#)
[request system storage cleanup interconnect-device device-name \(QFabric Systems\) on page 429](#)
[request system storage cleanup node-group group-name \(QFabric Systems\) on page 430](#)
[request system storage cleanup qfabric component device-name \(QFabric Systems\) on page 431](#)
[request system storage cleanup qfabric component device-name repository core \(QFabric Systems\) on page 431](#)
[request system storage cleanup qfabric component all \(QFabric Systems\) on page 431](#)

Output Fields [Table 49 on page 425](#) describes the output fields for the **request system storage cleanup** command. Output fields are listed in the approximate order in which they appear.

Table 49: request system storage cleanup Output Fields

Field Name	Field Description
List of files to delete:	Shows list of files available for deletion.
Size	Size of the core-dump file.
Date	Last core-dump file modification date and time.
Name	Name of the core-dump file.
Directory to delete:	Shows list of directories available for deletion.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the core repository, and the log files are located in the log repository. The default Repository scope is shared since both the core and log repositories are shared by all of the QFabric system devices.
Repository head:	Name of the top-level repository location.
Repository name:	Name of the repository: core or log .
Creating list of debug artifacts to be removed under:	Shows location of files available for deletion.
List of debug artifacts to be removed under:	Shows list of files available for deletion.

Sample Output

request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

Size	Date	Name
11.4K	Mar 8 15:00	/var/log/messages.1.gz
7245B	Feb 5 15:00	/var/log/messages.3.gz
11.8K	Feb 22 13:00	/var/log/messages.2.gz
3926B	Mar 16 13:57	/var/log/messages.0.gz
3962B	Feb 22 12:47	/var/log/sampled.1.gz
4146B	Mar 8 12:20	/var/log/sampled.0.gz
4708B	Dec 21 11:39	/var/log/sampled.2.gz
7068B	Jan 16 18:00	/var/log/messages.4.gz
13.7K	Dec 27 22:00	/var/log/messages.5.gz
890B	Feb 22 17:22	/var/tmp/sampled.pkts
65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz

request system storage cleanup

```
user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

Size	Date	Name
11.4K	Mar 8 15:00	/var/log/messages.1.gz
7245B	Feb 5 15:00	/var/log/messages.3.gz
11.8K	Feb 22 13:00	/var/log/messages.2.gz
3926B	Mar 16 13:57	/var/log/messages.0.gz
11.6K	Mar 8 15:00	/var/log/messages.5.gz
7254B	Feb 5 15:00	/var/log/messages.6.gz
12.9K	Feb 22 13:00	/var/log/messages.8.gz
3726B	Mar 16 13:57	/var/log/messages.7.gz
3962B	Feb 22 12:47	/var/log/sampled.1.gz
4146B	Mar 8 12:20	/var/log/sampled.0.gz
4708B	Dec 21 11:39	/var/log/sampled.2.gz
7068B	Jan 16 18:00	/var/log/messages.4.gz
13.7K	Dec 27 22:00	/var/log/messages.5.gz
890B	Feb 22 17:22	/var/tmp/sampled.pkts
65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz

Delete these files ? [yes,no] (yes)

request system storage cleanup director-group (QFabric Systems)

```
user@switch> request system storage cleanup director-group
List of files to delete:
```

Size	Date	Name
4.0K	2011-11-07 05:16:29	/tmp/2064.sfcauth
4.0K	2011-11-07 05:07:34	/tmp/30804.sfcauth
4.0K	2011-11-07 04:13:41	/tmp/26792.sfcauth
4.0K	2011-11-07 04:13:39	/tmp/26432.sfcauth
0	2011-11-07 07:45:40	/tmp/cluster_cleanup.log
1.3M	2011-11-07 07:39:11	/tmp/cn_monitor.20111107-052401.log
4.0K	2011-11-07 07:36:29	/tmp/clustat.28019.log
4.0K	2011-11-07 07:36:29	/tmp/clustat_x.28019.log
9.6M	2011-11-07 05:30:24	/tmp/sfc.2.log
4.0K	2011-11-07 05:28:11	/tmp/mgd-init.1320672491.log
248K	2011-11-07 05:19:24	/tmp/cn_monitor.20111107-045111.log
4.0K	2011-11-07 05:17:18	/tmp/clustat.3401.log
4.0K	2011-11-07 05:17:18	/tmp/clustat_x.3401.log
8.0K	2011-11-07 04:58:25	/tmp/mgd-init.1320670633.log
0	2011-11-07 04:54:01	/tmp/mysql_db_install_5.1.37.log
4.0K	2011-11-07 04:52:08	/tmp/cn_send.log
0	2011-11-07 04:52:00	/tmp/init_eth0.log
4.0K	2011-11-07 04:49:35	/tmp/install_interfaces.sh.log
4.0K	2011-11-07 04:48:15	/tmp/bootstrap.sh.log
160K	2011-11-07 04:47:43	/tmp/bootstrap_cleanup.log
38M	2011-11-07 04:42:42	/tmp/cn_monitor.20111104-110308.log
4.0K	2011-11-07 04:38:47	/tmp/clustat.30913.log
4.0K	2011-11-07 04:38:47	/tmp/clustat_x.30913.log
4.0K	2011-11-07 04:38:03	/tmp/dcf_upgrade.sh.remove.log
4.0K	2011-11-07 04:38:03	/tmp/peer_update.log
4.0K	2011-11-07 04:38:02	/tmp/dcf_upgrade.log
4.0K	2011-11-07 04:38:02	/tmp/perl_mark_upgrade.log

```

8.0K  2011-11-07 04:13:42 /tmp/install_dcf_rpm.log
4.0K  2011-11-07 04:13:06 /tmp/00_cleanup.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/ccif_patch_4410_4450.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/pcf-tools.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/initial.sh.1320667986.log
0      2011-11-07 04:13:06 /tmp/inventory.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/qf-db.sh.1320667986.log
4.0K  2011-11-07 04:13:06 /tmp/sfc.sh.1320667986.log
8.0K  2011-11-07 04:13:05 /tmp/jinstall-qfabric.log
8.0K  2011-11-04 11:10:24 /tmp/mgd-init.1320430192.log
4.0K  2011-11-04 11:07:03 /tmp/mysql_dcf_db_install.log
8.0K  2011-11-04 10:55:07 /tmp/ccif_patch_4410_4450.sh.1320429307.log
8.0K  2011-11-04 10:55:07 /tmp/initial.sh.1320429307.log
4.0K  2011-11-04 10:55:07 /tmp/inventory.sh.1320429307.log
8.0K  2011-11-04 10:55:07 /tmp/sfc.sh.1320429307.log
4.0K  2011-11-04 10:54:09 /tmp/ks-script-Ax0tz5.log
4.0K  2011-11-07 04:13:06 /tmp//sfc.sh.1320667986.log
8.0K  2011-11-04 10:55:07 /tmp//sfc.sh.1320429307.log

```

Directory to delete:

```

45M   2011-11-08 10:57:43 /tmp/sfc-captures

```

List of files to delete:

	Size	Date	Name
4.0K	2011-11-08 05:47:47	/tmp/5713.sfcauth	
4.0K	2011-11-08 05:14:32	/tmp/14494.sfcauth	
4.0K	2011-11-08 05:11:47	/tmp/9978.sfcauth	
4.0K	2011-11-08 05:09:37	/tmp/6128.sfcauth	
4.0K	2011-11-08 05:04:28	/tmp/29703.sfcauth	
4.0K	2011-11-07 11:59:10	/tmp/7811.sfcauth	
4.0K	2011-11-07 11:36:08	/tmp/32415.sfcauth	
4.0K	2011-11-07 11:30:30	/tmp/22406.sfcauth	
4.0K	2011-11-07 11:24:37	/tmp/12131.sfcauth	
4.0K	2011-11-07 10:48:42	/tmp/12687.sfcauth	
4.0K	2011-11-07 09:27:20	/tmp/31082.sfcauth	
4.0K	2011-11-07 07:33:58	/tmp/14633.sfcauth	
4.0K	2011-11-07 05:08:25	/tmp/15447.sfcauth	
4.0K	2011-11-07 04:12:29	/tmp/26874.sfcauth	
4.0K	2011-11-07 04:12:27	/tmp/26713.sfcauth	
4.0K	2011-11-07 03:49:17	/tmp/17691.sfcauth	
4.0K	2011-11-05 01:32:23	/tmp/5716.sfcauth	
4.0K	2011-11-07 08:00:17	/tmp/sfcsnmpd.log	
4.0K	2011-11-07 07:57:50	/tmp/cluster_cleanup.log	
824K	2011-11-07 07:38:37	/tmp/cn_monitor.20111107-053643.log	
4.0K	2011-11-07 07:36:30	/tmp/clustat.18399.log	
4.0K	2011-11-07 07:36:30	/tmp/clustat_x.18399.log	
4.0K	2011-11-07 07:35:47	/tmp/command_lock.log	
4.0K	2011-11-07 05:39:54	/tmp/mgd-init.1320673194.log	
92K	2011-11-07 05:19:25	/tmp/cn_monitor.20111107-050412.log	
4.0K	2011-11-07 05:17:20	/tmp/clustat.30115.log	
4.0K	2011-11-07 05:17:20	/tmp/clustat_x.30115.log	
8.0K	2011-11-07 05:08:07	/tmp/mgd-init.1320671241.log	
4.0K	2011-11-07 05:04:57	/tmp/cn_send.log	
0	2011-11-07 05:04:52	/tmp/init_eth0.log	
4.0K	2011-11-07 05:02:38	/tmp/install_interfaces.sh.log	
4.0K	2011-11-07 05:01:19	/tmp/bootstrap.sh.log	
160K	2011-11-07 05:00:47	/tmp/bootstrap_cleanup.log	
28M	2011-11-07 04:42:27	/tmp/cn_monitor.20111104-112954.log	
4.0K	2011-11-07 04:38:49	/tmp/clustat.6780.log	
4.0K	2011-11-07 04:38:49	/tmp/clustat_x.6780.log	

```

4.0K  2011-11-07 04:38:05 /tmp/issue_event.log
4.0K  2011-11-07 04:38:05 /tmp/peer_upgrade_reboot.log
12K   2011-11-07 04:38:05 /tmp/primary_update.log
4.0K  2011-11-07 04:38:04 /tmp/dcf_upgrade.sh.remove.log
4.0K  2011-11-07 04:38:04 /tmp/peer_rexec_upgrade.log
4.0K  2011-11-07 04:13:42 /tmp/peer_install_dcf_rpm.log
4.0K  2011-11-07 04:11:57 /tmp/dcf-tools.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/initial.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/inventory.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/qf-db.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/sfc.sh.1320667917.log
4.0K  2011-11-07 04:11:56 /tmp/00_cleanup.sh.1320667916.log
0     2011-11-07 04:11:56 /tmp/ccif_patch_4410_4450.sh.1320667916.log
8.0K  2011-11-07 04:11:56 /tmp/jinstall-qfabric.log
4.0K  2011-11-07 04:11:33 /tmp/dcf_upgrade.log
8.0K  2011-11-04 11:53:12 /tmp/mgd-init.1320432782.log
8.0K  2011-11-04 11:06:17 /tmp/ccif_patch_4410_4450.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/initial.sh.1320429977.log
4.0K  2011-11-04 11:06:17 /tmp/inventory.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/sfc.sh.1320429977.log
4.0K  2011-11-04 11:05:19 /tmp/ks-script-_tnWeb.log
4.0K  2011-11-07 04:11:57 /tmp//sfc.sh.1320667917.log
8.0K  2011-11-04 11:06:17 /tmp//sfc.sh.1320429977.log

```

Directory to delete:

```
49M   2011-11-08 10:45:20 /tmp/sfc-captures
```

request system storage cleanup infrastructure device-name (QFabric Systems)

```
user@switch> request system storage cleanup infrastructure FC-0
re0:
```

List of files to delete:

Size	Date	Name
139B	Nov 8 19:03	/var/log/default-log-messages.0.gz
5602B	Nov 8 19:03	/var/log/messages.0.gz
28.4K	Nov 8 10:15	/var/log/messages.1.gz
35.2K	Nov 7 13:45	/var/log/messages.2.gz
207B	Nov 7 16:02	/var/log/wtmp.0.gz
27B	Nov 7 12:14	/var/log/wtmp.1.gz
184.4M	Nov 7 12:16	/var/sw/pkg/jinstall-dc-re-11.3I20111104_1216_dc-builder-domestic-signed.tgz
124.0K	Nov 7 15:59	/var/tmp/gres-tp/env.dat
0B	Nov 7 12:57	/var/tmp/gres-tp/lock
155B	Nov 7 16:02	/var/tmp/krt_gencfg_filter.txt
0B	Nov 7 12:35	/var/tmp/last_ccif_update
1217B	Nov 7 12:15	/var/tmp/loader.conf.preinstall
184.4M	Nov 6 07:11	/var/tmp/mchassis-install.tgz
10.8M	Nov 7 12:16	/var/tmp/preinstall/bootstrap-install-11.3I20111104_1216_dc-builder.tar
57.4K	Nov 7 12:16	/var/tmp/preinstall/configs-11.3I20111104_1216_dc-builder.tgz
259B	Nov 7 12:16	/var/tmp/preinstall/install.conf
734.3K	Nov 4 13:46	/var/tmp/preinstall/jboot-dc-re-11.3I20111104_1216_dc-builder.tgz
177.8M	Nov 7 12:16	/var/tmp/preinstall/jbundle-dc-re-11.3I20111104_1216_dc-builder-domestic.tgz
124B	Nov 7 12:15	/var/tmp/preinstall/metatags


```
1217B Nov 7 12:16 /var/tmp/preinstall_boot_loader.conf
0B Nov 7 16:02 /var/tmp/rtsdb/if-rtsdb
```

request system storage cleanup interconnect-device device-name (QFabric Systems)

```
user@switch> request system storage cleanup interconnect IC-WS001
re1:
```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	128B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	9965B	Nov 8 19:06	/var/log/messages.0.gz
	15.8K	Nov 8 12:30	/var/log/messages.1.gz
	15.8K	Nov 8 11:00	/var/log/messages.2.gz
	15.7K	Nov 8 07:30	/var/log/messages.3.gz
	15.8K	Nov 8 04:00	/var/log/messages.4.gz
	15.7K	Nov 8 00:30	/var/log/messages.5.gz
	18.7K	Nov 7 21:00	/var/log/messages.6.gz
	17.6K	Nov 7 19:00	/var/log/messages.7.gz
	58.3K	Nov 7 16:00	/var/log/messages.8.gz
	20.3K	Nov 7 15:15	/var/log/messages.9.gz
	90B	Nov 7 15:41	/var/log/wtmp.0.gz
	57B	Nov 7 12:41	/var/log/wtmp.1.gz
	124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:41	/var/tmp/if-rtsdb/env.lck
	12.0K	Nov 7 15:41	/var/tmp/if-rtsdb/env.mem
	132.0K	Nov 7 15:55	/var/tmp/if-rtsdb/shm_usr1.mem
	2688.0K	Nov 7 15:41	/var/tmp/if-rtsdb/shm_usr2.mem
	2048.0K	Nov 7 15:41	/var/tmp/if-rtsdb/trace.mem
	730B	Nov 7 19:57	/var/tmp/juniper.conf+.gz
	155B	Nov 7 15:53	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 15:41	/var/tmp/rtsdb/if-rtsdb

```
re0:
```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	121B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	16.7K	Nov 8 19:06	/var/log/messages.0.gz
	22.2K	Nov 8 17:45	/var/log/messages.1.gz
	18.4K	Nov 8 17:00	/var/log/messages.2.gz
	21.6K	Nov 8 16:00	/var/log/messages.3.gz
	17.9K	Nov 8 14:30	/var/log/messages.4.gz
	19.4K	Nov 8 13:30	/var/log/messages.5.gz
	18.2K	Nov 8 12:30	/var/log/messages.6.gz
	20.4K	Nov 8 11:30	/var/log/messages.7.gz
	21.4K	Nov 8 10:15	/var/log/messages.8.gz
	21.0K	Nov 8 09:00	/var/log/messages.9.gz
	19.9K	Nov 8 08:13	/var/log/snmp-traps.0.gz
	203B	Nov 8 15:36	/var/log/wtmp.0.gz
	57B	Nov 7 12:41	/var/log/wtmp.1.gz
	124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:41	/var/tmp/if-rtsdb/env.lck

```

12.0K Nov 7 15:41 /var/tmp/if-rtbdb/env.mem
132.0K Nov 7 15:55 /var/tmp/if-rtbdb/shm_usr1.mem
2688.0K Nov 7 15:41 /var/tmp/if-rtbdb/shm_usr2.mem
2048.0K Nov 7 15:41 /var/tmp/if-rtbdb/trace.mem
727B Nov 7 15:54 /var/tmp/juniper.conf+.gz
155B Nov 7 15:55 /var/tmp/krt_gencfg_filter.txt
0B Nov 7 15:41 /var/tmp/rtbdb/if-rtbdb

```

request system storage cleanup node-group group-name (QFabric Systems)

```

user@switch> request system storage cleanup node-group NW-NG-0
BBAK0372:

```

List of files to delete:

Size	Date	Name
126B	Nov 8 19:07	/var/log/default-log-messages.0.gz
179B	Nov 7 13:32	/var/log/install.0.gz
22.9K	Nov 8 19:07	/var/log/messages.0.gz
26.5K	Nov 8 17:30	/var/log/messages.1.gz
20.5K	Nov 8 13:15	/var/log/messages.2.gz
33.2K	Nov 7 17:45	/var/log/messages.3.gz
35.5K	Nov 7 15:45	/var/log/messages.4.gz
339B	Nov 8 17:10	/var/log/wtmp.0.gz
58B	Nov 7 12:40	/var/log/wtmp.1.gz
124.0K	Nov 8 17:08	/var/tmp/gres-tp/env.dat
0B	Nov 7 12:39	/var/tmp/gres-tp/lock
0B	Nov 7 12:59	/var/tmp/if-rtbdb/env.lck
12.0K	Nov 8 17:09	/var/tmp/if-rtbdb/env.mem
2688.0K	Nov 8 17:09	/var/tmp/if-rtbdb/shm_usr1.mem
132.0K	Nov 8 17:09	/var/tmp/if-rtbdb/shm_usr2.mem
2048.0K	Nov 8 17:09	/var/tmp/if-rtbdb/trace.mem
1082B	Nov 8 17:09	/var/tmp/juniper.conf+.gz
155B	Nov 7 17:39	/var/tmp/krt_gencfg_filter.txt
0B	Nov 8 17:09	/var/tmp/rtbdb/if-rtbdb

EE3093:

List of files to delete:

Size	Date	Name
11B	Nov 8 17:33	/var/jail/tmp/alarmd.ts
119B	Nov 8 19:08	/var/log/default-log-messages.0.gz
180B	Nov 7 17:41	/var/log/install.0.gz
178B	Nov 7 13:32	/var/log/install.1.gz
2739B	Nov 8 19:08	/var/log/messages.0.gz
29.8K	Nov 8 18:45	/var/log/messages.1.gz
31.8K	Nov 8 17:15	/var/log/messages.2.gz
20.6K	Nov 8 16:00	/var/log/messages.3.gz
15.4K	Nov 8 10:15	/var/log/messages.4.gz
15.4K	Nov 8 02:15	/var/log/messages.5.gz
25.5K	Nov 7 20:45	/var/log/messages.6.gz
48.0K	Nov 7 17:45	/var/log/messages.7.gz
32.8K	Nov 7 13:45	/var/log/messages.8.gz
684B	Nov 8 17:02	/var/log/wtmp.0.gz
58B	Nov 7 12:40	/var/log/wtmp.1.gz
124.0K	Nov 7 17:34	/var/tmp/gres-tp/env.dat
0B	Nov 7 12:40	/var/tmp/gres-tp/lock
0B	Nov 7 12:59	/var/tmp/if-rtbdb/env.lck

```

12.OK Nov 7 17:39 /var/tmp/if-rtbdb/env.mem
2688.OK Nov 7 17:39 /var/tmp/if-rtbdb/shm_usr1.mem
132.OK Nov 7 17:40 /var/tmp/if-rtbdb/shm_usr2.mem
2048.OK Nov 7 17:39 /var/tmp/if-rtbdb/trace.mem
155B Nov 7 17:40 /var/tmp/krt_gencfg_filter.txt
0B Nov 7 17:39 /var/tmp/rtbdb/if-rtbdb

```

request system storage cleanup qfabric component device-name (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component A0001/YA0197
Repository type: regular
Repository head: /pbstorage
Creating list of debug artifacts to be removed under:
/pbstorage/rdumps/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.0.0.05162011123308.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.1.0.05162011123614.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.2.0.05162011123920.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/livekcore.05132011163930.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/tnetd.core.0.1057.05162011124500.gz ...
done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.05132011120528.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.kz ... done
Creating list of debug artifacts to be removed under: /pbstorage/rlogs/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rlogs/A0001/YA0197/kdumpinfo.05132011120528 ... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.0.1039.05122011234415.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.1.1039.05132011175544.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/tnetd.tarball.0.1057.05162011175453.tgz
... done

```

request system storage cleanup qfabric component device-name repository core (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component EE3093 repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps/EE3093
NOTE: core repository under /pbdata/export/rdumps/EE3093 empty

```

request system storage cleanup qfabric component all (QFabric Systems)


```

user@switch> request system storage cleanup qfabric component all
Repository scope: shared
Repository head: /pbdata/export
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps
NOTE: core repository under /pbdata/export/rdumps/all empty
Creating list of debug artifacts to be removed under: /pbdata/export/rlogs
List of debug artifacts to clean up ... (press control C to abort)
/pbdata/export/rlogs/73747cd8-0710-11e1-b6a4-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/77116f18-0710-11e1-a2a0-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/BBAK0372/install-11072011121538.log
/pbdata/export/rlogs/BBAK0394/install-11072011121532.log
/pbdata/export/rlogs/EE3093/install-11072011121536.log
/pbdata/export/rlogs/WS001/YN5999/install-11072011121644.log
/pbdata/export/rlogs/WS001/YW3803/install-11072011122429.log
/pbdata/export/rlogs/cd78871a-0710-11e1-878e-00e081c5297e/install-11072011125932.log
/pbdata/export/rlogs/d0afdale-0710-11e1-a1d0-00e081c5297e/install-11072011125930.log
/pbdata/export/rlogs/d0afdale-0710-11e1-a1d0-00e081c5297e/install-11072011133211.log
/pbdata/export/rlogs/d0afdale-0710-11e1-a1d0-00e081c5297e/install-11072011155302.log

```

/pbdata/export/rlogs/d31ab7a6-0710-11e1-ad1b-00e081c5297e/install-11072011125931.log
/pbdata/export/rlogs/d4d0f254-0710-11e1-90c3-00e081c5297e/install-11072011125932.log

request system zeroize

Syntax	request system zeroize <media>
Release Information	<p>Command introduced before Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 11.2 for EX Series switches.</p> <p>Option media added in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.2 for MX Series devices.</p> <p>Command introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p> NOTE: The media option is not available on the QFX Series.</p> <p>Remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.</p> <p>This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as root and start the Junos OS command-line interface (CLI) by typing cli at the prompt.</p> <p>To completely erase user-created data so that it is unrecoverable, use the media option.</p>
Options	<p>media—(Optional) In addition to removing all configuration and log files, the media option causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and the like. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the request system zeroize media operation can take considerably more time than the request system zeroize operation. However, the critical security parameters are all removed at the beginning of the process.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> <i>request system snapshot</i> request system snapshot on page 393 <i>Reverting to the Default Factory Configuration for the EX Series Switch</i> <i>Reverting to the Rescue Configuration for the EX Series Switch</i> Reverting to the Default Factory Configuration on page 173

- [Reverting to the Rescue Configuration on page 175](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 174](#)

List of Sample Output [request system zeroize on page 434](#)
[request system zeroize media on page 435](#)

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD: Version 6.0 (0x85)
DRAM: Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@juniper.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC
```

```

user@juniper.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
  JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...

```

request system zeroize media

```

user@host> request system zeroize media
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

warning: ipsec-key-management subsystem not running - not needed by configuration.
warning: zeroizing fpc0

{master:0}
root> Waiting (max 60 seconds) for system process `vnlru' to stop...done
...
Syncing disks, vnodes remaining...2 4 2 4 3 2 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 14m50s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC

```

```
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080<EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s2a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20055 free (31 frags, 2503 blocks, 0.0% fragmentation)
zeroizing /dev/da0s1a ...
. . .
zeroizing /dev/da0s3d ...
. . .
zeroizing /dev/da0s3e ...
. . .
zeroizing /dev/da0s4d ...
. . .
zeroizing /dev/da0s4e ...
. . .

syncing disks... All buffers synced.
Uptime: 3m40s
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
```



```

Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080 <EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s1a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20064 free (48 frags, 2502 blocks, 0.1% fragmentation)
zeroizing /dev/da0s2a ...
. . .
Creating initial configuration...mgd: error: Cannot open configuration file:
/config/juniper.conf
mgd: warning: activating factory configuration
mgd: commit complete
mgd: -----
mgd: Please login as 'root'. No password is required.
mgd: To start Initial Setup, type 'ezsetup' at the JUNOS prompt.
mgd: To start JUNOS CLI, type 'cli' at the JUNOS prompt.
mgd: -----
Setting initial options: debugger_on_panic=NO debugger_on_break=NO.
Starting optional daemons: .
Doing initial network setup:
. . .

Amnesiac (ttyu0)

```

restart

List of Syntax Syntax on page 438

Syntax (ACX Series Routers) on page 438

Syntax (EX Series Switches) on page 438

Syntax (Routing Matrix) on page 439

Syntax (J Series Routing Platform) on page 439

Syntax (TX Matrix Routers) on page 439

Syntax (TX Matrix Plus Routers) on page 439

Syntax (MX Series Routers) on page 439

Syntax (J Series Routers) on page 440

Syntax (QFX Series) on page 440

Syntax restart

```
<adaptive-services | ancpd-service | application-identification | audit-process |
auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
class-of-service | clksyncd-service | database-replication | datapath-trace-service
| dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
ecc-error-logging | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | iccp-service | idp-policy | immediately
| interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
| l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
| local-policy-decision-function | mac-validation | mib-process | mobile-ip | mountd-service
| mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
packet-triggered-subscribers | peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations |
root-system-domain-service | routing <logical-system logical-system-name> | sampling
| sbc-configuration-process | sdk-service | service-deployment | services | services pgcp
gateway gateway-name | snmp | soft | static-subscribers | statistics-service |
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
vrrp | web-management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

restart

```
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
| disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | immediately | interface-control |
ipsec-key-management | l2-learning | lacp | link-management | mib-process | mobile-ip |
mountd-service | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service
| ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
sdk-service | secure-neighbor-discovery | service-deployment | services | snmp | soft
| statistics-service | subscriber-management | subscriber-management-helper | tunnel-oamd
| vrrp>
```

Syntax (EX Series Switches)

restart

```
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
ethernet-switching | event-processing | firewall | general-authentication-service |
interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
```

	lldpd-service mib-process mounstd-service multicast-snooping pgm redundancy-interface-process remote-operations routing secure-neighbor-discovery service-deployment sflow-service snmp vrrp web-management>
Syntax (Routing Matrix)	restart <adaptive-services audit-process chassis-control class-of-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp> <all all-lcc lcc <i>number</i> > <gracefully immediately soft>
Syntax (J Series Routing Platform)	restart <adaptive-services audit-process chassis-control class-of-service dhcp dialer-services dlsr event-processing firewall interface-control ipsec-key-management isdn-signaling l2-learning l2tp-service mib-process network-access-service pgm ppp pppoe remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp usb-control web-management> <gracefully immediately soft>
Syntax (TX Matrix Routers)	restart <adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp statistics-service> <all-chassis all-lcc lcc <i>number</i> scc> <gracefully immediately soft>
Syntax (TX Matrix Plus Routers)	restart <adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp statistics-service> <all-chassis all-lcc all-sfc lcc <i>number</i> sfc <i>number</i> > <gracefully immediately soft>
Syntax (MX Series Routers)	restart <adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mobile-ip mounstd-service mpls-traceroute mspd multicast-snooping named-service nfsd-service

```
packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system logical-system-name> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway gateway-name |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

**Syntax (J Series
Routers)**

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsr | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsr | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |logical-system-name> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsr** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description Restart a Junos OS process.



CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dlsw—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lACP—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG,

and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management—(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function—(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation—(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc number—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with **0**.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation [• Overview of Junos OS CLI Operational Mode Commands on page 75](#)

List of Sample Output [restart interfaces on page 447](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```


rollback

Syntax	<code>rollback <number rescue></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Return to a previously committed configuration. The software saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the commit configuration command.</p> <p>The currently operational Junos OS configuration is stored in the file juniper.conf, and the last three committed configurations are stored in the files juniper.conf.1, juniper.conf.2, and juniper.conf.3. These four files are located in the directory /config, which is on the router's flash drive. The remaining 46 previous committed configurations, the files juniper.conf.4 through juniper.conf.49, are stored in the directory /var/db/config, which is on the router's hard disk.</p> <p>During rollback, the configuration you specify is loaded from the associated file. Only objects in the rollback configuration that differ from the previously loaded configuration are marked as changed (equivalent to load update).</p>
Options	<p>none (Optional)—Return to the most recently saved configuration.</p> <p>number—(Optional) Configuration to return to. The range of values is from 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. The default is 0.</p> <p>rescue—(Optional) Return to the rescue configuration.</p>
Required Privilege Level	rollback—To roll back to configurations other than the one most recently committed.
Related Documentation	<ul style="list-style-type: none"> • Returning to a Previously Committed Junos OS Configuration on page 1601 • Creating and Returning to a Rescue Configuration on page 1596

save

Syntax	<code>save <i>filename</i></code>
QFX Series	<code>save (dhcp-snooping <i>filename</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.</p> <p>When saving a file to a remote system, the software uses the scp/ssh protocol.</p>
Options	<p><i>filename</i>—Name of the saved file. You can specify a filename in one of the following ways:</p> <ul style="list-style-type: none">• <i>filename</i>—File in the user's home directory (the current directory) on the local flash drive.• <i>path/filename</i>—File on the local flash drive.• <i>/var/filename</i> or <i>/var/path/filename</i>—File on the local hard disk.• <i>a:filename</i> or <i>a:path/filename</i>—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.• <i>hostname:/path/filename</i>, <i>hostname:filename</i>, <i>hostname:path/filename</i>, or <i>scp://hostname/path/filename</i>—File on an scp/ssh client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify <i>hostname</i> as <i>username@hostname</i>.• <i>ftp://hostname/path/filename</i>—File on an FTP server. You can also specify <i>hostname</i> as <i>username @hostname</i> or <i>username:password @hostname</i>. The default path is the user's home directory. To specify an absolute path, the path must start with the string %2F; for example, <i>ftp://hostname/%2Fpath/filename</i>. To have the system prompt you for the password, specify <i>prompt</i> in place of the password. If a password is required, and you do not specify the password or <i>prompt</i>, an error message is displayed: <pre>user@host> file copy ftp://username@ftp.hostname.net//filename file copy ftp.hostname.net: Not logged in. user@host> file copy ftp://username:prompt@ftphostname.net//filename</pre><p>Password for <i>username@ftp.hostname.net</i>:</p>• <i>http://hostname/path/filename</i>—File on a Hypertext Transfer Protocol (HTTP) server. You can also specify <i>hostname</i> as <i>username@hostname</i> or <i>username:password@hostname</i>. If a password is required and you omit it, you are prompted for it.• <i>re0:/path/filename</i> or <i>re1:/path/filename</i>—File on a local Routing Engine.

Required Privilege Level configure—To enter configuration mode.

Related Documentation • *Deactivating and Reactivating Statements and Identifiers in a Junos Configuration*

show chassis alarms

List of Syntax	Syntax on page 452 Syntax (TX Matrix Routers) on page 452 Syntax (TX Matrix Plus Routers) on page 452 Syntax (MX Series Routers) on page 452 Syntax (MX2010 3D Universal Edge Routers) on page 452 Syntax (MX2020 3D Universal Edge Routers) on page 452 Syntax (QFX Series) on page 452 Syntax (PTX Series Packet Transport Switches) on page 452 Syntax (ACX Series Universal Access Routers) on page 452
Syntax	show chassis alarms
Syntax (TX Matrix Routers)	show chassis alarms <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis alarms <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis alarms <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis alarms
Syntax (MX2020 3D Universal Edge Routers)	show chassis alarms
Syntax (QFX Series)	show chassis alarms <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (PTX Series Packet Transport Switches)	show chassis alarms
Syntax (ACX Series Universal Access Routers)	show chassis alarms
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option for the TX Matrix Plus router introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.

	<p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display information about the conditions that have been configured to trigger alarms.
Options	<p>none—Display information about the conditions that have been configured to trigger alarms.</p> <p>all-members—(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number. Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. <p>local—(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.</p> <p>member <i>member-id</i>—(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> variable with a value of 0 or 1.</p> <p>node-device <i>name</i>—(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.</p> <p>scc—(TX Matrix router only) (Optional) Show information about the TX Matrix router (switch-card chassis).</p> <p>sfc <i>number</i>—(TX Matrix Plus router only) (Optional) Show information about the respective TX Matrix Plus router, which is the switch-fabric chassis. Replace <i>number</i> variable with 0.</p>
Additional Information	You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (**Plane/FPC/PFE**) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby sonic clock generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code:
257
Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

To understand what CM_ALARM error codes mean, you need to first identify the structure of the CM Alarm codes. A CM_ALARM code has the following structure:

Bits:	Error type:
1-31	Major (1)
0	Minor (0)

As per the above table, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

```
Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

First, you have to convert 559 to binary; that is **1000101111**. The LSB in this case is 1, which means that this is a major alarm. After removing the LSB, you are left with **100010111**, which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3

CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZERR	4
CMALARM_LCHIP_LOUT_HDRF_TOERRERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITYERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERRERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROPERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERRERR	9
CMALARM_LCHIP_UCODE_TIMEOUTERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROPERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZEERR	12
CMALARM_LCHIP_SRAM_PARITYERR	13
CMALARM_LCHIP_UCODE_OVFLWERR	14
CMALARM_LCHIP_LOUT_HDRF_MTUERR	15

Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECTERR	128

Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUTERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLWERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLWERR	258
CMALARM_NCHIP_RDDMA_SIZEERR	259
CMALARM_NCHIP_RDDMA_JBUS_CRCERR	260
CMALARM_NCHIP_WRDMA_PKTRERR	261
CMALARM_NCHIP_WRDMA_PKT_CRCERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUTERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLWERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLWERR	265

CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266
CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274
CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278
CMALARM_NCHIP_DBUF_CRC_ERR	279

Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512

Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602
CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608

CMALARM_ICHIP_WO_IP_INTER_ERR	609
CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626
CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM_NCHIP_DBUF_CRC_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice-versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257

Actual Error Code = $(257-1)/2 = 128$. Similarly, for minor alarms, Actual Error Code = $(\text{Error Code})/2$

Required Privilege Level

view

Related Documentation

- *Configuring an Alarm Entry and Its Attributes*
- *Chassis Conditions That Trigger Alarms*

List of Sample Output

[show chassis alarms \(Alarms Active\) on page 458](#)
[show chassis alarms \(No Alarms Active\) on page 458](#)
[show chassis alarms \(Fan Tray\) on page 458](#)
[show chassis alarms \(MX2020 Router\) on page 459](#)
[show chassis alarms \(MX2010 Router\) on page 459](#)
[show chassis alarms \(T4000 Router\) on page 459](#)
[show chassis alarms \(Unreachable Destinations Present on a T Series Router\) on page 459](#)

[show chassis alarms \(FPC Offline Due to Unreachable Destinations on a T Series Router\) on page 459](#)
[show chassis alarms \(SCG Absent on a T Series Router\) on page 460](#)
[show chassis alarms \(Alarms Active on a TX Matrix Router\) on page 460](#)
[show chassis alarms \(TX Matrix Plus router with 3D SIBs\) on page 460](#)
[show chassis alarms \(Alarms on a T4000 Router After the enhanced-mode Statement is Enabled\) on page 461](#)
[show chassis alarms \(Backup Routing Engine\) on page 461](#)
[show chassis alarms \(Alarms Active on the QFX Series\) on page 461](#)
[show chassis alarms node-device \(Alarms Active on the QFabric System\) on page 461](#)
[show chassis alarms \(Alarms Active on the QFabric System\) on page 462](#)
[show chassis alarms \(Alarms Active on an EX8200 Switch\) on page 462](#)
[show chassis alarms \(Alarms Active on a PTX5000 Packet Transport Switch\) on page 462](#)
[show chassis alarms \(Alarms Active on an ACX2000 Universal Access Router\) on page 463](#)

Output Fields Table 50 on page 458 lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 50: show chassis alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major .
Description	Information about the alarm.

Sample Output

show chassis alarms (Alarms Active)

```

user@host> show chassis alarms
3 alarms are currently active
Alarm time          Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed

```

show chassis alarms (No Alarms Active)

```

user@host> show chassis alarms
No alarms are currently active

```

show chassis alarms (Fan Tray)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time          Class  Description
2010-11-11 20:27:38 UTC Major Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC Minor Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC Major Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC Major Side Fan Tray 0 Failure

```

show chassis alarms (MX2020 Router)

```

user@host> show chassis alarms
1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online

```

show chassis alarms (MX2010 Router)

```

user@host> show chassis alarms
7 alarms currently active
Alarm time      Class Description
2012-08-07 00:46:06 PDT Major Fan Tray 2 Failure
2012-08-06 18:24:36 PDT Minor Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT Minor Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT Minor Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT Minor Loss of communication with Backup RE
2012-08-03 12:26:03 PDT Minor Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT Minor Redundant feed missing for PSM 7

```

show chassis alarms (T4000 Router)

```

user@host> show chassis alarms
9 alarms currently active
Alarm time      Class Description
2007-06-02 01:41:10 UTC Minor RE 0 Not Supported
2007-06-02 01:41:10 UTC Minor CB 0 Not Supported
2007-06-02 01:41:10 UTC Minor Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC Major SPMB 1 not online
2007-05-30 19:37:29 UTC Minor Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC Major PEM 1 Input Failure
2007-05-30 19:37:13 UTC Major PEM 0 Not OK
2007-05-30 19:37:03 UTC Major PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC Minor Backup RE Active

```

show chassis alarms (Unreachable Destinations Present on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time      Class Description
2011-08-30 18:43:53 PDT Major FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time      Class Description
2011-08-30 18:43:53 PDT Major FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online

```

```

2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (SCG Absent on a T Series Router)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time          Class Description
2011-01-23 21:42:46 PST Major SCG 0 NO EXT CLK MEAS-BKUP SCG ABS

```

show chassis alarms (Alarms Active on a TX Matrix Router)

```

user@host> show chassis alarms
scc-re0:
-----
8 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:53 PDT Minor LCC 0 Minor Errors
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:52 PDT Major SIB 2 Absent
2004-08-05 18:43:52 PDT Major SIB 1 Absent
2004-08-05 18:43:52 PDT Major SIB 0 Absent
2004-08-05 18:43:33 PDT Major LCC 2 Major Errors
2004-08-05 18:43:28 PDT Major LCC 0 Major Errors
2004-08-05 18:43:05 PDT Minor LCC 2 Minor Errors
lcc0-re0:
-----
5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:49 PDT Major SIB 2 Absent
2004-08-05 18:43:49 PDT Major SIB 1 Absent
2004-08-05 18:43:49 PDT Major SIB 0 Absent
2004-08-05 18:43:28 PDT Major PEM 0 Not OK
lcc2-re0:
-----
5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:35 PDT Minor SIB 3 Not Online
2004-08-05 18:43:33 PDT Major SIB 2 Absent
2004-08-05 18:43:33 PDT Major SIB 1 Absent
2004-08-05 18:43:33 PDT Major SIB 0 Absent
2004-08-05 18:43:05 PDT Minor PEM 1 Absent

```

show chassis alarms (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis alarms
sfc0-re0:
-----
Alarm time          Class Description
2012-07-19 10:07:32 UTC Minor SIB F13 0 Temperature Warm
2012-07-19 10:07:07 UTC Minor SIB F2S 0/6 Temperature Warm
2012-07-19 10:07:07 UTC Minor SIB F2S 0/4 Temperature Warm
2012-07-19 10:07:07 UTC Minor SIB F2S 0/2 Temperature Warm
2012-07-19 10:07:07 UTC Minor SIB F2S 0/0 Temperature Warm
2012-07-19 10:07:07 UTC Minor SIB F13 6 Temperature Warm
2012-07-19 10:06:42 UTC Minor SIB F2S 2/6 Temperature Warm
2012-07-19 10:06:42 UTC Minor SIB F2S 2/4 Temperature Warm
2012-07-19 10:06:42 UTC Minor SIB F2S 2/2 Temperature Warm

```



```

2012-07-19 10:06:42 UTC Minor SIB F2S 2/0 Temperature Warm
2012-07-19 10:06:42 UTC Minor SIB F13 3 Temperature Warm
2012-07-19 10:06:17 UTC Minor Temperature Warm
2012-07-19 10:06:17 UTC Minor SIB F2S 1/6 Temperature Warm
2012-07-19 10:06:17 UTC Minor SIB F2S 1/4 Temperature Warm
2012-07-19 10:06:17 UTC Minor SIB F2S 1/2 Temperature Warm
2012-07-19 10:06:17 UTC Minor SIB F2S 1/0 Temperature Warm
lcc0-re0:

```

```

-----
Alarm time          Class Description
2012-07-19 10:04:13 UTC Minor Temperature Warm
2012-07-19 10:04:13 UTC Minor SIB 2 Temperature Warm
2012-07-19 10:04:13 UTC Minor SIB 1 Temperature Warm
2012-07-19 10:04:13 UTC Minor SIB 0 Temperature Warm

```

```
lcc2-re0:
```

```

-----
Alarm time          Class Description
2012-07-19 10:04:18 UTC Minor Temperature Warm
2012-07-19 10:04:18 UTC Minor SIB 2 Temperature Warm
2012-07-19 10:04:18 UTC Minor SIB 1 Temperature Warm
2012-07-19 10:04:18 UTC Minor SIB 0 Temperature Warm

```

show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

On T4000 routers, when you include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router are online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation of the alarms.

```

user@host> show chassis alarms
2 alarms currently active
Alarm time          Class Description
2011-10-22 10:10:47 PDT Major FPC 1 misconfig
2011-10-22 10:10:46 PDT Major FPC 0 misconfig

```

show chassis alarms (Backup Routing Engine)

```

user@host> show chassis alarms
2 alarms are currently active
Alarm time          Class Description
2005-04-07 10:12:22 PDT Minor Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT Major Host 1 compact-flash missing in Boot List

```

show chassis alarms (Alarms Active on the QFX Series)

```

user@switch> show chassis alarms
1 alarms currently active
Alarm time          Class Description
2012-03-05 2:10:24 UTC Major FPC 0 PEM 0 Airflow not matching Chassis Airflow

```

show chassis alarms node-device (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms node-device ED3691
node-device ED3694
3 alarms currently active
Alarm time          Class Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down

```

```

2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms
IC-A0001:
-----
1 alarms currently active
Alarm time          Class Description
2011-08-24 16:04:15 UTC Minor Backup RE Active

ED3694:
-----
3 alarms currently active
Alarm time          Class Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

SNG-0:
-----

NW-NG-0:
-----
1 alarms currently active
Alarm time          Class Description
2011-08-24 15:49:27 UTC Major ED3691 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on an EX8200 Switch)

```

user@switch> show chassis alarms

6 alarms currently active
Alarm time          Class Description
2010-12-02 19:15:22 UTC Major Fan Tray Failure
2010-12-02 19:15:22 UTC Major Fan Tray Failure
2010-12-02 19:15:14 UTC Minor Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC Minor Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC Major PSU 1 Output Failure
2010-12-02 19:14:18 UTC Minor Loss of communication with Backup RE

```

show chassis alarms (Alarms Active on a PTX5000 Packet Transport Switch)

```

user@switch> show chassis alarms

23 alarms currently active
Alarm time          Class Description
2011-07-12 16:22:05 PDT Minor No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT Major PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT Minor No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT Major PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT Major PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT Minor No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT Major PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT Major PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT Minor Backup RE Active

```

show chassis alarms (Alarms Active on an ACX2000 Universal Access Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time          Class  Description
2012-05-22 11:19:09 UTC Major  xe-0/3/1: Link down
2012-05-22 11:19:09 UTC Major  xe-0/3/0: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/7: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/6: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/3: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/2: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/1: Link down
```

show chassis beacon

show chassis beacon

(QFX Series)

```
<cb slot-number>
<fpc slot-number>
<interconnect-device name (cb slot-number | fpc slot-number)>
<node-device name>
```

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the beacon LED status on a QFX3500 standalone switch, Node device, and an Interconnect device. You can also display the beacon LED status of the Control Boards and Flexible PIC Concentrators on the Interconnect device.

Options

cb slot-number— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Control Board on the Interconnect device.

fpc slot-number— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator (FPC) on the Interconnect device. (QFX3500 switches only) (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator on the standalone switch.

interconnect-device name— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Interconnect device.

node-device name— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Node device.

Required Privilege Level view

Related Documentation

- [request chassis beacon on page 358](#)

List of Sample Output

[show chassis beacon \(QFX Series\) on page 465](#)
[show chassis beacon interconnect-device \(QFabric System\) on page 465](#)
[show chassis beacon interconnect-device fpc \(QFabric System\) on page 465](#)
[show chassis beacon node-device \(QFabric System\) on page 465](#)
[show chassis beacon node-device fpc \(QFabric System\) on page 465](#)

Output Fields [Table 51 on page 464](#) lists the output fields for the **show chassis beacon** command. Output fields are listed in the approximate order in which they appear.

Table 51: show chassis led Output Fields

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.

Table 51: show chassis led Output Fields (*continued*)

Field Name	Field Description
Beacon State	Status of the beacon state: <ul style="list-style-type: none"> Off—The beacon is OFF. On—The beacon is ON.

Sample Output

show chassis beacon (QFX Series)

```
user@switch> show chassis beacon
Slot          Beacon State
FPC           0          OFF
```

show chassis beacon interconnect-device (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1
Chassis              OFF
CB 0                  OFF
CB 1                  OFF
FC 0 FPC 0           OFF
FC 1 FPC 1           OFF
RC 0 FPC 8           OFF
RC 1 FPC 9           OFF
```

show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                ON
```

show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device node1
node1                ON
```

show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                ON
```

show chassis ethernet-switch interconnect-device fpc

Syntax	show chassis ethernet-switch interconnect-device <i>name</i> fpc <detail> <port <i>number</i> > <slot <i>number</i> >
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFX3000-G QFabric systems only) Display Ethernet switch information for the front card Flexible Port Concentrators (FPCs) in an Interconnect device.
Options	<p>none—Display Ethernet switch information about each connected port on each online FPC in the Interconnect device.</p> <p>detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot in the Interconnect device.</p> <p>port <i>number</i>—(Optional) Display Ethernet switch information about a specific port on an FPC in the Interconnect device.</p> <p>slot <i>number</i>—(Optional) Display Ethernet switch information about an FPC in a specific slot in the Interconnect device.</p>
Required Privilege Level	view
List of Sample Output	show chassis ethernet-switch interconnect-device fpc on page 469 show chassis ethernet-switch interconnect-device fpc detail on page 472 show chassis ethernet-switch fpc detail slot on page 479 show chassis ethernet-switch fpc interconnect-device port on page 485 show chassis ethernet-switch fpc interconnect-device detail port on page 487
Output Fields	Table 52 on page 466 lists the output fields for the show chassis ethernet-switch interconnect-device fpc command. Output fields are listed in the approximate order in which they appear.

Table 52: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	<p>Information about the link between each port on the FPC's Ethernet switch and one of the following devices:</p> <ul style="list-style-type: none"> FWD-SWITCH-0 FWD-SWITCH-1 CB0 CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .

Table 52: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
Duplex is	Duplex type of the Ethernet link: full or half .
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted frames of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.

Table 52: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.

Table 52: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device fpc

```

user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full

```

Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 124638
RX Octets 86496

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 82191
RX Octets 58979

Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 145475
RX Octets 206828

Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 1
RX Octets 0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 82290
RX Octets 59443

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 40900
RX Octets 30013

Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets 89456
RX Octets 123189

```
Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

root@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          124697
RX Octets          86535

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82229
RX Octets          59009

Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          145544
RX Octets          206925

Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82327
RX Octets          59472

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
```

```
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          40918
RX Octets          30028
```

Link is good on XE port 28 connected to device: CB0

```
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          89500
RX Octets          123244
```

Link is good on XE port 29 connected to device: CB1

```
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0
```

show chassis ethernet-switch interconnect-device fpc detail

```
user@host> show chassis ethernet-switch interconnect-device IC-WS001 fpc detail
```

Port statistics for FC0 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

```
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 121716
TX Packets 256-511 Octets 2200
TX Packets 512-1023 Octets 823
TX Packets 1024-1518 Octets 2
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                124742
TX Multicast Packets      0
TX Broadcast Packets      1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           27391588
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 85924
RX Packets 256-511 Octets 555
RX Packets 512-1023 Octets 86
RX Packets 1024-1518 Octets 1
```

```

RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 86566
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 20380581
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets 0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 80374
TX Packets 256-511 Octets 1347
TX Packets 512-1023 Octets 532
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 82257
TX Multicast Packets 0
TX Broadcast Packets 1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 18146746
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 58410
RX Packets 256-511 Octets 522
RX Packets 512-1023 Octets 96
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 59030
RX Multicast Packets 0
RX Broadcast Packets 0

```

```
RX FCS Errors          0
RX Align Errors        0
RX Fragments           0
RX Symbol errors       0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets   0
RX Oversize Packets    0
RX Jabbers             0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter  0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter        13882179
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets   0
TX Packets 65-127 Octets 0
TX Packets 128-255 Octets 144334
TX Packets 256-511 Octets 1077
TX Packets 512-1023 Octets 182
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets              145596
TX Multicast Packets   0
TX Broadcast Packets   0
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets    0
TX FCS Error Counter   0
TX Fragment Counter    0
TX Byte Counter        34262760
RX Packets 64 Octets   0
RX Packets 65-127 Octets 1
RX Packets 128-255 Octets 202090
RX Packets 256-511 Octets 3547
RX Packets 512-1023 Octets 1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets              206998
RX Multicast Packets   0
RX Broadcast Packets   1
RX FCS Errors          0
RX Fragments           0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets   0
RX Oversize Packets    0
RX Jabbers             0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter        45538262
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets   0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 0
```

```

TX Packets 256-511 Octets    0
TX Packets 512-1023 Octets   0
TX Packets 1024-1518 Octets  0
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets 0
TX Octets                    1
TX Multicast Packets         0
TX Broadcast Packets         1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              72
RX Packets 64 Octets         0
RX Packets 65-127 Octets     0
RX Packets 128-255 Octets    0
RX Packets 256-511 Octets    0
RX Packets 512-1023 Octets   0
RX Packets 1024-1518 Octets  0
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets 0
RX Octets                    0
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Fragments                 0
RX MAC Control Packets       0
RX Out of Range Length       0
RX Undersize Packets         0
RX Oversize Packets         0
RX Jabbers                   0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              0

```

Port statistics for FC1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

```

TX Packets 64 Octets         0
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    80560
TX Packets 256-511 Octets    1279
TX Packets 512-1023 Octets   514
TX Packets 1024-1518 Octets  3
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms 0
TX Octets                    82357
TX Multicast Packets         0
TX Broadcast Packets         1
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0

```

TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18059906
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58733
RX Packets 256-511 Octets	639
RX Packets 512-1023 Octets	119
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	59494
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	13994432

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	39971
TX Packets 256-511 Octets	668
TX Packets 512-1023 Octets	290
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	40933
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0


```

TX Byte Counter          9050841
RX Packets 64 Octets     0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 29767
RX Packets 256-511 Octets 225
RX Packets 512-1023 Octets 44
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                30039
RX Multicast Packets     0
RX Broadcast Packets     0
RX FCS Errors            0
RX Align Errors          0
RX Fragments             0
RX Symbol errors         0
RX Unsupported opcodes   0
RX Out of Range Length   0
RX False Carrier Errors  0
RX Undersize Packets     0
RX Oversize Packets      0
RX Jabbers               0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter    0
RX Control Frame Counter 0
RX Pause Frame Counter   0
RX Byte Counter          7043738
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets     0
TX Packets 65-127 Octets 0
TX Packets 128-255 Octets 88500
TX Packets 256-511 Octets 864
TX Packets 512-1023 Octets 163
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                89533
TX Multicast Packets     0
TX Broadcast Packets     0
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets      0
TX FCS Error Counter     0
TX Fragment Counter      0
TX Byte Counter          21038170
RX Packets 64 Octets     0
RX Packets 65-127 Octets 1
RX Packets 128-255 Octets 120531
RX Packets 256-511 Octets 1947
RX Packets 512-1023 Octets 804
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                123289
RX Multicast Packets     0
RX Broadcast Packets     1
RX FCS Errors            0

```

```
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             27110675
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             72
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   0
RX Packets 256-511 Octets   0
RX Packets 512-1023 Octets  0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   0
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             0
```

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7

show chassis ethernet-switch fpc detail slot

```
user@qfabric> show chassis ethernet-switch fpc detail 0
re0:
```

```
-----
Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 121823
TX Packets 256-511 Octets 2200
TX Packets 512-1023 Octets 823
TX Packets 1024-1518 Octets 2
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                  124849
TX Multicast Packets      0
TX Broadcast Packets      1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           27414524
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 85998
RX Packets 256-511 Octets 557
RX Packets 512-1023 Octets 86
RX Packets 1024-1518 Octets 1
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                  86642
RX Multicast Packets      0
RX Broadcast Packets      0
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0
RX Unsupported opcodes    0
RX Out of Range Length    0
```

```
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              20398564
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets         0
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    80443
TX Packets 256-511 Octets    1347
TX Packets 512-1023 Octets   532
TX Packets 1024-1518 Octets  3
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms 0
TX Octets                    82326
TX Multicast Packets         0
TX Broadcast Packets         1
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames           0
TX Frame deferred Xmsns      0
TX Frame excessive deferl    0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              18161734
RX Packets 64 Octets         0
RX Packets 65-127 Octets     0
RX Packets 128-255 Octets    58460
RX Packets 256-511 Octets    523
RX Packets 512-1023 Octets   96
RX Packets 1024-1518 Octets  2
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    59081
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
```

```

RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             13894171
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    0
TX Packets 128-255 Octets   144458
TX Packets 256-511 Octets   1080
TX Packets 512-1023 Octets  182
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   145723
TX Multicast Packets        0
TX Broadcast Packets        0
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             34292735
RX Packets 64 Octets        0
RX Packets 65-127 Octets    1
RX Packets 128-255 Octets   202266
RX Packets 256-511 Octets   3547
RX Packets 512-1023 Octets  1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   207174
RX Multicast Packets        0
RX Broadcast Packets        1
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             45576186
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets    0
TX Packets 256-511 Octets    0
TX Packets 512-1023 Octets   0
TX Packets 1024-1518 Octets  0
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets 0
TX Octets                   1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0

```

TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

Port statistics for FC1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80629
TX Packets 256-511 Octets	1279
TX Packets 512-1023 Octets	514
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82426
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18074790
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58785
RX Packets 256-511 Octets	640
RX Packets 512-1023 Octets	119

```

RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 59547
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 14006842
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets 0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 40004
TX Packets 256-511 Octets 668
TX Packets 512-1023 Octets 290
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 40966
TX Multicast Packets 0
TX Broadcast Packets 1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 9058102
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 29794
RX Packets 256-511 Octets 225
RX Packets 512-1023 Octets 44
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 30066
RX Multicast Packets 0

```

```
RX Broadcast Packets      0
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0
RX Unsupported opcodes    0
RX Out of Range Length    0
RX False Carrier Errors   0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter     0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           7050000
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  0
TX Packets 128-255 Octets 88579
TX Packets 256-511 Octets 865
TX Packets 512-1023 Octets 163
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                 89613
TX Multicast Packets      0
TX Broadcast Packets      0
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           21056842
RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 120633
RX Packets 256-511 Octets 1947
RX Packets 512-1023 Octets 804
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                 123391
RX Multicast Packets      0
RX Broadcast Packets      1
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           27132820
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
```



```

TX Packets 128-255 Octets 0
TX Packets 256-511 Octets 0
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets 1
TX Multicast Packets 0
TX Broadcast Packets 1
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 72
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 0
RX Packets 256-511 Octets 0
RX Packets 512-1023 Octets 0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets 0
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 0

```

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7

show chassis ethernet-switch fpc interconnect-device port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 port 2
```

Summary for switch on FC0

Link is good on GE port 2 connected to device: FWD-SWITCH-0

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 319466

RX Octets 221869

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 210295

RX Octets 151164

Link is good on XE port 28 connected to device: CB0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 373033

RX Octets 529760

Link is good on XE port 29 connected to device: CB1

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 1

RX Octets 0

Summary for switch on FC1

Link is good on GE port 2 connected to device: FWD-SWITCH-0

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 210760

RX Octets 152617

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 104587

RX Octets 77315

Link is good on XE port 28 connected to device: CB0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

```

Flow Control RX is Disabled
TX Octets                229932
RX Octets                 315346

```

```

Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                 1
RX Octets                 0

```

show chassis ethernet-switch fpc interconnect-device detail port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 detail port 2
```

```
Port statistics for FC0 switch
```

```
Statistics for port 2 connected to device FWD-SWITCH-0:
```

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 311974
TX Packets 256-511 Octets 5552
TX Packets 512-1023 Octets 2084
TX Packets 1024-1518 Octets 2
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                 319613
TX Multicast Packets      0
TX Broadcast Packets      1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           70091196
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 220284
RX Packets 256-511 Octets 1486
RX Packets 512-1023 Octets 198
RX Packets 1024-1518 Octets 1
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                 221969
RX Multicast Packets      0
RX Broadcast Packets      0
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0

```

```
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             52192002
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   205595
TX Packets 256-511 Octets   3426
TX Packets 512-1023 Octets  1366
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                   210391
TX Multicast Packets        0
TX Broadcast Packets        1
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0
TX Frame deferred Xtns      0
TX Frame excessive deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             46380018
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   149866
RX Packets 256-511 Octets   1194
RX Packets 512-1023 Octets  173
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                   151235
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
```

```

RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            35496911
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets       0
TX Packets 65-127 Octets   0
TX Packets 128-255 Octets  370150
TX Packets 256-511 Octets  2680
TX Packets 512-1023 Octets 371
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  373204
TX Multicast Packets       0
TX Broadcast Packets       0
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets        0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            87688913
RX Packets 64 Octets       0
RX Packets 65-127 Octets   1
RX Packets 128-255 Octets  517569
RX Packets 256-511 Octets  8978
RX Packets 512-1023 Octets 3450
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  530003
RX Multicast Packets       0
RX Broadcast Packets       1
RX FCS Errors              0
RX Fragments               0
RX MAC Control Packets     0
RX Out of Range Length     0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            116471142
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets       0
TX Packets 65-127 Octets   1
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  1
TX Multicast Packets       0
TX Broadcast Packets       1

```

TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

show chassis ethernet-switch interconnect-device cb

Syntax	show chassis ethernet-switch interconnect-device <i>name</i> cb <detail> <port <i>number</i> > <slot <i>number</i> >
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFX3000-G QFabric systems only) Display Ethernet switch information for the Control Board (CB) ports in an Interconnect device.
Options	<p>none—Display Ethernet switch information about each connected port on each online CB in the Interconnect device.</p> <p>detail—(Optional) Display detailed status information for all CBs or for the CB in the specified slot in the Interconnect device.</p> <p>port <i>number</i>—(Optional) Display Ethernet switch information about a specific port on a CB in the Interconnect device.</p> <p>slot <i>number</i>—(Optional) Display Ethernet switch information about a CB in a specific slot in the Interconnect device.</p>
Required Privilege Level	view
List of Sample Output	show chassis ethernet-switch interconnect-device cb on page 495 show chassis ethernet-switch interconnect-device cb detail on page 496 show chassis ethernet-switch interconnect-device cb detail slot port on page 501
Output Fields	Table 52 on page 466 lists the output fields for the show chassis ethernet-switch interconnect-device cb command. Output fields are listed in the approximate order in which they appear.

Table 53: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: <ul style="list-style-type: none"> • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .
Duplex is	Duplex type of the Ethernet link: full or half .

Table 53: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
Flow Control TX is Enabled (or Disabled)	Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
Flow Control RX is Enabled (or Disabled)	Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted packets of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.

Table 53: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.

Table 53: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device cb

```

user@switch> show chassis ethernet-switch interconnect-device IC-WS001 cb
Displaying summary for switch 0
Link is down on XE port 1 connected to device: FPC7
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 2 connected to device: FPC6
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 3 connected to device: FPC5
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 5 connected to device: FPC4
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 7 connected to device: FPC3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 9 connected to device: FPC2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 10 connected to device: FPC1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                326358
  RX Octets                 237947

Link is good on XE port 11 connected to device: FPC0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                548249
  RX Octets                 386013

Link is down on XE port 20 connected to device: SFP3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 21 connected to device: SFP2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 22 connected to device: SFP1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                1
  RX Octets                11704758

Link is good on XE port 23 connected to device: SFP0
  Speed is 1000Mb

```

```
Duplex is full
Autonegotiate is Enabled
TX Octets          1500022
RX Octets          11629453
```

```
Link is good on XE port 24 connected to device: VCCPD
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
TX Octets          23332467
RX Octets          1500023
```

```
Link is good on GE port 25 connected to device: SFI
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
TX Octets          643918
RX Octets          894548
```

show chassis ethernet-switch interconnect-device cb detail

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb detail
Port statistics for CB switch
```

```
Link is down on XE port 1 connected to device: FPC7
```

```
Link is down on XE port 2 connected to device: FPC6
```

```
Link is down on XE port 3 connected to device: FPC5
```

```
Link is down on XE port 5 connected to device: FPC4
```

```
Link is down on XE port 7 connected to device: FPC3
```

```
Link is down on XE port 9 connected to device: FPC2
```

```
Statistics for port 10 connected to device FPC1:
```

```
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 319293
TX Packets 256-511 Octets 5043
TX Packets 512-1023 Octets 2072
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets 326415
TX Multicast Packets 0
TX Broadcast Packets 1
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 71659246
TX Packet OK Counter 326415
TX Pause Packet Counter 0
TX Unicast Counter 326414
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 235428
RX Packets 256-511 Octets 2134
```

```

RX Packets 512-1023 Octets  420
RX Packets 1024-1518 Octets  6
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                    237988
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Fragments                 0
RX MAC Control Packets       0
RX Out of Range Length       0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              55821504
RX Unicast Frame Count       237988
RX Packet OK Count           237988
Statistics for port 11 connected to device FPC0:
TX Packets 64 Octets         0
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    535483
TX Packets 256-511 Octets    9289
TX Packets 512-1023 Octets   3564
TX Packets 1024-1518 Octets  5
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets  0
TX Octets                    548342
TX Multicast Packets         0
TX Broadcast Packets         1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets          0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              120498414
TX Packet OK Counter         548342
TX Pause Packet Counter      0
TX Unicast Counter           548341
RX Packets 64 Octets         0
RX Packets 65-127 Octets     0
RX Packets 128-255 Octets    382931
RX Packets 256-511 Octets    2762
RX Packets 512-1023 Octets   386
RX Packets 1024-1518 Octets  3
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                    386082
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Fragments                 0
RX MAC Control Packets       0
RX Out of Range Length       0
RX Undersize Packets         0

```

RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	90717369
RX Unicast Frame Count	386082
RX Packet OK Count	386082

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	0
TX Packets 128-255 Octets	0
TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	0
TX Packets 1024-1518 Octets	1
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	1
TX Multicast Packets	1
TX Broadcast Packets	0
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1422
RX Packet OK Count	1
RX Packets 64 Octets	230013
RX Packets 65-127 Octets	174529
RX Packets 128-255 Octets	286735
RX Packets 256-511 Octets	343412
RX Packets 512-1023 Octets	172152
RX Packets 1024-1518 Octets	10500065
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11706906
RX Multicast Packets	11672320
RX Broadcast Packets	34460
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0

```

RX Jabbers                                0
RX 1519-1522 Good Vlan Frms              0
RX MTU Exceed Counter                     0
RX Control Frame Counter                  0
RX Pause Frame Counter                     0
RX Byte Counter                           2379464164
RX Packet OK Count                         11706906
Statistics for port 23 connected to device SFP0:
TX Packets 64 Octets                       3
TX Packets 65-127 Octets                   484733
TX Packets 128-255 Octets                  219112
TX Packets 256-511 Octets                  129014
TX Packets 512-1023 Octets                 503
TX Packets 1024-1518 Octets                666958
TX Packets 1519-2047 Octets                0
TX Packets 2048-4095 Octets                0
TX Packets 4096-9216 Octets                0
TX 1519-1522 Good Vlan Frms                0
TX Octets                                 1500323
TX Multicast Packets                       794098
TX Broadcast Packets                       1040
TX Single Collision Frames                 0
TX Mult. Collision Frames                  0
TX Late Collisions                         0
TX Excessive Collisions                   0
TX Collision Frames                       0
TX PAUSEMAC Ctrl Frames                   0
TX MAC Ctrl Frames                        0
TX Frame Deferred Xms                      0
TX Frame Excessive Deferl                 0
TX Oversize Packets                       0
TX Jabbers                                0
TX FCS Error Counter                      0
TX Fragment Counter                       0
TX Byte Counter                           1065466891
RX Packet OK Count                         1500323
RX Packets 64 Octets                       341563
RX Packets 65-127 Octets                   430810
RX Packets 128-255 Octets                  318279
RX Packets 256-511 Octets                  347147
RX Packets 512-1023 Octets                 184798
RX Packets 1024-1518 Octets                10008993
RX Packets 1519-2047 Octets                0
RX Packets 2048-4095 Octets                0
RX Packets 4096-9216 Octets                0
RX Octets                                 11631590
RX Multicast Packets                       10878484
RX Broadcast Packets                       33420
RX FCS Errors                             0
RX Align Errors                           0
RX Fragments                             0
RX Symbol Errors                          0
RX Unsupported Opcodes                    0
RX Out of Range Length                    0
RX False Carrier Errors                    0
RX Undersize Packets                      0
RX Oversize Packets                       0
RX Jabbers                                0
RX 1519-1522 Good Vlan Frms              0
RX MTU Exceed Counter                     0
RX Control Frame Counter                  0

```

```
RX Pause Frame Counter      0
RX Byte Counter             1720484325
RX Packet OK Count          11631591
Statistics for port 24 connected to device VCCPD:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1176546
TX Packets 128-255 Octets   604988
TX Packets 256-511 Octets   690561
TX Packets 512-1023 Octets  356942
TX Packets 1024-1518 Octets 20507438
TX Packets 1519-2047 Octets 278
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 278
TX Octets                   23336753
TX Multicast Packets        22549383
TX Broadcast Packets        67862
TX Single Collision Frames  0
TX Mult. Collision Frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision Frames         0
TX PAUSEMAC Ctrl Frames    0
TX MAC Ctrl Frames         0
TX Frame Deferred Xmsns     0
TX Frame Excessive Deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             4191296788
RX Packet OK Count          23336753
RX Packets 64 Octets        3
RX Packets 65-127 Octets    484673
RX Packets 128-255 Octets   219074
RX Packets 256-511 Octets   129100
RX Packets 512-1023 Octets  516
RX Packets 1024-1518 Octets 666959
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                   1500325
RX Multicast Packets        794099
RX Broadcast Packets        1040
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol Errors            0
RX Unsupported Opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             1071469739
RX Packet OK Count          1500325
Statistics for port 25 connected to device SFI:
```



```

TX Packets 64 Octets      12
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 618363
TX Packets 256-511 Octets 4896
TX Packets 512-1023 Octets 806
TX Packets 1024-1518 Octets 19950
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 0
TX Octets                  644028
TX Multicast Packets       4
TX Broadcast Packets       19954
TX Single Collision Frames 0
TX Mult. Collision Frames  0
TX Late Collisions         0
TX Excessive Collisions    0
TX Collision Frames        0
TX PAUSEMAC Ctrl Frames    0
TX MAC Ctrl Frames         0
TX Frame Deferred Xmsns    0
TX Frame Excessive Deferl  0
TX Oversize Packets        0
TX Jabbers                 0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            167039705
RX Packet OK Count         644028
RX Packets 64 Octets       0
RX Packets 65-127 Octets   0
RX Packets 128-255 Octets  854776
RX Packets 256-511 Octets  14332
RX Packets 512-1023 Octets 5636
RX Packets 1024-1518 Octets 19954
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                  894698
RX Multicast Packets       0
RX Broadcast Packets       19943
RX FCS Errors              0
RX Align Errors            0
RX Fragments               0
RX Symbol Errors           0
RX Unsupported Opcodes     0
RX Out of Range Length     0
RX False Carrier Errors    0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            212658920
RX Packet OK Count         894698

```

show chassis ethernet-switch interconnect-device cb detail slot port

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb slot 1 port 1
```

re0:

Port statistics for CB switch

Link is down on XE port 1 connected to device: FPC7

Link is down on XE port 2 connected to device: FPC6

Link is down on XE port 3 connected to device: FPC5

Link is down on XE port 5 connected to device: FPC4

Link is down on XE port 7 connected to device: FPC3

Link is down on XE port 9 connected to device: FPC2

Statistics for port 10 connected to device FPC1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	319366
TX Packets 256-511 Octets	5043
TX Packets 512-1023 Octets	2072
TX Packets 1024-1518 Octets	6
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	326488
TX Multicast Packets	0
TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	71675330
TX Packet OK Counter	326488
TX Pause Packet Counter	0
TX Unicast Counter	326487
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	235481
RX Packets 256-511 Octets	2134
RX Packets 512-1023 Octets	420
RX Packets 1024-1518 Octets	6
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	238041
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	55834224
RX Unicast Frame Count	238041

```

RX Packet OK Count          238041
Statistics for port 11 connected to device FPC0:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   535606
TX Packets 256-511 Octets   9289
TX Packets 512-1023 Octets  3564
TX Packets 1024-1518 Octets 5
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   548465
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             120525524
TX Packet OK Counter        548465
TX Pause Packet Counter     0
TX Unicast Counter          548464
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   383018
RX Packets 256-511 Octets   2762
RX Packets 512-1023 Octets  386
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   386169
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             90738249
RX Unicast Frame Count      386169
RX Packet OK Count          386169

```

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

```

TX Packets 64 Octets        0
TX Packets 65-127 Octets    0
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 1
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0

```

TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	1
TX Multicast Packets	1
TX Broadcast Packets	0
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1422
RX Packet OK Count	1
RX Packets 64 Octets	230071
RX Packets 65-127 Octets	174571
RX Packets 128-255 Octets	286812
RX Packets 256-511 Octets	343500
RX Packets 512-1023 Octets	172203
RX Packets 1024-1518 Octets	10502544
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11709701
RX Multicast Packets	11675110
RX Broadcast Packets	34465
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2383079858
RX Packet OK Count	11709701

Statistics for port 23 connected to device SFP0:

TX Packets 64 Octets	3
TX Packets 65-127 Octets	485048
TX Packets 128-255 Octets	219200
TX Packets 256-511 Octets	129053
TX Packets 512-1023 Octets	503
TX Packets 1024-1518 Octets	667127
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	1500934
TX Multicast Packets	794300

TX Broadcast Packets	1040
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1065764997
RX Packet OK Count	1500934
RX Packets 64 Octets	341648
RX Packets 65-127 Octets	431183
RX Packets 128-255 Octets	318367
RX Packets 256-511 Octets	347225
RX Packets 512-1023 Octets	184849
RX Packets 1024-1518 Octets	10011311
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11634583
RX Multicast Packets	10881071
RX Broadcast Packets	33425
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1723893006
RX Packet OK Count	11634583

Statistics for port 24 connected to device VCCPD:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1177102
TX Packets 128-255 Octets	605153
TX Packets 256-511 Octets	690727
TX Packets 512-1023 Octets	357044
TX Packets 1024-1518 Octets	20512235
TX Packets 1519-2047 Octets	278
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	278
TX Octets	23342539
TX Multicast Packets	22554760
TX Broadcast Packets	67872
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0

TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	4198344167
RX Packet OK Count	23342539
RX Packets 64 Octets	3
RX Packets 65-127 Octets	484985
RX Packets 128-255 Octets	219164
RX Packets 256-511 Octets	129139
RX Packets 512-1023 Octets	516
RX Packets 1024-1518 Octets	667128
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	1500935
RX Multicast Packets	794301
RX Broadcast Packets	1040
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1071770147
RX Packet OK Count	1500935

Statistics for port 25 connected to device SFI:

TX Packets 64 Octets	12
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	618503
TX Packets 256-511 Octets	4896
TX Packets 512-1023 Octets	806
TX Packets 1024-1518 Octets	19950
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	644168
TX Multicast Packets	4
TX Broadcast Packets	19954
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0

TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	167073305
RX Packet OK Count	644168
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	854972
RX Packets 256-511 Octets	14332
RX Packets 512-1023 Octets	5636
RX Packets 1024-1518 Octets	19954
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	894894
RX Multicast Packets	0
RX Broadcast Packets	19943
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	212702114
RX Packet OK Count	894894

show chassis environment

List of Syntax	Syntax on page 508
	Syntax (T320, T640, T1600, and T4000 Routers) on page 508
	Syntax (TX Matrix Routers) on page 508
	Syntax (TX Matrix Plus Routers) on page 508
	Syntax (MX Series Routers) on page 508
	Syntax (MX2020 3D Universal Edge Routers) on page 508
	Syntax (MX2010 3D Universal Edge Routers) on page 509
	Syntax (EX Series Switch) on page 509
	Syntax (EX Series Switch) on page 509
	Syntax (QFX Series) on page 509
	Syntax (PTX Series Packet Transport Switches) on page 509
	Syntax (ACX Series Universal Access Routers) on page 509
Syntax	show chassis environment
Syntax (T320, T640, T1600, and T4000 Routers)	show chassis environment <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> > <scg <i>scg-slot-number</i> > <sib <i>sib-slot-number</i> >
Syntax (TX Matrix Routers)	show chassis environment <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis environment <cb <i>cb-slot-number</i> > <cip <i>cip-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <lcc <i>number</i> > <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> > <scg <i>scg-slot-number</i> > <sfc <i>number</i> > <sib <i>sib-slot-number</i> >
Syntax (MX Series Routers)	show chassis environment <all-members> <local> <member <i>member-id</i> >
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment <adc <i>adc-slot-number</i> > <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <monitored>

	<psm <i>psm-slot-number</i> > <routing-engine <i>re-slot-number</i> > <sfb <i>sfb-slot-number</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment <adc <i>adc-slot-number</i> > <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <monitored> <psm <i>psm-slot-number</i> > <routing-engine <i>re-slot-number</i> > <sfb <i>sfb-slot-number</i> >
Syntax (EX Series Switch)	show chassis environment <all-members> <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <local> <member <i>member-id</i> > <routing-engine <i>re-slot-number</i> >
Syntax (EX Series Switch)	show chassis environment <all-members> <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <local> <member <i>member-id</i> > <power-supply-unit <i>psu-slot-number</i> > <routing-engine <i>slot-number</i> >
Syntax (QFX Series)	show chassis environment <cb <i>slot-number</i> <interconnect-device <i>name</i> >> <fpc <i>slot-number</i> <interconnect-device <i>name</i> >> <interconnect-device <i>name</i> <slot-number> <node-device <i>name</i> > <pem <i>slot-number</i> (interconnect-device <i>name</i> <i>slot-number</i>) (node-device <i>name</i>)> <routing-engine <i>name</i> <interconnect-device <i>name</i> <i>slot-number</i> >>
Syntax (PTX Series Packet Transport Switches)	show chassis environment <cb <i>cb-slot-number</i> > <ccg <i>ccg-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <monitored> <pdu <i>pdu-slot-number</i> > <routing-engine <i>re-slot-number</i> > <sib <i>sib-slot-number</i> >
Syntax (ACX Series Universal Access Routers)	show chassis environment <cb <i>cb-slot-number</i> > <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> >

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>monitored option added in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>pem option introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.</p>
Description	<p>Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.</p> <p>In addition on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.</p>
Options	<p>none—Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached routers.</p> <p>all-members—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.</p> <p>adc <i>adc-slot-number</i>—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace <i>adc-slot-number</i> with a value from 0 through 19. For MX2010 routers, replace <i>adc-slot-number</i> with a value from 0 through 9.</p> <p>cb <i>cb-slot-number</i>—(ACX Series Universal Access Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, PTX Series Packet Transport Switches, QFX Series, and T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace <i>cb-slot</i> with 0 or 1. For the EX Series switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> for information on CB slot numbering.</p> <p>cip <i>cip-slot-number</i>—(TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Connection Interface Panel (CIP). Replace the <i>cip-slot-number</i> variable with a value of 0 or 1.</p> <p>cb <i>interconnect-device name</i>—(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.</p> <p>cgc <i>cgc-slot-number</i>—(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace <i>cb-slot</i> with a value of 0 or 1.</p>

fpc *fpc-slot*—(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, PTX Series Packet Transport Switches, QFX Series, QFX3500 switches, QFabric systems, T Series routers, and TX Matrix Plus routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 routers, replace **fpc-slot** with a value from **0** through **9**. For MX2020 routers, replace **fpc-slot** with a value from **0** through **19**. For information about FPC numbering, see [show chassis environment fpc](#). On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see *EX Series Switches Hardware and CLI Terminology Mapping* for information on FPC numbering. On a TX Matrix Plus router with 3D SIBs replace **fpc-slot** with a value from **0** through **63**.

fpm—(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, PTX Series, Packet Transport Switches, T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

interconnect-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

monitored—(MX2020 routers and PTX Series Packet Transport Switches only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers and EX Series switches) (Optional) Display chassis environmental information for the local Virtual Chassis member.

member *member-id*—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace *member-id* variable with a value of **0** or **1**. For EX Series switches, see *member* for member ID values.

node-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Node device.

- pdu pdu-slot-number**—(PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.
- pem**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.
- pem pem-slot-number**—(ACX Series Universal Access Routers, M120, M320, and M40e routers, MX Series routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see [show chassis environment pem](#).
- psm psm-slot-number**—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace **psm-slot-number** with a value from 0 through 17. For MX2010 routers, replace **psm-slot-number** with a value from 0 through 8.
- psu psu-slot-number**—(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply. See *EX Series Switches Hardware and CLI Terminology Mapping* for detailed information.
- routing-engine**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device.
- routing-engine re-slot-number**—(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see [show chassis environment routing-engine](#).
- scg**—(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.
- scc**—(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (switch-card chassis).
- sfb sfb-slot-number**—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. Replace **sfb-slot-number** with a value from 0 through 7.
- sfc number**—(TX Matrix Plus routers only) (Optional) Display chassis environmental information about the respective TX Matrix Plus router (switch-fabric chassis). Replace **number** variable with 0.
- sib sib-slot-number**—(M320 routers, PTX Series Packet Transport Switches, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see *show chassis environment sib*.

Required Privilege view
Level

- Related Documentation**
- *show chassis environment adc*
 - [show chassis environment cb on page 561](#)
 - *show chassis environment ccg*
 - *show chassis environment cip*
 - [show chassis environment fpc on page 579](#)
 - *show chassis environment fpm*
 - *show chassis environment lcc*
 - *show chassis environment mcs*
 - *show chassis environment monitored*
 - *show chassis environment pcg*
 - *show chassis environment pdu*
 - [show chassis environment pem on page 604](#)
 - *show chassis environment psm*
 - *show chassis environment psu*
 - [show chassis environment routing-engine on page 613](#)
 - *show chassis environment scg*
 - *show chassis environment sfb*
 - *show chassis environment sib*
 - *show chassis environment sfc*

- List of Sample Output**
- [show chassis environment \(J2300 Router\) on page 515](#)
 - [show chassis environment \(J4300 or J6300 Router\) on page 515](#)
 - [show chassis environment \(M5 Router\) on page 515](#)
 - [show chassis environment \(M7i Router\) on page 516](#)
 - [show chassis environment \(M10 Router\) on page 516](#)
 - [show chassis environment \(M10i Router\) on page 516](#)
 - [show chassis environment \(M20 Router\) on page 517](#)
 - [show chassis environment \(M40 Router\) on page 517](#)
 - [show chassis environment \(M40e Router\) on page 517](#)
 - [show chassis environment \(M120 Router\) on page 518](#)
 - [show chassis environment \(M160 Router\) on page 519](#)
 - [show chassis environment \(M320 Router\) on page 520](#)
 - [show chassis environment \(MX240 Router\) on page 520](#)
 - [show chassis environment \(MX240 Router with Enhanced MX SCB\) on page 521](#)
 - [show chassis environment \(MX480 Router\) on page 522](#)
 - [show chassis environment \(MX480 Router with Enhanced MX SCB\) on page 523](#)
 - [show chassis environment \(MX960 Router\) on page 524](#)
 - [show chassis environment \(MX960 Router with Enhanced MX SCB\) on page 525](#)
 - [show chassis environment \(MX2020 Router\) on page 527](#)
 - [show chassis environment \(MX2010 Router\) on page 537](#)

[show chassis environment \(T320 Router\) on page 542](#)
[show chassis environment \(T640 Router\) on page 543](#)
[show chassis environment \(T4000 Router\) on page 544](#)
[show chassis environment \(TX Matrix Router\) on page 545](#)
[show chassis environment \(T1600 Router\) on page 547](#)
[show chassis environment \(TX Matrix Plus Router\) on page 548](#)
[show chassis environment \(TX Matrix Plus router with 3D SIBs\) on page 550](#)
[show chassis environment \(EX4200 Standalone Switch\) on page 553](#)
[show chassis environment \(EX8216 Switch\) on page 553](#)
[show chassis environment \(QFX Series\) on page 554](#)
[show chassis environment interconnect-device \(QFabric System\) on page 554](#)
[show chassis environment node-device \(QFabric System\) on page 556](#)
[show chassis environment pem node-device \(QFabric System\) on page 557](#)
[show chassis environment \(PTX5000 Packet Transport Switch\) on page 557](#)
[show chassis environment \(ACX2000 Universal Access Router\) on page 559](#)
[show chassis environment \(ACX4000 Universal Access Router\) on page 560](#)

Output Fields Table 54 on page 514 lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 54: show chassis environment Output Fields

Field Name	Field Description
Class	<p>Information about the category or class of chassis component:</p> <ul style="list-style-type: none"> • Power: Power information: <ul style="list-style-type: none"> • (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: OK, Testing, (during initial power-on), Failed, or Absent. • (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: OK, Testing, (during initial power-on), Check, Failed, or Absent. • (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: OK, Testing, (during initial power-on), Check, Failed, or Absent. • Temp: Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F). On PTX Series Packet Transport Switches and MX2010 and MX2020 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones. • Pic: On ACX4000 Routers, multiple temperature channels on a MIC. The status is: OK and the Measurement is in degrees Celsius (C) and Fahrenheit (F). • Fan: Fan status: OK, Testing (during initial power-on), Failed, or Absent. On PTX Series Packet Transport Switches and MX2010 and MX2020 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. Measurement indicates actual fan RPM (PTX and MX2010 and MX2020 Routers only). • Misc: Information about other components of the chassis. <ul style="list-style-type: none"> • On some routers, this field indicates the status of one or more additional components. • On the M40e, M160, and M320 router, Misc includes CIP (Connector Interface Panel). OK indicates that the CIP is present. Absent indicates that the CIP is not present. • On T Series routers, Misc includes CIP and SPMB (Switch Processor Mezzanine Board). OK indicates that the CIP or SPMB is present. Absent indicates that the CIP or SPMB is not present. • On PTX Series Packet Transport Switches, Misc includes the SPMB (Switch Processor Mezzanine Board). The SPMB is located on the control boards. OK indicates that the control board is present. Absent indicates that the control board is not present.

Table 54: show chassis environment Output Fields (*continued*)

Field Name	Field Description
Item	(MX2010 and MX2020 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).
Status	<p>(MX2010 and MX2020 Routers) Status of the specified chassis component. For example, if the Class is Fan, the fan status can be:</p> <ul style="list-style-type: none"> • OK: The fans are operational. • Testing: The fans are being tested during initial power-on. • Failed: The fans have failed or the fans are not spinning. • Absent: The fan tray is not installed. <p>If the Class is Power, the power supply status can be:</p> <ul style="list-style-type: none"> • OK: The power component is operational. • Testing: The power component is being tested during initial power-on. • Check: There is insufficient power---that is, fewer than the minimum required feeds are connected. • Failed: The inputs leads have failed. • Absent: The power component is not installed.
Measurement	(MX2010 and MX2020 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.

Sample Output

show chassis environment (J2300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp  Routing Engine    OK      40 degrees C / 104 degrees F
Fan   Fan              OK

```

show chassis environment (J4300 or J6300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp  Routing Engine    OK      41 degrees C / 105 degrees F
Fan   Fan 0            OK
      Fan 1            OK

```

show chassis environment (M5 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Power Power Supply A    OK
      Power Supply B    Absent
Temp  FPC 0             OK      30 degrees C / 86 degrees F
      FEB              OK      33 degrees C / 91 degrees F
      PS Intake         OK      27 degrees C / 80 degrees F
      PS Exhaust        OK      27 degrees C / 80 degrees F
      Routing Engine    OK      34 degrees C / 93 degrees F
Fans  Left Fan 1       OK      Spinning at normal speed

```

	Left Fan 2	OK	Spinning at normal speed
	Left Fan 3	OK	Spinning at normal speed
	Left Fan 4	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M7i Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply 0	OK	
	Power Supply 1	Absent	
Temp	Intake	OK	22 degrees C / 71 degrees F
	FPC 0	OK	23 degrees C / 73 degrees F
	Power Supplies	OK	23 degrees C / 73 degrees F
	CFEB Intake	OK	24 degrees C / 75 degrees F
	CFEB Exhaust	OK	29 degrees C / 84 degrees F
	Routing Engine	OK	26 degrees C / 78 degrees F
Fans	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed
	Fan 3	OK	Spinning at normal speed
	Fan 4	OK	Spinning at normal speed

show chassis environment (M10 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Failed	
Temp	FPC 0	OK	36 degrees C / 96 degrees F
	FPC 1	OK	35 degrees C / 95 degrees F
	FEB	OK	34 degrees C / 93 degrees F
	PS Intake	OK	31 degrees C / 87 degrees F
	PS Exhaust	OK	34 degrees C / 93 degrees F
	Routing Engine	OK	35 degrees C / 95 degrees F
Fans	Left Fan 1	OK	Spinning at normal speed
	Left Fan 2	OK	Spinning at normal speed
	Left Fan 3	OK	Spinning at normal speed
	Left Fan 4	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M10i Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply 0	OK	
	Power Supply 1	OK	
	Power Supply 2	Absent	
	Power Supply 3	Absent	
Temp	Intake	OK	26 degrees C / 78 degrees F
	FPC 0	OK	27 degrees C / 80 degrees F
	FPC 1	OK	28 degrees C / 82 degrees F
	Lower Power Supplies	OK	29 degrees C / 84 degrees F
	Upper Power Supplies	OK	28 degrees C / 82 degrees F
	CFEB Intake	OK	27 degrees C / 80 degrees F
	CFEB Exhaust	OK	36 degrees C / 96 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed

Fan Tray 0 Fan 3	OK	Spinning at normal speed
Fan Tray 0 Fan 4	OK	Spinning at normal speed
Fan Tray 0 Fan 5	OK	Spinning at normal speed
Fan Tray 0 Fan 6	OK	Spinning at normal speed
Fan Tray 0 Fan 7	OK	Spinning at normal speed
Fan Tray 0 Fan 8	OK	Spinning at normal speed
Fan Tray 1 Fan 1	Absent	
Fan Tray 1 Fan 2	Absent	
Fan Tray 1 Fan 3	Absent	
Fan Tray 1 Fan 4	Absent	
Fan Tray 1 Fan 5	Absent	
Fan Tray 1 Fan 6	Absent	
Fan Tray 1 Fan 7	Absent	
Fan Tray 1 Fan 8	Absent	

show chassis environment (M20 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply A      OK
      Power Supply B      Absent
Temp  FPC 0                OK          28 degrees C / 82 degrees F
      FPC 1                OK          27 degrees C / 80 degrees F
      Power Supply A      OK          22 degrees C / 71 degrees F
      Power Supply B      Absent
      SSB 0                OK          30 degrees C / 86 degrees F
      Backplane            OK          22 degrees C / 71 degrees F
      Routing Engine 0     OK          26 degrees C / 78 degrees F
      Routing Engine 1     Testing
Fans  Rear Fan             OK          Spinning at normal speed
      Front Upper Fan      OK          Spinning at normal speed
      Front Middle Fan     OK          Spinning at normal speed
      Front Bottom Fan     OK          Spinning at normal speed
Misc  Craft Interface      OK

```

show chassis environment (M40 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply A      OK
      Power Supply B      Absent
Temp  FPC 3                OK          24 degrees C / 75 degrees F
      FPC 6                OK          26 degrees C / 78 degrees F
      SCB                  OK          26 degrees C / 78 degrees F
      Backplane @ A1       OK          28 degrees C / 82 degrees F
      Backplane @ A2       OK          23 degrees C / 73 degrees F
      Routing Engine        OK          26 degrees C / 78 degrees F
Fans  Top Impeller         OK          Spinning at normal speed
      Bottom impeller      OK          Spinning at normal speed
      Rear Left Fan        OK          Spinning at normal speed
      Rear Center Fan      OK          Spinning at normal speed
      Rear Right Fan       OK          Spinning at normal speed
Misc  Craft Interface      OK

```

show chassis environment (M40e Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power PEM 0               OK

```

	PEM 1	Absent	
Temp	PCG 0	OK	44 degrees C / 111 degrees F
	PCG 1	OK	47 degrees C / 116 degrees F
	Routing Engine 0	OK	40 degrees C / 104 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	MCS 0	OK	45 degrees C / 113 degrees F
	MCS 1	OK	42 degrees C / 107 degrees F
	SFM 0 SPP	OK	40 degrees C / 104 degrees F
	SFM 0 SPR	OK	44 degrees C / 111 degrees F
	SFM 1 SPP	OK	43 degrees C / 109 degrees F
	SFM 1 SPR	OK	45 degrees C / 113 degrees F
	FPC 0	OK	38 degrees C / 100 degrees F
	FPC 1	OK	40 degrees C / 104 degrees F
	FPC 2	OK	38 degrees C / 100 degrees F
	FPC 4	OK	34 degrees C / 93 degrees F
	FPC 5	OK	43 degrees C / 109 degrees F
	FPC 6	OK	41 degrees C / 105 degrees F
	FPC 7	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	28 degrees C / 82 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
Misc	Fan Tray Front Right	OK	Spinning at normal speed
	CIP	OK	

show chassis environment (M120 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	
	PEM 1	OK	
	Routing Engine 0	OK	43 degrees C / 109 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 1 Intake	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 1 Exhaust B	OK	35 degrees C / 95 degrees F
	FEB 3 Intake	OK	35 degrees C / 95 degrees F
	FEB 3 Exhaust A	OK	37 degrees C / 98 degrees F
	FEB 3 Exhaust B	OK	39 degrees C / 102 degrees F
	FEB 4 Intake	OK	33 degrees C / 91 degrees F
	FEB 4 Exhaust A	OK	39 degrees C / 102 degrees F
	FEB 4 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 2 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust B	OK	31 degrees C / 87 degrees F
	FPC 3 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 3 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 4 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 4 Exhaust B	OK	30 degrees C / 86 degrees F
Fans	Front Top Tray Fan 1	OK	Spinning at normal speed
	Front Top Tray Fan 2	OK	Spinning at normal speed
	Front Top Tray Fan 3	OK	Spinning at normal speed
	Front Top Tray Fan 4	OK	Spinning at normal speed
	Front Top Tray Fan 5	OK	Spinning at normal speed

Front Top Tray Fan 6	OK	Spinning at normal speed
Front Top Tray Fan 7	OK	Spinning at normal speed
Front Top Tray Fan 8	OK	Spinning at normal speed
Front Bottom Tray Fan 1	OK	Spinning at normal speed
Front Bottom Tray Fan 2	OK	Spinning at normal speed
Front Bottom Tray Fan 3	OK	Spinning at normal speed
Front Bottom Tray Fan 4	OK	Spinning at normal speed
Front Bottom Tray Fan 5	OK	Spinning at normal speed
Front Bottom Tray Fan 6	OK	Spinning at normal speed
Front Bottom Tray Fan 7	OK	Spinning at normal speed
Front Bottom Tray Fan 8	OK	Spinning at normal speed
Rear Top Tray Fan 1	OK	Spinning at normal speed
Rear Top Tray Fan 2	OK	Spinning at normal speed
Rear Top Tray Fan 3	OK	Spinning at normal speed
Rear Top Tray Fan 4	OK	Spinning at normal speed
Rear Top Tray Fan 5	OK	Spinning at normal speed
Rear Top Tray Fan 6	OK	Spinning at normal speed
Rear Top Tray Fan 7	OK	Spinning at normal speed
Rear Top Tray Fan 8	OK	Spinning at normal speed
Rear Bottom Tray Fan 1	OK	Spinning at normal speed
Rear Bottom Tray Fan 2	OK	Spinning at normal speed
Rear Bottom Tray Fan 3	OK	Spinning at normal speed
Rear Bottom Tray Fan 4	OK	Spinning at normal speed
Rear Bottom Tray Fan 5	OK	Spinning at normal speed
Rear Bottom Tray Fan 6	OK	Spinning at normal speed
Rear Bottom Tray Fan 7	OK	Spinning at normal speed
Rear Bottom Tray Fan 8	OK	Spinning at normal speed

show chassis environment (M160 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power PEM 0                OK          PEM 1          Absent
Temp  PCG 0                OK          45 degrees C / 113 degrees F
      PCG 1                Absent
      Routing Engine 0      OK          35 degrees C / 95 degrees F
      Routing Engine 1      Absent
      MCS 0                OK          50 degrees C / 122 degrees F
      SFM 0 SPP             OK          47 degrees C / 116 degrees F
      SFM 0 SPR             OK          49 degrees C / 120 degrees F
      SFM 1 SPP             OK          50 degrees C / 122 degrees F
      SFM 1 SPR             OK          50 degrees C / 122 degrees F
      SFM 2 SPP             OK          51 degrees C / 123 degrees F
      SFM 2 SPR             OK          52 degrees C / 125 degrees F
      SFM 3 SPP             OK          52 degrees C / 125 degrees F
      SFM 3 SPR             OK          48 degrees C / 118 degrees F
      FPC 0                OK          45 degrees C / 113 degrees F
      FPC 6                OK          43 degrees C / 109 degrees F
      FPM CMB              OK          31 degrees C / 87 degrees F
      FPM Display          OK          33 degrees C / 91 degrees F
Fans  Rear Bottom Blower   OK          Spinning at normal speed
      Rear Top Blower      OK          Spinning at normal speed
      Front Top Blower     OK          Spinning at normal speed
      Fan Tray Rear Left   OK          Spinning at normal speed
      Fan Tray Rear Right  OK          Spinning at normal speed
      Fan Tray Front Left  OK          Spinning at normal speed
      Fan Tray Front Right OK          Spinning at normal speed
Misc  CIP                  OK

```

show chassis environment (M320 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Temp  PEM 0                 Absent
      PEM 1                 Absent
      PEM 2                 OK
      PEM 3                 OK
      Routing Engine 0      OK          33 degrees C / 91 degrees F
      Routing Engine 1      OK          32 degrees C / 89 degrees F
      CB 0                  OK          36 degrees C / 96 degrees F
      CB 1                  OK          36 degrees C / 96 degrees F
      SIB 0                 OK          38 degrees C / 100 degrees F
      SIB 1                 OK          29 degrees C / 84 degrees F
      SIB 2                 OK          38 degrees C / 100 degrees F
      SIB 3                 OK          41 degrees C / 105 degrees F
      FPC 0 Intake          OK          28 degrees C / 82 degrees F
      FPC 0 Exhaust         OK          40 degrees C / 104 degrees F
      FPC 1 Intake          OK          29 degrees C / 84 degrees F
      FPC 1 Exhaust         OK          39 degrees C / 102 degrees F
      FPC 2 Intake          OK          28 degrees C / 82 degrees F
      FPC 2 Exhaust         OK          38 degrees C / 100 degrees F
      FPC 3 Intake          OK          28 degrees C / 82 degrees F
      FPC 3 Exhaust         OK          39 degrees C / 102 degrees F
      FPC 6 Intake          OK          27 degrees C / 80 degrees F
      FPC 6 Exhaust         OK          39 degrees C / 102 degrees F
      FPC 7 Intake          OK          27 degrees C / 80 degrees F
      FPC 7 Exhaust         OK          42 degrees C / 107 degrees F
      FPM GBUS              OK          30 degrees C / 86 degrees F
Fan   Top Left Front fan    OK          Spinning at normal speed
      Top Right Rear fan    OK          Spinning at normal speed
      Top Right Front fan   OK          Spinning at normal speed
      Top Left Rear fan     OK          Spinning at normal speed
      Bottom Left Front fan  OK          Spinning at normal speed
      Bottom Right Rear fan  OK          Spinning at normal speed
      Bottom Right Front fan OK          Spinning at normal speed
      Bottom Left Rear fan   OK          Spinning at normal speed
      Rear Fan 1 (TOP)       OK          Spinning at normal speed
      Rear Fan 2             OK          Spinning at normal speed
      Rear Fan 3             OK          Spinning at normal speed
      Rear Fan 4             OK          Spinning at normal speed
      Rear Fan 5             OK          Spinning at normal speed
      Rear Fan 6             OK          Spinning at normal speed
      Rear Fan 7 (Bottom)    OK          Spinning at normal speed
Misc  CIP                   OK

```

show chassis environment (MX240 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Temp  PEM 0                 OK          40 degrees C / 104 degrees F
      PEM 1                 OK          45 degrees C / 113 degrees F
      PEM 2                 Absent
      PEM 3                 Absent
      Routing Engine 0      OK          39 degrees C / 102 degrees F
      Routing Engine 1      OK          37 degrees C / 98 degrees F
      CB 0 Intake           OK          36 degrees C / 96 degrees F
      CB 0 Exhaust A        OK          34 degrees C / 93 degrees F
      CB 0 Exhaust B        OK          38 degrees C / 100 degrees F
      CB 0 ACBC             OK          37 degrees C / 98 degrees F

```

CB 0 SF A	OK	49 degrees C / 120 degrees F
CB 0 SF B	OK	41 degrees C / 105 degrees F
CB 1 Intake	OK	37 degrees C / 98 degrees F
CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
CB 1 ACBC	OK	38 degrees C / 100 degrees F
CB 1 SF A	OK	47 degrees C / 116 degrees F
CB 1 SF B	OK	41 degrees C / 105 degrees F
FPC 1 Intake	OK	33 degrees C / 91 degrees F
FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 2 Intake	OK	32 degrees C / 89 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans Front Fan	OK	Spinning at normal speed
Middle Fan	OK	Spinning at normal speed
Rear Fan	OK	Spinning at normal speed

show chassis environment (MX240 Router with Enhanced MX SCB)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	40 degrees C / 104 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F
	CB 0 XF B	OK	41 degrees C / 105 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F
	CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust B	OK	39 degrees C / 102 degrees F

CB 1 ACBC	OK	38 degrees C / 100 degrees F
CB 1 XF A	OK	47 degrees C / 116 degrees F
CB 1 XF B	OK	41 degrees C / 105 degrees F
FPC 1 Intake	OK	33 degrees C / 91 degrees F
FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 2 Intake	OK	32 degrees C / 89 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans Front Fan	OK	Spinning at normal speed
Middle Fan	OK	Spinning at normal speed
Rear Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router)

user@host> show chassis environment		
Class	Item	Status Measurement
Temp	PEM 0	OK 35 degrees C / 95 degrees F
	PEM 1	OK 40 degrees C / 104 degrees F
	PEM 2	Absent
	PEM 3	Absent
	Routing Engine 0	OK 44 degrees C / 111 degrees F
	Routing Engine 1	OK 45 degrees C / 113 degrees F
	CB 0 Intake	OK 36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK 38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK 39 degrees C / 102 degrees F
	CB 0 ACBC	OK 37 degrees C / 98 degrees F
	CB 0 SF A	OK 51 degrees C / 123 degrees F
	CB 0 SF B	OK 44 degrees C / 111 degrees F
	CB 1 Intake	OK 36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK 39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK 40 degrees C / 104 degrees F
	CB 1 ACBC	OK 37 degrees C / 98 degrees F
	CB 1 SF A	OK 50 degrees C / 122 degrees F
	CB 1 SF B	OK 43 degrees C / 109 degrees F
	FPC 0 Intake	OK 36 degrees C / 96 degrees F
	FPC 0 Exhaust A	OK 39 degrees C / 102 degrees F

	FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
	FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	37 degrees C / 98 degrees F
	FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
	FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
	FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
	FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans	Top Rear Fan	OK	Spinning at normal speed
	Bottom Rear Fan	OK	Spinning at normal speed
	Top Middle Fan	OK	Spinning at normal speed
	Bottom Middle Fan	OK	Spinning at normal speed
	Top Front Fan	OK	Spinning at normal speed
	Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	51 degrees C / 123 degrees F
	CB 0 XF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F
	CB 1 XF A	OK	50 degrees C / 122 degrees F
	CB 1 XF B	OK	43 degrees C / 109 degrees F
	FPC 0 Intake	OK	36 degrees C / 96 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
	FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F

	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	37 degrees C / 98 degrees F
	FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
	FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
	FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
	FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans	Top Rear Fan	OK	Spinning at normal speed
	Bottom Rear Fan	OK	Spinning at normal speed
	Top Middle Fan	OK	Spinning at normal speed
	Bottom Middle Fan	OK	Spinning at normal speed
	Top Front Fan	OK	Spinning at normal speed
	Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX960 Router)

user@host> show chassis environment

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	Check	
	PEM 3	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	37 degrees C / 98 degrees F
	Routing Engine 1	Absent	
	CB 0 Intake	OK	24 degrees C / 75 degrees F
	CB 0 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 Intake	Absent	
	CB 1 Exhaust A	Absent	
	CB 1 Exhaust B	Absent	
	CB 1 ACBC	Absent	
	CB 1 SF A	Absent	
	CB 1 SF B	Absent	
	CB 2 Intake	Absent	
	CB 2 Exhaust A	Absent	
	CB 2 Exhaust B	Absent	
	CB 2 ACBC	Absent	
	CB 2 SF A	Absent	
	CB 2 SF B	Absent	
	FPC 4 Intake	OK	24 degrees C / 75 degrees F
	FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 4 Exhaust B	OK	38 degrees C / 100 degrees F

	FPC 7 Intake	OK	24 degrees C / 75 degrees F
	FPC 7 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust B	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	Failed	
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	Failed	
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX960 Router with Enhanced MX SCB)

user@host> show chassis environment

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	50 degrees C / 122 degrees F
	PEM 2	OK	50 degrees C / 122 degrees F
	PEM 3	OK	50 degrees C / 122 degrees F
	Routing Engine 0	OK	42 degrees C / 107 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	39 degrees C / 102 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	CB 0 ACBC	OK	38 degrees C / 100 degrees F
	CB 0 XF A	OK	53 degrees C / 127 degrees F
	CB 0 XF B	OK	47 degrees C / 116 degrees F
	CB 1 Intake	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust A	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust B	OK	41 degrees C / 105 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	52 degrees C / 125 degrees F
	CB 1 XF B	OK	47 degrees C / 116 degrees F
	CB 2 Intake	OK	32 degrees C / 89 degrees F
	CB 2 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 2 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 2 ACBC	OK	33 degrees C / 91 degrees F
	CB 2 XF A	OK	51 degrees C / 123 degrees F
	CB 2 XF B	OK	50 degrees C / 122 degrees F
	FPC 0 Intake	OK	35 degrees C / 95 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	50 degrees C / 122 degrees F
	FPC 0 I3 2 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	41 degrees C / 105 degrees F
	FPC 0 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F

FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 1 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 1 LU 0 TCAM TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 TCAM Chip	OK	57 degrees C / 134 degrees F
FPC 1 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 1 MQ 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TCAM TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 TCAM Chip	OK	52 degrees C / 125 degrees F
FPC 1 LU 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 Chip	OK	53 degrees C / 127 degrees F
FPC 1 MQ 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 2 Intake	OK	35 degrees C / 95 degrees F
FPC 2 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust B	OK	54 degrees C / 129 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	59 degrees C / 138 degrees F
FPC 2 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 2 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 2 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 2 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 IA 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 0 Chip	OK	46 degrees C / 114 degrees F
FPC 2 IA 1 TSensor	OK	45 degrees C / 113 degrees F
FPC 2 IA 1 Chip	OK	49 degrees C / 120 degrees F
FPC 3 Intake	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 3 I3 0 TSensor	OK	48 degrees C / 118 degrees F
FPC 3 I3 0 Chip	OK	52 degrees C / 125 degrees F
FPC 3 I3 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 3 I3 1 Chip	OK	48 degrees C / 118 degrees F
FPC 3 IA 0 TSensor	OK	41 degrees C / 105 degrees F
FPC 3 IA 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 Intake	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust A	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 Chip	OK	54 degrees C / 129 degrees F
FPC 5 LU 1 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 1 Chip	OK	61 degrees C / 141 degrees F
FPC 5 LU 2 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 5 LU 3 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 5 MQ 3 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 7 Intake	OK	36 degrees C / 96 degrees F

FPC 7 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
FPC 7 QX 0 TSensor	OK	42 degrees C / 107 degrees F
FPC 7 QX 0 Chip	OK	47 degrees C / 116 degrees F
FPC 7 LU 0 TCAM TSensor	OK	42 degrees C / 107 degrees F
FPC 7 LU 0 TCAM Chip	OK	44 degrees C / 111 degrees F
FPC 7 LU 0 TSensor	OK	42 degrees C / 107 degrees F
FPC 7 LU 0 Chip	OK	46 degrees C / 114 degrees F
FPC 7 MQ 0 TSensor	OK	42 degrees C / 107 degrees F
FPC 7 MQ 0 Chip	OK	45 degrees C / 113 degrees F
FPC 8 Intake	OK	33 degrees C / 91 degrees F
FPC 8 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 8 Exhaust B	OK	36 degrees C / 96 degrees F
FPC 8 I3 0 TSensor	OK	38 degrees C / 100 degrees F
FPC 8 I3 0 Chip	OK	43 degrees C / 109 degrees F
FPC 8 BDS 0 TSensor	OK	37 degrees C / 98 degrees F
FPC 8 BDS 0 Chip	OK	36 degrees C / 96 degrees F
FPC 8 IA 0 TSensor	OK	37 degrees C / 98 degrees F
FPC 8 IA 0 Chip	OK	37 degrees C / 98 degrees F
FPC 10 Intake	OK	38 degrees C / 100 degrees F
FPC 10 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 10 Exhaust B	OK	41 degrees C / 105 degrees F
FPC 10 I3 0 TSensor	OK	40 degrees C / 104 degrees F
FPC 10 I3 0 Chip	OK	42 degrees C / 107 degrees F
FPC 10 I3 1 TSensor	OK	40 degrees C / 104 degrees F
FPC 10 I3 1 Chip	OK	44 degrees C / 111 degrees F
FPC 10 I3 2 TSensor	OK	42 degrees C / 107 degrees F
FPC 10 I3 2 Chip	OK	43 degrees C / 109 degrees F
FPC 10 I3 3 TSensor	OK	39 degrees C / 102 degrees F
FPC 10 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 10 IA 0 TSensor	OK	36 degrees C / 96 degrees F
FPC 10 IA 0 Chip	OK	36 degrees C / 96 degrees F
FPC 10 IA 1 TSensor	OK	43 degrees C / 109 degrees F
FPC 10 IA 1 Chip	OK	42 degrees C / 107 degrees F
Fans Top Fan Tray Temp	OK	37 degrees C / 98 degrees F
Top Tray Fan 1	OK	Spinning at normal speed
Top Tray Fan 2	OK	Spinning at normal speed
Top Tray Fan 3	OK	Spinning at normal speed
Top Tray Fan 4	OK	Spinning at normal speed
Top Tray Fan 5	OK	Spinning at normal speed
Top Tray Fan 6	OK	Spinning at normal speed
Bottom Fan Tray Temp	OK	28 degrees C / 82 degrees F
Bottom Tray Fan 1	OK	Spinning at normal speed
Bottom Tray Fan 2	OK	Spinning at normal speed
Bottom Tray Fan 3	OK	Spinning at normal speed
Bottom Tray Fan 4	OK	Spinning at normal speed
Bottom Tray Fan 5	OK	Spinning at normal speed
Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX2020 Router)

```

user@host> show chassis environment

```

Class	Item	Status	Measurement
Temp	PSM 0	Absent	
	PSM 1	Absent	
	PSM 2	OK	41 degrees C / 105 degrees F
	PSM 3	OK	39 degrees C / 102 degrees F
	PSM 4	OK	39 degrees C / 102 degrees F
	PSM 5	OK	38 degrees C / 100 degrees F
	PSM 6	OK	38 degrees C / 100 degrees F
	PSM 7	OK	38 degrees C / 100 degrees F

PSM 8	OK	37 degrees C / 98 degrees F
PSM 9	Absent	
PSM 10	Absent	
PSM 11	OK	47 degrees C / 116 degrees F
PSM 12	OK	45 degrees C / 113 degrees F
PSM 13	OK	44 degrees C / 111 degrees F
PSM 14	OK	44 degrees C / 111 degrees F
PSM 15	OK	43 degrees C / 109 degrees F
PSM 16	OK	42 degrees C / 107 degrees F
PSM 17	OK	41 degrees C / 105 degrees F
PDM 0	OK	
PDM 1	Absent	
PDM 2	Absent	
PDM 3	OK	
CB 0 IntakeA-Zone0	OK	45 degrees C / 113 degrees F
CB 0 IntakeB-Zone1	OK	34 degrees C / 93 degrees F
CB 0 IntakeC-Zone0	OK	48 degrees C / 118 degrees F
CB 0 ExhaustA-Zone0	OK	45 degrees C / 113 degrees F
CB 0 ExhaustB-Zone1	OK	37 degrees C / 98 degrees F
CB 0 TCBC-Zone0	OK	41 degrees C / 105 degrees F
CB 1 IntakeA-Zone0	OK	46 degrees C / 114 degrees F
CB 1 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
CB 1 IntakeC-Zone0	OK	49 degrees C / 120 degrees F
CB 1 ExhaustA-Zone0	OK	46 degrees C / 114 degrees F
CB 1 ExhaustB-Zone1	OK	41 degrees C / 105 degrees F
CB 1 TCBC-Zone0	OK	46 degrees C / 114 degrees F
SPMB 0 Intake	OK	33 degrees C / 91 degrees F
SPMB 1 Intake	OK	42 degrees C / 107 degrees F
Routing Engine 0	OK	35 degrees C / 95 degrees F
Routing Engine 0 CPU	OK	34 degrees C / 93 degrees F
Routing Engine 1	OK	44 degrees C / 111 degrees F
Routing Engine 1 CPU	OK	42 degrees C / 107 degrees F
SFB 0 Intake-Zone0	OK	55 degrees C / 131 degrees F
SFB 0 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 0 IntakeA-Zone0	OK	50 degrees C / 122 degrees F
SFB 0 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 0 Exhaust-Zone0	OK	52 degrees C / 125 degrees F
SFB 0 SFB-XF2-Zone1	OK	61 degrees C / 141 degrees F
SFB 0 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 0 SFB-XF0-Zone0	OK	68 degrees C / 154 degrees F
SFB 1 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 1 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 1 Exhaust-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 SFB-XF2-Zone1	OK	62 degrees C / 143 degrees F
SFB 1 SFB-XF1-Zone0	OK	67 degrees C / 152 degrees F
SFB 1 SFB-XF0-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 2 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 2 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 2 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 2 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 SFB-XF0-Zone0	OK	70 degrees C / 158 degrees F
SFB 3 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 3 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 3 IntakeA-Zone0	OK	52 degrees C / 125 degrees F
SFB 3 IntakeB-Zone1	OK	41 degrees C / 105 degrees F
SFB 3 Exhaust-Zone0	OK	53 degrees C / 127 degrees F

SFB 3 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 3 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 3 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 4 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 4 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 4 IntakeA-Zone0	OK	54 degrees C / 129 degrees F
SFB 4 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
SFB 4 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 4 SFB-XF2-Zone1	OK	64 degrees C / 147 degrees F
SFB 4 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 4 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 5 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 5 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 5 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 5 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 5 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 5 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 5 SFB-XF0-Zone0	OK	74 degrees C / 165 degrees F
SFB 6 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 6 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 6 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 6 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 6 SFB-XF0-Zone0	OK	72 degrees C / 161 degrees F
SFB 7 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 7 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 7 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 7 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 7 SFB-XF2-Zone1	OK	68 degrees C / 154 degrees F
SFB 7 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 7 SFB-XF0-Zone0	OK	73 degrees C / 163 degrees F
FPC 0 Intake	OK	41 degrees C / 105 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 0 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 0 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 0 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 0 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 0 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 0 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 0 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 1 Intake	OK	40 degrees C / 104 degrees F
FPC 1 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 1 Exhaust B	OK	58 degrees C / 136 degrees F
FPC 1 LU 0 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 1 Chip	OK	58 degrees C / 136 degrees F

FPC 1 LU 2 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 LU 3 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 1 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 1 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 Intake	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 2 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 2 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 2 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 2 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 2 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 2 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 2 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 3 Intake	OK	40 degrees C / 104 degrees F
FPC 3 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 3 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 3 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 3 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 3 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 3 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 3 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 3 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 3 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 3 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 4 Intake	OK	40 degrees C / 104 degrees F
FPC 4 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 4 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 4 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 4 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 4 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 4 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 3 Chip	OK	53 degrees C / 127 degrees F

FPC 4 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 4 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 4 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 5 Intake	OK	41 degrees C / 105 degrees F
FPC 5 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 5 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 5 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 0 Chip	OK	63 degrees C / 145 degrees F
FPC 5 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 5 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 5 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 5 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 5 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 6 Intake	OK	42 degrees C / 107 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 6 LU 0 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 6 LU 1 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 6 LU 2 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 6 LU 3 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 3 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 6 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 6 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 7 Intake	OK	41 degrees C / 105 degrees F
FPC 7 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 7 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 7 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 7 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 7 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 7 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 0 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 1 Chip	OK	54 degrees C / 129 degrees F

FPC 7 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 7 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 8 Intake	OK	41 degrees C / 105 degrees F
FPC 8 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 8 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 8 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 8 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 8 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 2 Chip	OK	55 degrees C / 131 degrees F
FPC 8 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 8 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 8 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 8 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 9 Intake	OK	42 degrees C / 107 degrees F
FPC 9 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 9 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 9 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 0 Chip	OK	65 degrees C / 149 degrees F
FPC 9 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 1 Chip	OK	67 degrees C / 152 degrees F
FPC 9 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 9 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 0 Chip	OK	55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 10 Intake	OK	44 degrees C / 111 degrees F
FPC 10 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 10 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 10 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 0 Chip	OK	55 degrees C / 131 degrees F
FPC 10 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 1 Chip	OK	59 degrees C / 138 degrees F
FPC 10 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 2 Chip	OK	52 degrees C / 125 degrees F
FPC 10 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip	OK	47 degrees C / 116 degrees F

FPC 11 Intake	OK	30 degrees C / 86 degrees F
FPC 11 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 11 Exhaust B	OK	30 degrees C / 86 degrees F
FPC 11 LU 0 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 0 Chip	OK	58 degrees C / 136 degrees F
FPC 11 LU 1 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 11 LU 2 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 11 LU 3 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip	OK	57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 11 MQ 3 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 3 Chip	OK	52 degrees C / 125 degrees F
FPC 12 Intake	OK	40 degrees C / 104 degrees F
FPC 12 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 12 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 12 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 12 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 12 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 12 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 3 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 12 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 13 Intake	OK	40 degrees C / 104 degrees F
FPC 13 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 13 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 13 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 13 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 13 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 13 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 13 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 13 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 13 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 14 Intake	OK	40 degrees C / 104 degrees F
FPC 14 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 14 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 14 LU 0 TSen	OK	50 degrees C / 122 degrees F

FPC 14 LU 0 Chip	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 Chip	OK	54 degrees C / 129 degrees F
FPC 14 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 14 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 14 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 14 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 14 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 14 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 15 Intake	OK	44 degrees C / 111 degrees F
FPC 15 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 15 Exhaust B	OK	60 degrees C / 140 degrees F
FPC 15 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 Chip	OK	58 degrees C / 136 degrees F
FPC 15 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 3 Chip	OK	63 degrees C / 145 degrees F
FPC 15 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XM 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 XF 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XF 0 Chip	OK	68 degrees C / 154 degrees F
FPC 15 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 15 PLX Switch Chip	OK	56 degrees C / 132 degrees F
FPC 16 Intake	OK	42 degrees C / 107 degrees F
FPC 16 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 16 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 16 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 16 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 16 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 16 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 16 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 16 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 16 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 17 Intake	OK	43 degrees C / 109 degrees F
FPC 17 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 17 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 17 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 17 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 1 Chip	OK	60 degrees C / 140 degrees F
FPC 17 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 17 LU 3 TSen	OK	54 degrees C / 129 degrees F

FPC 17 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 17 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 17 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 17 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 17 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 3 Chip	OK	51 degrees C / 123 degrees F
FPC 18 Intake	OK	44 degrees C / 111 degrees F
FPC 18 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 18 Exhaust B	OK	57 degrees C / 134 degrees F
FPC 18 LU 0 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 18 LU 1 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 18 LU 2 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 18 LU 3 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 3 Chip	OK	55 degrees C / 131 degrees F
FPC 18 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 0 Chip	OK	54 degrees C / 129 degrees F
FPC 18 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip	OK	50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip	OK	53 degrees C / 127 degrees F
FPC 19 Intake	OK	48 degrees C / 118 degrees F
FPC 19 Exhaust A	OK	56 degrees C / 132 degrees F
FPC 19 Exhaust B	OK	64 degrees C / 147 degrees F
FPC 19 LU 0 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 19 LU 1 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 1 Chip	OK	70 degrees C / 158 degrees F
FPC 19 LU 2 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 2 Chip	OK	61 degrees C / 141 degrees F
FPC 19 LU 3 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 3 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip	OK	60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip	OK	57 degrees C / 134 degrees F
ADC 0 Intake	OK	40 degrees C / 104 degrees F
ADC 0 Exhaust	OK	52 degrees C / 125 degrees F
ADC 0 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 0 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 1 Intake	OK	38 degrees C / 100 degrees F
ADC 1 Exhaust	OK	50 degrees C / 122 degrees F
ADC 1 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 1 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 2 Intake	OK	37 degrees C / 98 degrees F
ADC 2 Exhaust	OK	52 degrees C / 125 degrees F
ADC 2 ADC-XF1	OK	53 degrees C / 127 degrees F
ADC 2 ADC-XF0	OK	61 degrees C / 141 degrees F
ADC 3 Intake	OK	40 degrees C / 104 degrees F
ADC 3 Exhaust	OK	51 degrees C / 123 degrees F

ADC 3 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 3 ADC-XF0	OK	64 degrees C / 147 degrees F
ADC 4 Intake	OK	39 degrees C / 102 degrees F
ADC 4 Exhaust	OK	51 degrees C / 123 degrees F
ADC 4 ADC-XF1	OK	60 degrees C / 140 degrees F
ADC 4 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 5 Intake	OK	38 degrees C / 100 degrees F
ADC 5 Exhaust	OK	54 degrees C / 129 degrees F
ADC 5 ADC-XF1	OK	56 degrees C / 132 degrees F
ADC 5 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 6 Intake	OK	39 degrees C / 102 degrees F
ADC 6 Exhaust	OK	52 degrees C / 125 degrees F
ADC 6 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 6 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 7 Intake	OK	39 degrees C / 102 degrees F
ADC 7 Exhaust	OK	54 degrees C / 129 degrees F
ADC 7 ADC-XF1	OK	62 degrees C / 143 degrees F
ADC 7 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 8 Intake	OK	39 degrees C / 102 degrees F
ADC 8 Exhaust	OK	52 degrees C / 125 degrees F
ADC 8 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 8 ADC-XF0	OK	65 degrees C / 149 degrees F
ADC 9 Intake	OK	41 degrees C / 105 degrees F
ADC 9 Exhaust	OK	51 degrees C / 123 degrees F
ADC 9 ADC-XF1	OK	63 degrees C / 145 degrees F
ADC 9 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 10 Intake	OK	48 degrees C / 118 degrees F
ADC 10 Exhaust	OK	53 degrees C / 127 degrees F
ADC 10 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 10 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 12 Intake	OK	49 degrees C / 120 degrees F
ADC 12 Exhaust	OK	54 degrees C / 129 degrees F
ADC 12 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 12 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 13 Intake	OK	49 degrees C / 120 degrees F
ADC 13 Exhaust	OK	57 degrees C / 134 degrees F
ADC 13 ADC-XF1	OK	66 degrees C / 150 degrees F
ADC 13 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 14 Intake	OK	51 degrees C / 123 degrees F
ADC 14 Exhaust	OK	59 degrees C / 138 degrees F
ADC 14 ADC-XF1	OK	69 degrees C / 156 degrees F
ADC 14 ADC-XF0	OK	74 degrees C / 165 degrees F
ADC 15 Intake	OK	50 degrees C / 122 degrees F
ADC 15 Exhaust	OK	59 degrees C / 138 degrees F
ADC 15 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 15 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 16 Intake	OK	52 degrees C / 125 degrees F
ADC 16 Exhaust	OK	58 degrees C / 136 degrees F
ADC 16 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 16 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 17 Intake	OK	52 degrees C / 125 degrees F
ADC 17 Exhaust	OK	59 degrees C / 138 degrees F
ADC 17 ADC-XF1	OK	69 degrees C / 156 degrees F
ADC 17 ADC-XF0	OK	71 degrees C / 159 degrees F
ADC 18 Intake	OK	53 degrees C / 127 degrees F
ADC 18 Exhaust	OK	59 degrees C / 138 degrees F
ADC 18 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 18 ADC-XF0	OK	73 degrees C / 163 degrees F
ADC 19 Intake	OK	50 degrees C / 122 degrees F
ADC 19 Exhaust	OK	59 degrees C / 138 degrees F
ADC 19 ADC-XF1	OK	68 degrees C / 154 degrees F

	ADC 19 ADC-XF0	OK	72 degrees C / 161 degrees F
Fans	Fan Tray 0 Fan 1	OK	7440 RPM
	Fan Tray 0 Fan 2	OK	7200 RPM
	Fan Tray 0 Fan 3	OK	6960 RPM
	Fan Tray 0 Fan 4	OK	7200 RPM
	Fan Tray 0 Fan 5	OK	7080 RPM
	Fan Tray 0 Fan 6	OK	6840 RPM
	Fan Tray 1 Fan 1	OK	6840 RPM
	Fan Tray 1 Fan 2	OK	6960 RPM
	Fan Tray 1 Fan 3	OK	6960 RPM
	Fan Tray 1 Fan 4	OK	7080 RPM
	Fan Tray 1 Fan 5	OK	6960 RPM
	Fan Tray 1 Fan 6	OK	6960 RPM
	Fan Tray 2 Fan 1	OK	8640 RPM
	Fan Tray 2 Fan 2	OK	8640 RPM
	Fan Tray 2 Fan 3	OK	8760 RPM
	Fan Tray 2 Fan 4	OK	8760 RPM
	Fan Tray 2 Fan 5	OK	8640 RPM
	Fan Tray 2 Fan 6	OK	8640 RPM
	Fan Tray 3 Fan 1	OK	8520 RPM
	Fan Tray 3 Fan 2	OK	8520 RPM
	Fan Tray 3 Fan 3	OK	8640 RPM
	Fan Tray 3 Fan 4	OK	8640 RPM
	Fan Tray 3 Fan 5	OK	8520 RPM
	Fan Tray 3 Fan 6	OK	8520 RPM

show chassis environment (MX2010 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PSM 0	OK	7 degrees C / 44 degrees F
	PSM 1	OK	7 degrees C / 44 degrees F
	PSM 2	OK	7 degrees C / 44 degrees F
	PSM 3	OK	6 degrees C / 42 degrees F
	PSM 4	OK	6 degrees C / 42 degrees F
	PSM 5	OK	6 degrees C / 42 degrees F
	PSM 6	OK	6 degrees C / 42 degrees F
	PSM 7	OK	7 degrees C / 44 degrees F
	PSM 8	OK	7 degrees C / 44 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	CB 0 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 IntakeB-Zone1	OK	7 degrees C / 44 degrees F
	CB 0 IntakeC-Zone0	OK	22 degrees C / 71 degrees F
	CB 0 ExhaustA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 ExhaustB-Zone1	OK	9 degrees C / 48 degrees F
	CB 0 TCBC-Zone0	OK	11 degrees C / 51 degrees F
	CB 1 IntakeA-Zone0	OK	9 degrees C / 48 degrees F
	CB 1 IntakeB-Zone1	OK	5 degrees C / 41 degrees F
	CB 1 IntakeC-Zone0	OK	20 degrees C / 68 degrees F
	CB 1 ExhaustA-Zone0	OK	12 degrees C / 53 degrees F
	CB 1 ExhaustB-Zone1	OK	7 degrees C / 44 degrees F
	CB 1 TCBC-Zone0	OK	10 degrees C / 50 degrees F
	SPMB 0 Intake	OK	5 degrees C / 41 degrees F
	SPMB 1 Intake	OK	4 degrees C / 39 degrees F
	Routing Engine 0	OK	9 degrees C / 48 degrees F
	Routing Engine 0 CPU	OK	9 degrees C / 48 degrees F
	Routing Engine 1	OK	6 degrees C / 42 degrees F
	Routing Engine 1 CPU	OK	6 degrees C / 42 degrees F
	SFB 0 Intake-Zone0	OK	26 degrees C / 78 degrees F
	SFB 0 Exhaust-Zone1	OK	17 degrees C / 62 degrees F

SFB 0 IntakeA-Zone0	OK	16 degrees C / 60 degrees F
SFB 0 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 0 Exhaust-Zone0	OK	18 degrees C / 64 degrees F
SFB 0 SFB-XF2-Zone1	OK	25 degrees C / 77 degrees F
SFB 0 SFB-XF1-Zone0	OK	23 degrees C / 73 degrees F
SFB 0 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 1 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone1	OK	15 degrees C / 59 degrees F
SFB 1 IntakeA-Zone0	OK	20 degrees C / 68 degrees F
SFB 1 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 1 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 1 SFB-XF2-Zone1	OK	26 degrees C / 78 degrees F
SFB 1 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 2 Intake-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 2 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 2 IntakeB-Zone1	OK	9 degrees C / 48 degrees F
SFB 2 Exhaust-Zone0	OK	16 degrees C / 60 degrees F
SFB 2 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 2 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 SFB-XF0-Zone0	OK	26 degrees C / 78 degrees F
SFB 4 Intake-Zone0	OK	28 degrees C / 82 degrees F
SFB 4 Exhaust-Zone1	OK	16 degrees C / 60 degrees F
SFB 4 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 4 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 4 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 4 SFB-XF2-Zone1	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 5 Intake-Zone0	OK	22 degrees C / 71 degrees F
SFB 5 Exhaust-Zone1	OK	14 degrees C / 57 degrees F
SFB 5 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 5 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 5 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 5 SFB-XF2-Zone1	OK	22 degrees C / 71 degrees F
SFB 5 SFB-XF1-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF0-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 6 IntakeA-Zone0	OK	19 degrees C / 66 degrees F
SFB 6 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 6 Exhaust-Zone0	OK	20 degrees C / 68 degrees F
SFB 6 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 SFB-XF1-Zone0	OK	32 degrees C / 89 degrees F
SFB 6 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 7 Intake-Zone0	OK	25 degrees C / 77 degrees F
SFB 7 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 7 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
SFB 7 IntakeB-Zone1	OK	8 degrees C / 46 degrees F
SFB 7 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 7 SFB-XF2-Zone1	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
FPC 0 Intake	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust A	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust B	OK	14 degrees C / 57 degrees F
FPC 0 LU 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 0 LU 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 1 Chip	OK	27 degrees C / 80 degrees F

FPC 0 LU 2 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 0 LU 3 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 3 Chip	OK	23 degrees C / 73 degrees F
FPC 0 XM 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 0 Chip	OK	33 degrees C / 91 degrees F
FPC 0 XM 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 0 PLX Switch TSen	OK	28 degrees C / 82 degrees F
FPC 0 PLX Switch Chip	OK	26 degrees C / 78 degrees F
FPC 1 Intake	OK	10 degrees C / 50 degrees F
FPC 1 Exhaust A	OK	24 degrees C / 75 degrees F
FPC 1 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 1 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 0 Chip	OK	31 degrees C / 87 degrees F
FPC 1 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 1 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 2 Chip	OK	25 degrees C / 77 degrees F
FPC 1 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 1 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 1 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 1 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 2 Intake	OK	9 degrees C / 48 degrees F
FPC 2 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 2 Exhaust B	OK	10 degrees C / 50 degrees F
FPC 2 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 2 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 Chip	OK	17 degrees C / 62 degrees F
FPC 2 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 2 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 2 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch Chip	OK	20 degrees C / 68 degrees F
FPC 3 Intake	OK	12 degrees C / 53 degrees F
FPC 3 Exhaust A	OK	16 degrees C / 60 degrees F
FPC 3 Exhaust B	OK	26 degrees C / 78 degrees F
FPC 3 LU 0 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 3 LU 1 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 3 LU 2 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 2 Chip	OK	22 degrees C / 71 degrees F
FPC 3 LU 3 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 3 MQ 0 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 0 Chip	OK	18 degrees C / 64 degrees F
FPC 3 MQ 1 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 3 MQ 2 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 2 Chip	OK	17 degrees C / 62 degrees F

FPC 3 MQ 3 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 4 Intake	OK	11 degrees C / 51 degrees F
FPC 4 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 4 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 4 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 4 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 4 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 4 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 4 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 4 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 4 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 4 PLX Switch Chip	OK	23 degrees C / 73 degrees F
FPC 5 Intake	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust A	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust B	OK	12 degrees C / 53 degrees F
FPC 5 LU 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 0 Chip	OK	28 degrees C / 82 degrees F
FPC 5 LU 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 5 LU 3 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 5 XM 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 0 Chip	OK	36 degrees C / 96 degrees F
FPC 5 XM 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 5 PLX Switch TSen	OK	27 degrees C / 80 degrees F
FPC 5 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 6 Intake	OK	12 degrees C / 53 degrees F
FPC 6 Exhaust A	OK	17 degrees C / 62 degrees F
FPC 6 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 6 LU 0 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 0 Chip	OK	29 degrees C / 84 degrees F
FPC 6 LU 1 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 1 Chip	OK	30 degrees C / 86 degrees F
FPC 6 LU 2 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 2 Chip	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 6 MQ 0 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 0 Chip	OK	19 degrees C / 66 degrees F
FPC 6 MQ 1 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 6 MQ 2 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 6 MQ 3 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 7 Intake	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 7 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 TSen	OK	26 degrees C / 78 degrees F

FPC 7 LU 1 Chip	OK	29 degrees C / 84 degrees F
FPC 7 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 7 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 3 Chip	OK	24 degrees C / 75 degrees F
FPC 7 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 7 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 1 Chip	OK	32 degrees C / 89 degrees F
FPC 7 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 7 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 8 Intake	OK	10 degrees C / 50 degrees F
FPC 8 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 8 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 8 LU 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 8 LU 1 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 1 Chip	OK	23 degrees C / 73 degrees F
FPC 8 LU 2 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 8 LU 3 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 8 XM 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XM 0 Chip	OK	29 degrees C / 84 degrees F
FPC 8 XF 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XF 0 Chip	OK	38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen	OK	20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 9 Intake	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust A	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 9 LU 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 0 Chip	OK	24 degrees C / 75 degrees F
FPC 9 LU 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 9 LU 2 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 2 Chip	OK	16 degrees C / 60 degrees F
FPC 9 LU 3 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 9 XM 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 0 Chip	OK	32 degrees C / 89 degrees F
FPC 9 XM 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 1 Chip	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip	OK	21 degrees C / 69 degrees F
ADC 0 Intake	OK	12 degrees C / 53 degrees F
ADC 0 Exhaust	OK	20 degrees C / 68 degrees F
ADC 0 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 0 ADC-XF0	OK	32 degrees C / 89 degrees F
ADC 1 Intake	OK	11 degrees C / 51 degrees F
ADC 1 Exhaust	OK	21 degrees C / 69 degrees F
ADC 1 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 1 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 2 Intake	OK	14 degrees C / 57 degrees F
ADC 2 Exhaust	OK	21 degrees C / 69 degrees F
ADC 2 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 2 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 3 Intake	OK	13 degrees C / 55 degrees F
ADC 3 Exhaust	OK	19 degrees C / 66 degrees F
ADC 3 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 3 ADC-XF0	OK	31 degrees C / 87 degrees F

	ADC 4 Intake	OK	9 degrees C / 48 degrees F
	ADC 4 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 4 ADC-XF1	OK	28 degrees C / 82 degrees F
	ADC 4 ADC-XF0	OK	35 degrees C / 95 degrees F
	ADC 5 Intake	OK	12 degrees C / 53 degrees F
	ADC 5 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 5 ADC-XF1	OK	28 degrees C / 82 degrees F
	ADC 5 ADC-XF0	OK	34 degrees C / 93 degrees F
	ADC 6 Intake	OK	11 degrees C / 51 degrees F
	ADC 6 Exhaust	OK	21 degrees C / 69 degrees F
	ADC 6 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 6	ADC-XF0	OK	35 degrees C / 95 degrees F
	ADC 7 Intake	OK	14 degrees C / 57 degrees F
	ADC 7 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 7 ADC-XF1	OK	26 degrees C / 78 degrees F
	ADC 7 ADC-XF0	OK	34 degrees C / 93 degrees F
	ADC 8 Intake	OK	14 degrees C / 57 degrees F
	ADC 8 Exhaust	OK	21 degrees C / 69 degrees F
	ADC 8 ADC-XF1	OK	24 degrees C / 75 degrees F
	ADC 8 ADC-XF0	OK	31 degrees C / 87 degrees F
	ADC 9 Intake	OK	10 degrees C / 50 degrees F
	ADC 9 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 9 ADC-XF1	OK	28 degrees C / 82 degrees F
	ADC 9 ADC-XF0	OK	36 degrees C / 96 degrees F
Fans	Fan Tray 0 Fan 1	OK	3480 RPM
	Fan Tray 0 Fan 2	OK	3480 RPM
	Fan Tray 0 Fan 3	OK	3480 RPM
	Fan Tray 0 Fan 4	OK	3360 RPM
	Fan Tray 0 Fan 5	OK	3360 RPM
	Fan Tray 0 Fan 6	OK	3480 RPM
	Fan Tray 1 Fan 1	OK	3360 RPM
	Fan Tray 1 Fan 2	OK	3360 RPM
	Fan Tray 1 Fan 3	OK	3360 RPM
	Fan Tray 1 Fan 4	OK	3480 RPM
	Fan Tray 1 Fan 5	OK	3480 RPM
	Fan Tray 1 Fan 6	OK	3480 RPM
	Fan Tray 2 Fan 1	OK	3360 RPM
	Fan Tray 2 Fan 2	OK	3360 RPM
	Fan Tray 2 Fan 3	OK	3480 RPM
	Fan Tray 2 Fan 4	OK	3480 RPM
	Fan Tray 2 Fan 5	OK	3360 RPM
	Fan Tray 2 Fan 6	OK	3480 RPM
	Fan Tray 3 Fan 1	OK	3360 RPM
	Fan Tray 3 Fan 2	OK	3360 RPM
	Fan Tray 3 Fan 3	OK	3480 RPM
	Fan Tray 3 Fan 4	OK	3480 RPM
	Fan Tray 3 Fan 5	OK	3480 RPM
	Fan Tray 3 Fan 6	OK	3360 RPM

show chassis environment (T320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	SCG 0	OK	28 degrees C / 82 degrees F
	SCG 1	OK	28 degrees C / 82 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F

	SIB 0	OK	33 degrees C / 91 degrees F
	SIB 1	OK	33 degrees C / 91 degrees F
	SIB 2	OK	34 degrees C / 93 degrees F
	FPC 0 Top	OK	38 degrees C / 100 degrees F
	FPC 0 Bottom	OK	32 degrees C / 89 degrees F
	FPC 1 Top	OK	38 degrees C / 100 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPC 2 Top	OK	36 degrees C / 96 degrees F
	FPC 2 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	26 degrees C / 78 degrees F
	FPM Display	OK	29 degrees C / 84 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Middle fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (T640 Router)

```

user@host> show chassis environment

```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	22 degrees C / 71 degrees F
	SCG 0	OK	30 degrees C / 86 degrees F
	SCG 1	OK	30 degrees C / 86 degrees F
	Routing Engine 0	Present	
	Routing Engine 1	OK	27 degrees C / 80 degrees F
	CB 0	Present	
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	Absent	
	SIB 1	Absent	
	SIB 2	Absent	
	SIB 3	Absent	
	SIB 4	Absent	
	FPC 4 Top	Testing	
	FPC 4 Bottom	Testing	
	FPC 5 Top	Testing	
	FPC 5 Bottom	Testing	
	FPC 6 Top	Testing	
	FPC 6 Bottom	Testing	
	FPM GBUS	OK	23 degrees C / 73 degrees F
	FPM Display	Absent	
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed

Top Right Front fan	OK	Spinning at normal speed
Top Right Middle fan	OK	Spinning at normal speed
Top Right Rear fan	OK	Spinning at normal speed
Bottom Left Front fan	OK	Spinning at normal speed
Bottom Left Middle fan	OK	Spinning at normal speed
Bottom Left Rear fan	OK	Spinning at normal speed
Bottom Right Front fan	OK	Spinning at normal speed
Bottom Right Middle fan	OK	Spinning at normal speed
Bottom Right Rear fan	OK	Spinning at normal speed
Fourth Blower from top	OK	Spinning at normal speed
Bottom Blower	OK	Spinning at normal speed
Middle Blower	OK	Spinning at normal speed
Top Blower	OK	Spinning at normal speed
Second Blower from top	OK	Spinning at normal speed
Misc CIP	OK	
SPMB 0	OK	
SPMB 1	OK	

show chassis environment (T4000 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	SCG 0	OK	33 degrees C / 91 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU	OK	50 degrees C / 122 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	46 degrees C / 114 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	OK	42 degrees C / 107 degrees F
	SIB 1	OK	42 degrees C / 107 degrees F
	SIB 2	OK	42 degrees C / 107 degrees F
	SIB 3	OK	43 degrees C / 109 degrees F
	SIB 4	OK	45 degrees C / 113 degrees F
	FPC 0 Fan Intake	OK	34 degrees C / 93 degrees F
	FPC 0 Fan Exhaust	OK	48 degrees C / 118 degrees F
	FPC 0 PMB	OK	47 degrees C / 116 degrees F
	FPC 0 LMB0	OK	50 degrees C / 122 degrees F
	FPC 0 LMB1	OK	41 degrees C / 105 degrees F
	FPC 0 LMB2	OK	35 degrees C / 95 degrees F
	FPC 0 PFE1 LU2	OK	46 degrees C / 114 degrees F
	FPC 0 PFE1 LU0	OK	41 degrees C / 105 degrees F
	FPC 0 PFE0 LU0	OK	57 degrees C / 134 degrees F
	FPC 0 XF1	OK	46 degrees C / 114 degrees F
	FPC 0 XF0	OK	52 degrees C / 125 degrees F
	FPC 0 XM1	OK	41 degrees C / 105 degrees F
	FPC 0 XM0	OK	50 degrees C / 122 degrees F
	FPC 0 PFE0 LU1	OK	56 degrees C / 132 degrees F
	FPC 0 PFE0 LU2	OK	45 degrees C / 113 degrees F
	FPC 0 PFE1 LU1	OK	37 degrees C / 98 degrees F
	FPC 3 Fan Intake	OK	36 degrees C / 96 degrees F
	FPC 3 Fan Exhaust	OK	51 degrees C / 123 degrees F
	FPC 3 PMB	OK	43 degrees C / 109 degrees F
	FPC 3 LMB0	OK	57 degrees C / 134 degrees F
	FPC 3 LMB1	OK	54 degrees C / 129 degrees F
	FPC 3 LMB2	OK	38 degrees C / 100 degrees F
	FPC 3 PFE1 LU2	OK	63 degrees C / 145 degrees F

	FPC 3 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 3 PFE0 LU0	OK	69 degrees C / 156 degrees F
	FPC 3 XF1	OK	62 degrees C / 143 degrees F
	FPC 3 XF0	OK	63 degrees C / 145 degrees F
	FPC 3 XM1	OK	43 degrees C / 109 degrees F
	FPC 3 XM0	OK	67 degrees C / 152 degrees F
	FPC 3 PFE0 LU1	OK	63 degrees C / 145 degrees F
	FPC 3 PFE0 LU2	OK	66 degrees C / 150 degrees F
	FPC 3 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPC 5 Top	OK	39 degrees C / 102 degrees F
	FPC 5 Bottom	OK	38 degrees C / 100 degrees F
	FPC 6 Fan Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Fan Exhaust	OK	49 degrees C / 120 degrees F
	FPC 6 PMB	OK	40 degrees C / 104 degrees F
	FPC 6 LMB0	OK	60 degrees C / 140 degrees F
	FPC 6 LMB1	OK	58 degrees C / 136 degrees F
	FPC 6 LMB2	OK	40 degrees C / 104 degrees F
	FPC 6 PFE1 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 6 PFE0 LU0	OK	71 degrees C / 159 degrees F
	FPC 6 XF1	OK	58 degrees C / 136 degrees F
	FPC 6 XF0	OK	65 degrees C / 149 degrees F
	FPC 6 XM1	OK	39 degrees C / 102 degrees F
	FPC 6 XM0	OK	66 degrees C / 150 degrees F
	FPC 6 PFE0 LU1	OK	69 degrees C / 156 degrees F
	FPC 6 PFE0 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU1	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	27 degrees C / 80 degrees F
Fans	Top Left Front fan	OK	Spinning at high speed
	Top Left Middle fan	OK	Spinning at high speed
	Top Left Rear fan	OK	Spinning at high speed
	Top Right Front fan	OK	Spinning at high speed
	Top Right Middle fan	OK	Spinning at high speed
	Top Right Rear fan	OK	Spinning at high speed
	Bottom Left Front fan	OK	Spinning at high speed
	Bottom Left Middle fan	OK	Spinning at high speed
	Bottom Left Rear fan	OK	Spinning at high speed
	Bottom Right Front fan	OK	Spinning at high speed
	Bottom Right Middle fan	OK	Spinning at high speed
	Bottom Right Rear fan	OK	Spinning at high speed
	Rear Tray Top fan	OK	Spinning at high speed
	Rear Tray Second fan	OK	Spinning at high speed
	Rear Tray Third fan	OK	Spinning at high speed
	Rear Tray Fourth fan	OK	Spinning at high speed
	Rear Tray Fifth fan	OK	Spinning at high speed
	Rear Tray Sixth fan	OK	Spinning at high speed
	Rear Tray Seventh fan	OK	Spinning at high speed
	Rear Tray Bottom fan	OK	Spinning at high speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Router)

```
user@host> show chassis environment
scc-re0:
```

Class Item		Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	29 degrees C / 84 degrees F

	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	34 degrees C / 93 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	44 degrees C / 111 degrees F
	SIB 0 (B)	OK	44 degrees C / 111 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	32 degrees C / 89 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP 0	OK	
	CIP 1	OK	
	SPMB 0	OK	
	SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	35 degrees C / 95 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	36 degrees C / 96 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	40 degrees C / 104 degrees F
	SIB 0 (B)	OK	51 degrees C / 123 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	31 degrees C / 87 degrees F
	FPC 1 Top	OK	34 degrees C / 93 degrees F
	FPC 1 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed

```

Bottom Right Front fan OK      Spinning at normal speed
Bottom Right Middle fan OK     Spinning at normal speed
Bottom Right Rear fan OK       Spinning at normal speed
Rear Tray Top fan OK           Spinning at normal speed
Rear Tray Second fan OK        Spinning at normal speed
Rear Tray Third fan OK         Spinning at normal speed
Rear Tray Fourth fan OK        Spinning at normal speed
Rear Tray Fifth fan OK         Spinning at normal speed
Rear Tray Sixth fan OK         Spinning at normal speed
Rear Tray Seventh fan OK       Spinning at normal speed
Rear Tray Bottom fan OK        Spinning at normal speed
Misc CIP OK
SPMB 0 OK
SPMB 1 OK

```

```
lcc2-re0:
```

```

-----
Class Item      Status      Measurement
Temp PEM 0      OK          29 degrees C / 84 degrees F
      PEM 1      Absent
      SCG 0      OK          32 degrees C / 89 degrees F
      SCG 1      Absent
      Routing Engine 0 OK          31 degrees C / 87 degrees F
      Routing Engine 1 OK          32 degrees C / 89 degrees F
      CB 0       OK          30 degrees C / 86 degrees F
      SIB 0      OK          38 degrees C / 100 degrees F
      SIB 0 (B)  OK          49 degrees C / 120 degrees F
      FPC 0 Top  OK          45 degrees C / 113 degrees F
      FPC 0 Bottom OK          33 degrees C / 91 degrees F
      FPC 1 Top  OK          37 degrees C / 98 degrees F
      FPC 1 Bottom OK          33 degrees C / 91 degrees F
      FPM GBUS   OK          30 degrees C / 86 degrees F
      FPM Display OK          34 degrees C / 93 degrees F
Fans  Top Left Front fan OK      Spinning at normal speed
      Top Left Middle fan OK      Spinning at normal speed
...

```

show chassis environment (T1600 Router)

```

user@host> show chassis environment
Class Item      Status      Measurement
Temp PEM 0      OK          27 degrees C / 80 degrees F
      PEM 1      Absent
      SCG 0      OK          31 degrees C / 87 degrees F
      SCG 1      OK          35 degrees C / 95 degrees F
      Routing Engine 0 OK          30 degrees C / 86 degrees F
      Routing Engine 1 OK          30 degrees C / 86 degrees F
      CB 0       OK          31 degrees C / 87 degrees F
      CB 1       OK          31 degrees C / 87 degrees F
      SIB 0      OK          41 degrees C / 105 degrees F
      SIB 0 (B)  OK          34 degrees C / 93 degrees F
      SIB 1      OK          0 degrees C / 32 degrees F
      SIB 1 (B)  OK          0 degrees C / 32 degrees F
      SIB 2      OK          0 degrees C / 32 degrees F
      SIB 2 (B)  OK          0 degrees C / 32 degrees F
      SIB 3      OK          0 degrees C / 32 degrees F
      SIB 3 (B)  OK          0 degrees C / 32 degrees F
      SIB 4      OK          0 degrees C / 32 degrees F
      SIB 4 (B)  OK          0 degrees C / 32 degrees F
      FPC 0 Top  OK          49 degrees C / 120 degrees F
      FPC 0 Bottom OK          50 degrees C / 122 degrees F

```

	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Plus Router)

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	28 degrees C / 82 degrees F
	PEM 1	Absent	
	Routing Engine 0	OK	27 degrees C / 80 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	CB 0 Intake	OK	26 degrees C / 78 degrees F
	CB 0 Exhaust A	OK	25 degrees C / 77 degrees F
	CB 0 Exhaust B	OK	25 degrees C / 77 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	26 degrees C / 78 degrees F
	SIB F13 0	OK	47 degrees C / 116 degrees F
	SIB F13 0 (B)	OK	48 degrees C / 118 degrees F
	SIB F13 1	OK	38 degrees C / 100 degrees F
	SIB F13 1 (B)	OK	37 degrees C / 98 degrees F
	SIB F2S 0/0	OK	27 degrees C / 80 degrees F
	SIB F2S 0/2	OK	28 degrees C / 82 degrees F
	SIB F2S 0/4	OK	27 degrees C / 80 degrees F
	SIB F2S 0/6	OK	28 degrees C / 82 degrees F
	SIB F2S 1/0	OK	26 degrees C / 78 degrees F
	SIB F2S 1/2	OK	26 degrees C / 78 degrees F
	SIB F2S 1/4	OK	26 degrees C / 78 degrees F
	SIB F2S 1/6	OK	26 degrees C / 78 degrees F
	SIB F2S 2/0	OK	25 degrees C / 77 degrees F
	SIB F2S 2/2	OK	25 degrees C / 77 degrees F
	SIB F2S 2/4	OK	23 degrees C / 73 degrees F
	CIP 0 Intake	OK	23 degrees C / 73 degrees F
	CIP 0 Exhaust A	OK	24 degrees C / 75 degrees F

	CIP 0 Exhaust B	OK	24 degrees C / 75 degrees F
	CIP 1 Intake	OK	24 degrees C / 75 degrees F
	CIP 1 Exhaust A	OK	25 degrees C / 77 degrees F
	CIP 1 Exhaust B	OK	25 degrees C / 77 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	OK	Spinning at normal speed
Misc	SPMB 0	OK	
	SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	27 degrees C / 80 degrees F
	PEM 1	Absent	
	SCG 0	OK	31 degrees C / 87 degrees F

	SCG 1	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	30 degrees C / 86 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	31 degrees C / 87 degrees F
	CB 1	OK	31 degrees C / 87 degrees F
	SIB 0	OK	41 degrees C / 105 degrees F
	SIB 0 (B)	OK	34 degrees C / 93 degrees F
	SIB 1	OK	0 degrees C / 32 degrees F
	SIB 1 (B)	OK	0 degrees C / 32 degrees F
	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Check	30 degrees C / 86 degrees F
	PEM 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	28 degrees C / 82 degrees F
	Routing Engine 0 CPU	OK	42 degrees C / 107 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust A	OK	28 degrees C / 82 degrees F
	CB 0 Exhaust B	OK	30 degrees C / 86 degrees F
	CB 1 Intake	OK	31 degrees C / 87 degrees F

	CB 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 1 Exhaust B	OK	31 degrees C / 87 degrees F
	SIB F13 0 Board	OK	44 degrees C / 111 degrees F
	SIB F13 0 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F13 3 Board	OK	45 degrees C / 113 degrees F
	SIB F13 3 XF Junction	OK	60 degrees C / 140 degrees F
	SIB F13 6 Board	OK	47 degrees C / 116 degrees F
	SIB F13 6 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F2S 0/0 Board	OK	32 degrees C / 89 degrees F
	SIB F2S 0/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/2 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/2 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 0/4 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/6 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/0 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 1/0 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/2 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 1/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/4 XF Junction	OK	35 degrees C / 95 degrees F
	SIB F2S 1/6 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 1/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 2/0 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 2/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/2 Board	OK	28 degrees C / 82 degrees F
	SIB F2S 2/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 2/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/6 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/6 XF Junction	OK	41 degrees C / 105 degrees F
	CIP 0 Intake	OK	25 degrees C / 77 degrees F
	CIP 0 Exhaust A	OK	26 degrees C / 78 degrees F
	CIP 0 Exhaust B	OK	26 degrees C / 78 degrees F
	CIP 1 Intake	OK	26 degrees C / 78 degrees F
	CIP 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CIP 1 Exhaust B	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed

	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	Check	
Misc	SPMB 0	OK	
	SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Check	29 degrees C / 84 degrees F
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	32 degrees C / 89 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	49 degrees C / 120 degrees F
	CB 0	OK	34 degrees C / 93 degrees F
	CB 1	OK	34 degrees C / 93 degrees F
	SIB 0	OK	39 degrees C / 102 degrees F
	SIB 0 (B)	Absent	
	SIB 1	OK	39 degrees C / 102 degrees F
	SIB 1 (B)	Absent	
	SIB 2	OK	39 degrees C / 102 degrees F
	SIB 2 (B)	Absent	
	FPC 4 Top	OK	43 degrees C / 109 degrees F
	FPC 4 Bottom	OK	43 degrees C / 109 degrees F
	FPC 7 Fan Intake	OK	35 degrees C / 95 degrees F
	FPC 7 Fan Exhaust	OK	50 degrees C / 122 degrees F
	FPC 7 PMB	OK	50 degrees C / 122 degrees F
	FPC 7 LMB0	OK	55 degrees C / 131 degrees F
	FPC 7 LMB1	OK	49 degrees C / 120 degrees F
	FPC 7 LMB2	OK	39 degrees C / 102 degrees F
	FPC 7 PFE1 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 7 PFE0 LU0	OK	62 degrees C / 143 degrees F
	FPC 7 XF1	OK	52 degrees C / 125 degrees F
	FPC 7 XF0	OK	61 degrees C / 141 degrees F
	FPC 7 XM1	OK	39 degrees C / 102 degrees F

	FPC 7 XMO	OK	56 degrees C / 132 degrees F
	FPC 7 PFE0 LU1	OK	60 degrees C / 140 degrees F
	FPC 7 PFE0 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray fan 1 (Top)	OK	Spinning at normal speed
	Rear Tray fan 2	OK	Spinning at normal speed
	Rear Tray fan 3	OK	Spinning at normal speed
	Rear Tray fan 4	OK	Spinning at normal speed
Misc	Rear Tray fan 5	OK	Spinning at normal speed
	Rear Tray fan 6	OK	Spinning at normal speed
	Rear Tray fan 7	OK	Spinning at normal speed
	Rear Tray fan 8	OK	Spinning at normal speed
	Rear Tray fan 9	OK	Spinning at normal speed
	Rear Tray fan 10	OK	Spinning at normal speed
	Rear Tray fan 11	OK	Spinning at normal speed
	Rear Tray fan 12	OK	Spinning at normal speed
	Rear Tray fan 13	OK	Spinning at normal speed
	Rear Tray fan 14	OK	Spinning at normal speed
	Rear Tray fan 15	OK	Spinning at normal speed
	Rear Tray fan 16 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (EX4200 Standalone Switch)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	Absent	
Temp	FPC 0 CPU	OK	41 degrees C / 105 degrees F
	FPC 0 EX-PFE1	OK	42 degrees C / 107 degrees F
	FPC 0 EX-PFE2	OK	46 degrees C / 114 degrees F
	FPC 0 GEPHY Front Left	OK	25 degrees C / 77 degrees F
	FPC 0 GEPHY Front Right	OK	27 degrees C / 80 degrees F
	FPC 0 Uplink Conn	OK	29 degrees C / 84 degrees F
Fans	FPC 0 Fan 1	OK	Spinning at normal speed
	FPC 0 Fan 2	OK	Spinning at normal speed
	FPC 0 Fan 3	OK	Spinning at normal speed

show chassis environment (EX8216 Switch)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	PSU 0	OK	
	PSU 1	OK	
	PSU 2	OK	

	PSU 3	Check	
	PSU 4	Absent	
	PSU 5	Absent	
Temp	CB 0 Intake	OK	23 degrees C / 73 degrees F
	CB 0 Exhaust	OK	26 degrees C / 78 degrees F
	CB 1 Intake	OK	22 degrees C / 71 degrees F
	CB 1 Exhaust	OK	25 degrees C / 77 degrees F
	FPC 4 Intake	OK	49 degrees C / 120 degrees F
	FPC 4 Exhaust	OK	59 degrees C / 138 degrees F
	SIB 5 Intake	OK	25 degrees C / 77 degrees F
	SIB 5 Exhaust	OK	35 degrees C / 95 degrees F
	SIB 6 Intake	OK	25 degrees C / 77 degrees F
Fans	SIB 6 Exhaust	OK	38 degrees C / 100 degrees F
	Top Fan 1	OK	Spinning at normal speed
	Top Fan 2	OK	Spinning at normal speed
	Top Fan 3	OK	Spinning at normal speed
	Top Fan 4	OK	Spinning at normal speed
	Top Fan 5	OK	Spinning at normal speed
	Top Fan 6	OK	Spinning at normal speed
	Top Fan 7	OK	Spinning at normal speed
	Top Fan 8	OK	Spinning at normal speed
	Top Fan 9	OK	Spinning at normal speed
	Bottom Fan 1	OK	Spinning at normal speed
	Bottom Fan 2	OK	Spinning at normal speed
	Bottom Fan 3	OK	Spinning at normal speed
	Bottom Fan 4	OK	Spinning at normal speed
	Bottom Fan 5	OK	Spinning at normal speed
	Bottom Fan 6	OK	Spinning at normal speed
	Bottom Fan 7	OK	Spinning at normal speed
	Bottom Fan 8	OK	Spinning at normal speed
	Bottom Fan 9	OK	Spinning at normal speed

show chassis environment (QFX Series)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	OK	
Temp	FPC 0 Sensor TopLeft I	OK	26 degrees C / 78 degrees F
	FPC 0 Sensor TopRight I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor TopLeft E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopRight E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle I	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Bottom I	OK	34 degrees C / 93 degrees F
	FPC 0 Sensor Bottom E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Die Temp	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Mgmt Brd I	OK	24 degrees C / 75 degrees F
Fans	FPC 0 Sensor Switch I	OK	28 degrees C / 82 degrees F
	FPC 0 Fan 1 (left)	Failed	
	FPC 0 Fan 2 (right)	OK	Spinning at normal speed
	FPC 0 Fan 3 (middle)	OK	Spinning at normal speed

show chassis environment interconnect-device (QFabric System)

```
user@switch> show chassis environment interconnect-device IC-A0004
```

Class	Item	Status	Measurement
	CB 0		
	CB 0 L Intake	OK	30 degrees C / 86 degrees F
	CB 0 R Intake	OK	31 degrees C / 87 degrees F
	CB 0 L Exhaust	OK	32 degrees C / 89 degrees F

CB 0 R Exhaust	OK	33 degrees C / 91 degrees F
Routing Engine 0 CPU temp	OK	51 degrees C / 123 degrees F
CB 1		
CB 1 L Intake	OK	27 degrees C / 80 degrees F
CB 1 R Intake	OK	29 degrees C / 84 degrees F
CB 1 L Exhaust	OK	31 degrees C / 87 degrees F
CB 1 R Exhaust	OK	32 degrees C / 89 degrees F
Routing Engine 1 CPU temp	OK	40 degrees C / 104 degrees F
FC 0 FPC 0		
FPC 0 L Intake	OK	25 degrees C / 77 degrees F
FPC 0 R Intake	OK	28 degrees C / 82 degrees F
FPC 0 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 0 R Exhaust	OK	29 degrees C / 84 degrees F
FC 7 FPC 7		
FPC 7 L Intake	OK	25 degrees C / 77 degrees F
FPC 7 R Intake	OK	26 degrees C / 78 degrees F
FPC 7 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 7 R Exhaust	OK	29 degrees C / 84 degrees F
RC 0 FPC 8		
FPC 8 L Intake	OK	25 degrees C / 77 degrees F
FPC 8 R Intake	OK	26 degrees C / 78 degrees F
FPC 8 L Exhaust	OK	32 degrees C / 89 degrees F
FPC 8 R Exhaust	OK	30 degrees C / 86 degrees F
RC 7 FPC 15		
FPC 15 L Intake	OK	24 degrees C / 75 degrees F
FPC 15 R Intake	OK	25 degrees C / 77 degrees F
FPC 15 L Exhaust	OK	33 degrees C / 91 degrees F
FPC 15 R Exhaust	OK	31 degrees C / 87 degrees F
Fans TFT 0 Fan 0	OK	Spinning at normal speed
Fans TFT 0 Fan 1	OK	Spinning at normal speed
Fans TFT 0 Fan 2	OK	Spinning at normal speed
Fans TFT 0 Fan 3	OK	Spinning at normal speed
Fans TFT 0 Fan 4	OK	Spinning at normal speed
Fans TFT 0 Fan 5	OK	Spinning at normal speed
Fans BFT 1 Fan 0	OK	Spinning at normal speed
Fans BFT 1 Fan 1	OK	Spinning at normal speed
Fans BFT 1 Fan 2	OK	Spinning at normal speed
Fans BFT 1 Fan 3	Check	
Fans BFT 1 Fan 4	OK	Spinning at normal speed
Fans BFT 1 Fan 5	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 0	OK	Spinning at normal speed

Fans	SFT 2	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 2	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 2	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 3	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 3	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 3	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 3	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 1	OK	Spinning at normal speed
Power	PEM 0			OK	30 degrees C / 86 degrees F
Power	PEM 1			OK	30 degrees C / 86 degrees F
Power	PEM 2			OK	30 degrees C / 86 degrees F
Power	PEM 3			Absent	
Power	PEM 4			Absent	
Power	PEM 5			Absent	

show chassis environment node-device (QFabric System)

```

user@switch> show chassis environment node-device node1
Class Item                               Status Measurement
Power node1 Power Supply 0              Absent
        node1 Power Supply 1            Absent
Fans   node1 Fan Tray 0                  Testing
        node1 Fan Tray 1                Testing
        node1 Fan Tray 2                Testing

```


show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State          Check
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12          10       120     18
FPC 0 PEM 1 status:
  State          Online
  Airflow        Back to Front
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  11          10       110     17

```

show chassis environment (PTX5000 Packet Transport Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Temp  PDU 0                               OK
      PDU 0 PSM 0                         OK          36 degrees C / 96 degrees F
      PDU 0 PSM 1                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 2                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 3                         OK          37 degrees C / 98 degrees F
      PDU 1                               Absent
      CCG 0                               OK          44 degrees C / 111 degrees F
      CCG 1                               OK          44 degrees C / 111 degrees F
      Routing Engine 0                     OK          62 degrees C / 143 degrees F
      Routing Engine 0 CPU                  OK          75 degrees C / 167 degrees F
      Routing Engine 1                     OK          51 degrees C / 123 degrees F
      Routing Engine 1 CPU                  OK          64 degrees C / 147 degrees F
      CB 0 Intake                          OK          38 degrees C / 100 degrees F
      CB 0 Exhaust A                      OK          46 degrees C / 114 degrees F
      CB 0 Exhaust B                      OK          42 degrees C / 107 degrees F
      CB 1 Intake                          OK          35 degrees C / 95 degrees F
      CB 1 Exhaust A                      OK          39 degrees C / 102 degrees F
      CB 1 Exhaust B                      OK          36 degrees C / 96 degrees F
      SIB 0 Intake                        OK          39 degrees C / 102 degrees F
      SIB 0 Exhaust                      OK          37 degrees C / 98 degrees F
      SIB 0 Junction                     OK          43 degrees C / 109 degrees F
      SIB 1 Intake                        OK          39 degrees C / 102 degrees F
      SIB 1 Exhaust                      OK          36 degrees C / 96 degrees F
      SIB 1 Junction                     OK          46 degrees C / 114 degrees F
      SIB 2 Intake                        OK          37 degrees C / 98 degrees F
      SIB 2 Exhaust                      OK          37 degrees C / 98 degrees F
      SIB 2 Junction                     OK          42 degrees C / 107 degrees F
      SIB 3 Intake                        OK          40 degrees C / 104 degrees F
      SIB 3 Exhaust                      OK          40 degrees C / 104 degrees F
      SIB 3 Junction                     OK          45 degrees C / 113 degrees F
      SIB 4 Intake                        OK          47 degrees C / 116 degrees F
      SIB 4 Exhaust                      OK          44 degrees C / 111 degrees F
      SIB 4 Junction                     OK          58 degrees C / 136 degrees F
      SIB 5 Intake                        OK          58 degrees C / 136 degrees F
      SIB 5 Exhaust                      OK          43 degrees C / 109 degrees F
      SIB 5 Junction                     OK          71 degrees C / 159 degrees F
      SIB 6 Intake                        OK          57 degrees C / 134 degrees F
      SIB 6 Exhaust                      OK          42 degrees C / 107 degrees F
      SIB 6 Junction                     OK          65 degrees C / 149 degrees F

```

SIB 7 Intake	OK	58 degrees C / 136 degrees F
SIB 7 Exhaust	OK	42 degrees C / 107 degrees F
SIB 7 Junction	OK	66 degrees C / 150 degrees F
SIB 8 Intake	OK	57 degrees C / 134 degrees F
SIB 8 Exhaust	OK	42 degrees C / 107 degrees F
SIB 8 Junction	OK	70 degrees C / 158 degrees F
FPC 0 PMB	OK	35 degrees C / 95 degrees F
FPC 0 Intake	OK	33 degrees C / 91 degrees F
FPC 0 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 TL0	OK	48 degrees C / 118 degrees F
FPC 0 TQ0	OK	53 degrees C / 127 degrees F
FPC 0 TL1	OK	56 degrees C / 132 degrees F
FPC 0 TQ1	OK	58 degrees C / 136 degrees F
FPC 0 TL2	OK	55 degrees C / 131 degrees F
FPC 0 TQ2	OK	56 degrees C / 132 degrees F
FPC 0 TL3	OK	59 degrees C / 138 degrees F
FPC 0 TQ3	OK	59 degrees C / 138 degrees F
FPC 2 PMB	OK	35 degrees C / 95 degrees F
FPC 2 Intake	OK	34 degrees C / 93 degrees F
FPC 2 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 TL0	OK	53 degrees C / 127 degrees F
FPC 2 TQ0	OK	53 degrees C / 127 degrees F
FPC 2 TL1	OK	57 degrees C / 134 degrees F
FPC 2 TQ1	OK	58 degrees C / 136 degrees F
FPC 2 TL2	OK	54 degrees C / 129 degrees F
FPC 2 TQ2	OK	59 degrees C / 138 degrees F
FPC 2 TL3	OK	60 degrees C / 140 degrees F
FPC 2 TQ3	OK	64 degrees C / 147 degrees F
PIC 2/0 Ambient	OK	49 degrees C / 120 degrees F
FPC 3 PMB	OK	34 degrees C / 93 degrees F
FPC 3 Intake	OK	35 degrees C / 95 degrees F
FPC 3 Exhaust A	OK	54 degrees C / 129 degrees F
FPC 3 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 3 TL0	OK	49 degrees C / 120 degrees F
FPC 3 TQ0	OK	55 degrees C / 131 degrees F
FPC 3 TL1	OK	56 degrees C / 132 degrees F
FPC 3 TQ1	OK	58 degrees C / 136 degrees F
FPC 3 TL2	OK	56 degrees C / 132 degrees F
FPC 3 TQ2	OK	59 degrees C / 138 degrees F
FPC 3 TL3	OK	62 degrees C / 143 degrees F
FPC 3 TQ3	OK	63 degrees C / 145 degrees F
PIC 3/1	Absent	
FPC 5 PMB	OK	35 degrees C / 95 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 TL0	OK	54 degrees C / 129 degrees F
FPC 5 TQ0	OK	52 degrees C / 125 degrees F
FPC 5 TL1	OK	61 degrees C / 141 degrees F
FPC 5 TQ1	OK	60 degrees C / 140 degrees F
FPC 5 TL2	OK	55 degrees C / 131 degrees F
FPC 5 TQ2	OK	55 degrees C / 131 degrees F
FPC 5 TL3	OK	59 degrees C / 138 degrees F
FPC 5 TQ3	OK	58 degrees C / 136 degrees F
PIC 5/0 Ambient	OK	51 degrees C / 123 degrees F
PIC 5/1 Ambient	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/0	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/1	OK	36 degrees C / 96 degrees F
FPC 6 PMB	OK	36 degrees C / 96 degrees F

	FPC 6 Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
	FPC 6 Exhaust B	OK	39 degrees C / 102 degrees F
	FPC 6 TL0	OK	44 degrees C / 111 degrees F
	FPC 6 TQ0	OK	54 degrees C / 129 degrees F
	FPC 6 TL1	OK	59 degrees C / 138 degrees F
	FPC 6 TQ1	OK	58 degrees C / 136 degrees F
	FPC 6 TL2	OK	60 degrees C / 140 degrees F
	FPC 6 TQ2	OK	57 degrees C / 134 degrees F
	FPC 6 TL3	OK	65 degrees C / 149 degrees F
	FPC 6 TQ3	OK	60 degrees C / 140 degrees F
	FPC 7 PMB	OK	35 degrees C / 95 degrees F
	FPC 7 Intake	OK	33 degrees C / 91 degrees F
	FPC 7 Exhaust A	OK	53 degrees C / 127 degrees F
	FPC 7 Exhaust B	OK	40 degrees C / 104 degrees F
	FPC 7 TL0	OK	46 degrees C / 114 degrees F
	FPC 7 TQ0	OK	58 degrees C / 136 degrees F
	FPC 7 TL1	OK	53 degrees C / 127 degrees F
	FPC 7 TQ1	OK	59 degrees C / 138 degrees F
	FPC 7 TL2	OK	56 degrees C / 132 degrees F
	FPC 7 TQ2	OK	61 degrees C / 141 degrees F
	FPC 7 TL3	OK	63 degrees C / 145 degrees F
	FPC 7 TQ3	OK	63 degrees C / 145 degrees F
	FPM I2CS	OK	37 degrees C / 98 degrees F
Fans	Fan Tray 0 Fan 1	OK	3042 RPM
	Fan Tray 0 Fan 2	OK	3042 RPM
	Fan Tray 0 Fan 3	OK	3000 RPM
	Fan Tray 0 Fan 4	OK	3042 RPM
	Fan Tray 0 Fan 5	OK	3000 RPM
	Fan Tray 0 Fan 6	OK	3042 RPM
	Fan Tray 0 Fan 7	OK	3085 RPM
	Fan Tray 0 Fan 8	OK	3042 RPM
	Fan Tray 0 Fan 9	OK	3042 RPM
	Fan Tray 0 Fan 10	OK	3085 RPM
	Fan Tray 0 Fan 11	OK	3085 RPM
	Fan Tray 0 Fan 12	OK	3128 RPM
	Fan Tray 0 Fan 13	OK	3128 RPM
	Fan Tray 0 Fan 14	OK	3042 RPM
	Fan Tray 1 Fan 1	OK	2299 RPM
	Fan Tray 1 Fan 2	OK	2399 RPM
	Fan Tray 1 Fan 3	OK	2299 RPM
	Fan Tray 1 Fan 4	OK	2266 RPM
	Fan Tray 1 Fan 5	OK	2266 RPM
	Fan Tray 1 Fan 6	OK	2366 RPM
	Fan Tray 2 Fan 1	OK	2199 RPM
	Fan Tray 2 Fan 2	OK	2133 RPM
	Fan Tray 2 Fan 3	OK	2366 RPM
	Fan Tray 2 Fan 4	OK	2233 RPM
	Fan Tray 2 Fan 5	OK	2399 RPM
	Fan Tray 2 Fan 6	OK	2233 RPM
Misc	SPMB 0 Intake	OK	50 degrees C / 122 degrees F
	SPMB 1 Intake	OK	40 degrees C / 104 degrees F

show chassis environment (ACX2000 Universal Access Router)

```

user@host> show chassis environment
Class Item                Status    Measurement
PCB Left                 OK        44 degrees C / 111 degrees F
SFP+ Xcvr                OK        50 degrees C / 122 degrees F
FEB                       OK        70 degrees C / 158 degrees F
PCB Up                   OK        63 degrees C / 145 degrees F

```

PCB Mid	OK	66 degrees C / 150 degrees F
Telecom Mod	OK	65 degrees C / 149 degrees F
Routing Engine	OK	54 degrees C / 129 degrees F
Heater off		

show chassis environment (ACX4000 Universal Access Router)

On the ACX4000 router, the MIC output of the **show chassis environment** command varies depending on the number of temperature channels present in the installed MIC.

user@host> show chassis environment

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	PCB Bottom	OK	30 degrees C / 86 degrees F
	PCB Middle	OK	34 degrees C / 93 degrees F
	BCM56445	OK	33 degrees C / 91 degrees F
	SFP+ Xcvr	OK	32 degrees C / 89 degrees F
	Fan tray inlet	OK	39 degrees C / 102 degrees F
	Exhaust	OK	30 degrees C / 86 degrees F
	Routing Engine	OK	32 degrees C / 89 degrees F
	Heater off		
Pic	PIC 0/0 Channel 0	OK	28 degrees C / 82 degrees F
	PIC 0/0 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 0/0 Channel 2	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 9	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 0	OK	33 degrees C / 91 degrees F
	PIC 1/0 Channel 1	OK	31 degrees C / 87 degrees F
	PIC 1/0 Channel 2	OK	30 degrees C / 86 degrees F
	PIC 1/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 0	OK	31 degrees C / 87 degrees F
	PIC 1/1 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 1/1 Channel 2	OK	28 degrees C / 82 degrees F
	PIC 1/1 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 8	OK	0 degrees C / 32 degrees F
Fans	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed

show chassis environment cb

List of Syntax	Syntax on page 561 Syntax (TX Matrix Routers) on page 561 Syntax (TX Matrix Plus Routers) on page 561 Syntax (MX Series Routers) on page 561 Syntax (MX2010 3D Universal Edge Routers) on page 561 Syntax (MX2020 3D Universal Edge Routers) on page 561 Syntax (QFabric System) on page 561
Syntax	show chassis environment cb <slot>
Syntax (TX Matrix Routers)	show chassis environment cb <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment cb <lcc number sfc number > <slot>
Syntax (MX Series Routers)	show chassis environment cb <slot> <all-members> <local> <member member-id>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment cb <slot>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment cb <slot>
Syntax (QFabric System)	show chassis environment cb <slot interconnect-device interconnect-device-name> < interconnect-device interconnect-device-name slot>
Release Information	<p>Command introduced before Junos Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	(M120, M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Switches only) Display environmental information about the Control Boards (CBs). For information about the meaning of "CBs" on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> .

Options **none**—Display environmental information about all CBs. For a TX Matrix router, display environmental information about all CBs on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all CBs on the TX Matrix Plus router and its attached T1600 or T4000 routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the CBs on all the members of the Virtual Chassis configuration.

interconnect-device—(QFabric systems only) Display environmental information about the CBs on the Interconnect device.

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information about the CBs on the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display environmental information about the CBs on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display environmental information about the CBs in the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus router only) (Optional) Display environmental information about the CBs in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about the specified CB. On routers and PTX Series Packet Transport Switches, replace *slot* with **0** or **1**. On EX Series switches replace *slot* with **0**, **1**, or **2**. On QFX Series switches, replace *slot* with **0** or **1**.

Required Privilege Level

view

Related Documentation

- [request chassis cb on page 360](#)
- *Switching Control Board Redundancy*
- *Routing Engine and Switching Control Board Redundancy Configuration Statements*

List of Sample Output

- [show chassis environment cb \(M120 Router\) on page 564](#)
- [show chassis environment cb \(M320 Router\) on page 564](#)
- [show chassis environment cb \(MX80 Router\) on page 565](#)
- [show chassis environment cb \(MX240 Router\) on page 565](#)
- [show chassis environment cb \(MX240 Router with Enhanced MX SCB\) on page 566](#)
- [show chassis environment cb \(MX480 Router\) on page 566](#)
- [show chassis environment cb \(MX480 Router with Enhanced MX SCB\) on page 567](#)
- [show chassis environment cb \(MX960 Router\) on page 567](#)
- [show chassis environment cb \(MX960 Router with Enhanced MX SCB\) on page 567](#)
- [show chassis environment cb \(MX2020 Router\) on page 568](#)
- [show chassis environment cb \(MX2010 Router\) on page 569](#)
- [show chassis environment cb \(T4000 Core Router\) on page 569](#)
- [show chassis environment cb \(TX Matrix Router\) on page 570](#)
- [show chassis environment cb \(TX Matrix Plus Router\) on page 571](#)
- [show chassis environment cb \(EX8200 Switch\) on page 574](#)
- [show chassis environment cb \(EX8208 Switch\) on page 575](#)
- [show chassis environment cb \(PTX5000 Packet Transport Switch\) on page 577](#)
- [show chassis environment cb \(QFabric System\) on page 577](#)

Output Fields Table 55 on page 563 lists the output fields for the **show chassis environment cb** command. Output fields are listed in the approximate order in which they appear.

Table 55: show chassis environment cb Output Fields

Field Name	Field Description
State	<p>Status of the CB. If two CBs are installed and online, one is functioning as the master, and the other is the standby.</p> <ul style="list-style-type: none"> • Online—CB is online and running. • Offline—CB is powered down. <p>NOTE: On the EX8208 switch, the installation can include three CBs. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>.</p>
Temperature	<p>Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB.</p> <ul style="list-style-type: none"> • Temperature Intake—Measures the temperature of the air intake to cool the power supplies. • Temperature Exhaust—Measures the temperature of the hot air exhaust. <p>NOTE: On the MX2010 and MX2020 routers, the intake temperature measures the temperature of the air intake to cool the Control Board (CB). The MX2010 and MX2020 routers include intake and exhaust temperatures for multiple zones (Intake A, Intake B, Intake C, Exhaust A, Exhaust B, and TCBC).</p>
Power	<p>Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.</p>
BUS Revision	<p>Revision level of the generic bus device. (Not on switches.)</p>
FPGA Revision	<p>Revision level of the field-programmable gate array (FPGA). (Not on switches.)</p>

Table 55: show chassis environment cb Output Fields (*continued*)

Field Name	Field Description
PMBus device (on MX240, MX480, and MX960 routers with Enhanced MX SCB)	Enhanced SCB on MX 240, MX480, and MX960 routers allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for routers with Enhanced MX SCB: <ul style="list-style-type: none"> • Expected voltage • Measured voltage • Measured current • Calculated power

Sample Output

show chassis environment cb (M120 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                Online Master
  Temperature          33 degrees C / 91 degrees F
  Power
    1.2 V              1214 mV
    1.5 V              1495 mV
    2.5 V              2494 mV
    3.3 V              3319 mV
    5.0 V              5085 mV
    3.3 V bias         3296 mV
  Bus Revision         12
  FPGA Revision        17
CB 1 status:
  State                Online Standby
  Temperature          34 degrees C / 93 degrees F
  Power
    1.2 V              1195 mV
    1.5 V              1495 mV
    2.5 V              2504 mV
    3.3 V              3312 mV
    5.0 V              5111 mV
    3.3 V bias         3296 mV
  Bus Revision         12
  FPGA Revision        17

```

show chassis environment cb (M320 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                Online Master
  Temperature          29 degrees C / 84 degrees F
  Power:
    1.8 V              1805 mV
    2.5 V              2501 mV
    3.3 V              3293 mV
    4.6 V              4725 mV
    5.0 V              5032 mV
    12.0 V             11975 mV
    3.3 V bias         3286 mV

```



```

      8.0 V bias          7589 mV
    BUS Revision          40
    FPGA Revision         7
  CB 1 status:
    State                  Online Standby
    Temperature            32 degrees C / 89 degrees F
    Power:
      1.8 V                1802 mV
      2.5 V                2482 mV
      3.3 V                3289 mV
      4.6 V                4720 mV
      5.0 V                5001 mV
      12.0 V              11946 mV
      3.3 V bias          3274 mV
      8.0 V bias          7562 mV
    BUS Revision          40
    FPGA Revision         7

```

show chassis environment cb (MX80 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                  Online Master
  Temperature            36 degrees C / 96 degrees F
  Power 1
    1.0 V                1034 mV
    1.0 V MQ             1037 mV
    1.0 V LU             1005 mV
    1.2 V                1218 mV
    1.5 V                1524 mV
    1.8 V                1814 mV
    2.5 V                2558 mV
    3.3 V                3296 mV
    5.0 V                5233 mV
    5.0 V bias           5207 mV
    12.0 V              12162 mV

```

show chassis environment cb (MX240 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                  Online Standby
  Temperature            37 degrees C / 98 degrees F
  Power 1
    1.2 V                1208 mV
    1.5 V                1521 mV
    1.8 V                1811 mV
    2.5 V                2513 mV
    3.3 V                3332 mV
    5.0 V                5059 mV
    12.0 V              12162 mV
    1.25 V               1260 mV
    3.3 V SM3            3306 mV
    5.0 V RE             5085 mV
    12.0 V RE            11872 mV
  Power 2
    11.3 V bias PEM      11272 mV
    4.6 V bias MidPlane  4827 mV
    11.3 V bias FPD      11272 mV
    11.3 V bias POE 0    11292 mV
    11.3 V bias POE 1    11253 mV

```

```
Bus Revision          42
FPGA Revision         1
```

show chassis environment cb (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
CB 0 status:
State                Online Standby
Temperature           37 degrees C / 98 degrees F
Power 1
  1.2 V              1208 mV
  1.5 V              1521 mV
  1.8 V              1811 mV
  2.5 V              2513 mV
  3.3 V              3332 mV
  5.0 V              5059 mV
  12.0 V             12162 mV
  1.25 V             1260 mV
  3.3 V SM3          3306 mV
  5.0 V RE           5085 mV
  12.0 V RE          11872 mV
Power 2
  11.3 V bias PEM    11272 mV
  4.6 V bias MidPlane 4827 mV
  11.3 V bias FPD    11272 mV
  11.3 V bias POE 0   11292 mV
  11.3 V bias POE 1   11253 mV
Bus Revision         42
FPGA Revision        1
PMBus
device              Expected voltage    Measured voltage    Measured current    Calculated power
XF ASIC A           1000 mV      997 mV      11031 mA     10997 mW
XF ASIC B           1000 mV      996 mV      12125 mA     12076 mW
```

show chassis environment cb (MX480 Router)

```
user@host> show chassis environment cb
CB 0 status:
State                Online Master
Temperature           41 degrees C / 105 degrees F
Power 1
  1.2 V              1202 mV
  1.5 V              1511 mV
  1.8 V              1798 mV
  2.5 V              2507 mV
  3.3 V              3312 mV
  5.0 V              5027 mV
  12.0 V             12200 mV
  1.25 V             1260 mV
  3.3 V SM3          3293 mV
  5 V RE             5040 mV
  12 V RE            11910 mV
Power 2
  11.3 V bias PEM    11156 mV
  4.6 V bias MidPlane 4801 mV
  11.3 V bias FPD    11214 mV
  11.3 V bias POE 0   11098 mV
  11.3 V bias POE 1   11330 mV
Bus Revision         42
FPGA Revision        1
```

show chassis environment cb (MX480 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         41 degrees C / 105 degrees F
Power 1
  1.2 V                             1202 mV
  1.5 V                             1511 mV
  1.8 V                             1798 mV
  2.5 V                             2507 mV
  3.3 V                             3312 mV
  5.0 V                             5027 mV
  12.0 V                             12200 mV
  1.25 V                             1260 mV
  3.3 V SM3                          3293 mV
  5 V RE                             5040 mV
  12 V RE                             11910 mV
Power 2
  11.3 V bias PEM                    11156 mV
  4.6 V bias MidPlane                4801 mV
  11.3 V bias FPD                    11214 mV
  11.3 V bias POE 0                  11098 mV
  11.3 V bias POE 1                  11330 mV
Bus Revision                         42
FPGA Revision                        1
PMBus                               Expected Measured Measured Calculated
device                               voltage voltage current power
  XF ASIC A                         1000 mV   997 mV  11031 mA  10997 mW
  XF ASIC B                         1000 mV   996 mV  12125 mA  12076 mW

```

show chassis environment cb (MX960 Router)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         24 degrees C / 75 degrees F
Power 1
  1.2 V                             1965 mV
  1.5 V                             2465 mV
  1.8 V                             2990 mV
  2.5 V                             3296 mV
  3.3 V                             3296 mV
  5.0 V                             6593 mV
  12.0 V                             13187 mV
  3.3 V bias                         3296 mV
  1.25 V                             1994 mV
  3.3 V SM3                          3296 mV
  5 V RE                             6593 mV
  12 V RE                             13174 mV
Power 2                             Sensor failure
Bus Revision                         4
FPGA Revision                        3

```

show chassis environment cb (MX960 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         24 degrees C / 75 degrees F
Power 1

```

1.2 V	1965 mV			
1.5 V	2465 mV			
1.8 V	2990 mV			
2.5 V	3296 mV			
3.3 V	3296 mV			
5.0 V	6593 mV			
12.0 V	13187 mV			
3.3 V bias	3296 mV			
1.25 V	1994 mV			
3.3 V SM3	3296 mV			
5 V RE	6593 mV			
12 V RE	13174 mV			
Power 2	Sensor failure			
Bus Revision	4			
FPGA Revision	3			
PMBus	Expected	Measured	Measured	Calculated
device	voltage	voltage	current	power
XF ASIC A	1000 mV	997 mV	11031 mA	10997 mW
XF ASIC B	1000 mV	996 mV	12125 mA	12076 mW

show chassis environment cb (MX2020 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State Online Master
  IntakeA-Zone0 Temperature 44 degrees C / 111 degrees F
  IntakeB-Zone1 Temperature 34 degrees C / 93 degrees F
  IntakeC-Zone0 Temperature 45 degrees C / 113 degrees F
  ExhaustA-Zone0 Temperature 43 degrees C / 109 degrees F
  ExhaustB-Zone1 Temperature 36 degrees C / 96 degrees F
  TCBC-Zone0 Temperature 39 degrees C / 102 degrees F
  Power 1
    1.0 V 1011 mV
    1.2 V 1208 mV
    1.8 V 1801 mV
    2.5 V 2552 mV
    3.3 V 3312 mV
    5.0 V 5040 mV
    5.0 V RE 4988 mV
    12.0 V 12065 mV
    12.0 V RE 12046 mV
  Bus Revision 99
  FPGA Revision 270
CB 1 status:
  State Online Standby
  IntakeA-Zone0 Temperature 45 degrees C / 113 degrees F
  IntakeB-Zone1 Temperature 41 degrees C / 105 degrees F
  IntakeC-Zone0 Temperature 46 degrees C / 114 degrees F
  ExhaustA-Zone0 Temperature 44 degrees C / 111 degrees F
  ExhaustB-Zone1 Temperature 41 degrees C / 105 degrees F
  TCBC-Zone0 Temperature 45 degrees C / 113 degrees F
  Power 1
    1.0 V 1008 mV
    1.2 V 1208 mV
    1.8 V 1798 mV
    2.5 V 2539 mV
    3.3 V 3325 mV
    5.0 V 5033 mV
    5.0 V RE 4950 mV
    12.0 V 12046 mV
    12.0 V RE 11968 mV

```

```

Bus Revision          99
FPGA Revision         0

```

show chassis environment cb (MX2010 Router)

```

user@host> show chassis environment cb
CB 0 status:
State                Online Master
IntakeA-Zone0 Temperature 36 degrees C / 96 degrees F
IntakeB-Zone1 Temperature 30 degrees C / 86 degrees F
IntakeC-Zone0 Temperature 38 degrees C / 100 degrees F
ExhaustA-Zone0 Temperature 36 degrees C / 96 degrees F
ExhaustB-Zone1 Temperature 32 degrees C / 89 degrees F
TCBC-Zone0 Temperature 34 degrees C / 93 degrees F
Power 1
  1.0 V              1015 mV
  1.2 V              1205 mV
  1.8 V              1804 mV
  2.5 V              2552 mV
  3.3 V              3325 mV
  5.0 V              5020 mV
  5.0 V RE           4988 mV
  12.0 V             12104 mV
  12.0 V RE          12026 mV
Bus Revision         100
FPGA Revision        270
CB 1 status:
State                Online
IntakeA-Zone0 Temperature 35 degrees C / 95 degrees F
IntakeB-Zone1 Temperature 28 degrees C / 82 degrees F
IntakeC-Zone0 Temperature 37 degrees C / 98 degrees F
ExhaustA-Zone0 Temperature 34 degrees C / 93 degrees F
ExhaustB-Zone1 Temperature 29 degrees C / 84 degrees F
TCBC-Zone0 Temperature 32 degrees C / 89 degrees F
Power 1
  1.0 V              1011 mV
  1.2 V              1208 mV
  1.8 V              1788 mV
  2.5 V              2526 mV
  3.3 V              3319 mV
  5.0 V              5046 mV
  5.0 V RE           4975 mV
  12.0 V             12046 mV
  12.0 V RE          12007 mV
Bus Revision         100
FPGA Revision        0

```

show chassis environment cb (T4000 Core Router)

```

user@host> show chassis environment cb
CB 0 status:
State                Online Master
Temperature          33 degrees C / 91 degrees F
Power 1
  1.8 V              1805 mV
  2.5 V              2523 mV
  3.3 V              3324 mV
  3.3 V bias         3296 mV
  4.6 V              4680 mV
  5.0 V              4893 mV
  8.0 V bias         7572 mV

```

```
12.0 V          11916 mV
Power 2
1.0 V           993 mV
1.2 V           1210 mV
3.3 V RE        3330 mV
Bus Revision    51
FPGA Revision   5
CB 1 status:
State           Online Standby
Temperature     33 degrees C / 91 degrees F
Power 1
1.8 V           1810 mV
2.5 V           2496 mV
3.3 V           3308 mV
3.3 V bias      3286 mV
4.6 V           4692 mV
5.0 V           4954 mV
8.0 V bias      7282 mV
12.0 V          11926 mV
Power 2
1.0 V           993 mV
1.2 V           1185 mV
3.3 V RE        3316 mV
Bus Revision    51
FPGA Revision   5
```

show chassis environment cb (TX Matrix Router)

```
user@host> show chassis environment cb
```

```
-----
CB 0 status:
State           Online Master
Temperature     32 degrees C / 89 degrees F
Power:
1.8 V           1797 mV
2.5 V           2477 mV
3.3 V           3311 mV
4.6 V           4727 mV
5.0 V           5015 mV
12.0 V          12185 mV
3.3 V bias      3304 mV
8.0 V bias      7870 mV
BUS Revision    40
FPGA Revision   1
CB 1 status:
State           Online Standby
...
```

```
1cc0-re0:
```

```
-----
CB 0 status:
State           Online Master
Temperature     32 degrees C / 89 degrees F
Power:
1.8 V           1787 mV
2.5 V           2473 mV
3.3 V           3306 mV
4.6 V           4793 mV
5.0 V           5025 mV
12.0 V          12156 mV
3.3 V bias      3289 mV
```

```

      8.0 V bias          7609 mV
    BUS Revision          40
    FPGA Revision         5
  CB 1 status:
    State                  Online Standby
    ....
    BUS Revision          40
    FPGA Revision         5

```

```
1cc2-re0:
```

```

-----
  CB 0 status:
    State                  Online Master
    ...
  CB 1 status:
    State                  Online Standby
    ...

```

show chassis environment cb (TX Matrix Plus Router)

```
user@host> show chassis environment cb
```

```
sfc0-re0:
```

```

-----
  CB 0 status:
    State                  Online Master
    Temperature            38 degrees C / 100 degrees F
    Power 1
      1.0 V                1005 mV
      1.1 V                1108 mV
      1.2 V                1205 mV
      1.25 V               1269 mV
      1.5 V                1508 mV
      1.8 V                1814 mV
      2.5 V                2507 mV
      3.3 V                3306 mV
      3.3 V bias           3300 mV
      9.0 V                9058 mV
      9.0 V RE             9107 mV
    Power 2
      3.9 V                3963 mV
      5.0 V                5020 mV
      9.0 V                9087 mV
    Bus Revision           79
    FPGA Revision          23
  CB 1 status:
    State                  Online Standby
    Temperature            39 degrees C / 102 degrees F
    Power 1
      1.0 V                1002 mV
      1.1 V                1105 mV
      1.2 V                1198 mV
      1.25 V               1276 mV
      1.5 V                1504 mV
      1.8 V                1804 mV
      2.5 V                2507 mV
      3.3 V                3300 mV
      3.3 V bias           3293 mV
      9.0 V                9039 mV
      9.0 V RE             9049 mV
    Power 2
      3.9 V                3892 mV

```

5.0 V	5040 mV
9.0 V	9058 mV
Bus Revision	79
FPGA Revision	23

lcc0-re0:

CB 0 status:

State	Online Master
Temperature	39 degrees C / 102 degrees F
Power 1	
1.8 V	1799 mV
2.5 V	2499 mV
3.3 V	3327 mV
3.3 V bias	3299 mV
4.6 V	4673 mV
5.0 V	4918 mV
8.0 V bias	7308 mV
12.0 V	11887 mV
Power 2	
1.0 V	996 mV
1.2 V	1199 mV
3.3 V RE	3319 mV
Bus Revision	51
FPGA Revision	3

CB 1 status:

State	Online Standby
Temperature	40 degrees C / 104 degrees F
Power 1	
1.8 V	1800 mV
2.5 V	2496 mV
3.3 V	3322 mV
3.3 V bias	3284 mV
4.6 V	4680 mV
5.0 V	4954 mV
8.0 V bias	7284 mV
12.0 V	11902 mV
Power 2	
1.0 V	998 mV
1.2 V	1205 mV
3.3 V RE	3327 mV
Bus Revision	51
FPGA Revision	3

lcc1-re0:

CB 0 status:

State	Online Master
Temperature	41 degrees C / 105 degrees F
Power 1	
1.8 V	1804 mV
2.5 V	2517 mV
3.3 V	3300 mV
3.3 V bias	3284 mV
4.6 V	4681 mV
5.0 V	4927 mV
8.0 V bias	7357 mV
12.0 V	11907 mV
Power 2	
1.0 V	991 mV
1.2 V	1202 mV


```

    3.3 V RE          3301 mV
    Bus Revision      51
    FPGA Revision     3
    CB 1 status:
    State              Online Standby
    Temperature        40 degrees C / 104 degrees F
    Power 1
    1.8 V              1805 mV
    2.5 V              2528 mV
    3.3 V              3324 mV
    3.3 V bias         3289 mV
    4.6 V              4694 mV
    5.0 V              4959 mV
    8.0 V bias         7311 mV
    12.0 V             11926 mV
    Power 2
    1.0 V              998 mV
    1.2 V              1200 mV
    3.3 V RE          3313 mV
    Bus Revision      51
    FPGA Revision     3

```

```
lcc2-re0:
```

```

-----
    CB 0 status:
    State              Online Master
    Temperature        41 degrees C / 105 degrees F
    Power 1
    1.8 V              1805 mV
    2.5 V              2494 mV
    3.3 V              3333 mV
    3.3 V bias         3296 mV
    4.6 V              4673 mV
    5.0 V              4901 mV
    8.0 V bias         7343 mV
    12.0 V             11916 mV
    Power 2
    1.0 V              993 mV
    1.2 V              1213 mV
    3.3 V RE          3328 mV
    Bus Revision      51
    FPGA Revision     3
    CB 1 status:
    State              Online Standby
    Temperature        41 degrees C / 105 degrees F
    Power 1
    1.8 V              1804 mV
    2.5 V              2523 mV
    3.3 V              3334 mV
    3.3 V bias         3291 mV
    4.6 V              4697 mV
    5.0 V              4969 mV
    8.0 V bias         7308 mV
    12.0 V             11936 mV
    Power 2
    1.0 V              996 mV
    1.2 V              1200 mV
    3.3 V RE          3328 mV
    Bus Revision      51
    FPGA Revision     3

```

lcc3-re0:

CB 0 status:

State	Online Master
Temperature	37 degrees C / 98 degrees F
Power 1	
1.8 V	1809 mV
2.5 V	2510 mV
3.3 V	3296 mV
3.3 V bias	3291 mV
4.6 V	4670 mV
5.0 V	4905 mV
8.0 V bias	7211 mV
12.0 V	11882 mV
Power 2	
1.0 V	996 mV
1.2 V	1188 mV
3.3 V RE	3326 mV
Bus Revision	51
FPGA Revision	5

CB 1 status:

State	Online Standby
Temperature	38 degrees C / 100 degrees F
Power 1	
1.8 V	1813 mV
2.5 V	2510 mV
3.3 V	3322 mV
3.3 V bias	3289 mV
4.6 V	4692 mV
5.0 V	4967 mV
8.0 V bias	7194 mV
12.0 V	11916 mV
Power 2	
1.0 V	996 mV
1.2 V	1205 mV
3.3 V RE	3273 mV
Bus Revision	51
FPGA Revision	5

show chassis environment cb (EX8200 Switch)

user@host> show chassis environment cb

CB 0 status:

State	Online Master
Temperature Intake	20 degrees C / 68 degrees F
Temperature Exhaust	24 degrees C / 75 degrees F
Power 1	
1.1 V	1086 mV
1.2 V	1179 mV
1.2 V *	1182 mV
1.2 V *	1182 mV
1.25 V	1211 mV
1.5 V	1472 mV
1.8 V	1756 mV
2.5 V	2449 mV
3.3 V	3254 mV
3.3 V bias	3300 mV
5.0 V	4911 mV
12.0 V	11891 mV
Power 2	

```

3.3 V bias *          3615 mV
3.3 V bias *          3615 mV
3.3 V bias *          3567 mV
3.3 V bias *          3664 mV
4.3 V bias *          4224 mV
4.3 V bias *          4215 mV
4.3 V bias *          4224 mV
4.3 V bias *          4205 mV
4.3 V bias *          4195 mV
4.3 V bias *          4215 mV
5.0 V bias            4920 mV
CB 1 status:
State                  Online Standby
Temperature Intake     19 degrees C / 66 degrees F
Temperature Exhaust    23 degrees C / 73 degrees F
Power 1
1.1 V                  1082 mV
1.2 V                  1169 mV
1.2 V *                1179 mV
1.2 V *                1179 mV
1.25 V                1214 mV
1.5 V                  1482 mV
1.8 V                  1759 mV
2.5 V                  2481 mV
3.3 V                  3248 mV
3.3 V bias             3306 mV
5.0 V                  4911 mV
12.0 V                 11910 mV
Power 2
3.3 V bias *           3644 mV
3.3 V bias *           3664 mV
3.3 V bias *           3586 mV
3.3 V bias *           3654 mV
4.3 V bias *           4224 mV
4.3 V bias *           4215 mV
4.3 V bias *           4224 mV
4.3 V bias *           4205 mV
4.3 V bias *           4244 mV
4.3 V bias *           4215 mV
5.0 V bias             4930 mV
CB 2 status:
State                  Online
Temperature Intake     19 degrees C / 66 degrees F
Temperature Exhaust    23 degrees C / 73 degrees F
Power 1
1.2 V                  1195 mV
1.5 V                  1511 mV
1.8 V                  1804 mV
2.5 V                  2526 mV
3.3 V                  3300 mV
3.3 V bias             3306 mV
12.0 V                 12220 mV

```

show chassis environment cb (EX8208 Switch)

```

user@host> show chassis environment cb
CB 0 status:
State                  Online Master
Temperature Intake     20 degrees C / 68 degrees F
Temperature Exhaust    24 degrees C / 75 degrees F
Power 1

```

1.1 V	1086 mV
1.2 V	1179 mV
1.2 V *	1182 mV
1.2 V *	1182 mV
1.25 V	1211 mV
1.5 V	1466 mV
1.8 V	1759 mV
2.5 V	2455 mV
3.3 V	3261 mV
3.3 V bias	3300 mV
5.0 V	4930 mV
12.0 V	11891 mV
Power 2	
3.3 V bias *	3606 mV
3.3 V bias *	3615 mV
3.3 V bias *	3567 mV
3.3 V bias *	3673 mV
4.3 V bias *	4224 mV
4.3 V bias *	4215 mV
4.3 V bias *	4234 mV
4.3 V bias *	4205 mV
4.3 V bias *	4186 mV
4.3 V bias *	4215 mV
5.0 V bias	4940 mV
CB 1 status:	
State	Online Standby
Temperature Intake	19 degrees C / 66 degrees F
Temperature Exhaust	23 degrees C / 73 degrees F
Power 1	
1.1 V	1086 mV
1.2 V	1169 mV
1.2 V *	1179 mV
1.2 V *	1179 mV
1.25 V	1211 mV
1.5 V	1479 mV
1.8 V	1759 mV
2.5 V	2475 mV
3.3 V	3235 mV
3.3 V bias	3306 mV
5.0 V	4930 mV
12.0 V	11891 mV
Power 2	
3.3 V bias *	3644 mV
3.3 V bias *	3664 mV
3.3 V bias *	3586 mV
3.3 V bias *	3654 mV
4.3 V bias *	4215 mV
4.3 V bias *	4224 mV
4.3 V bias *	4215 mV
4.3 V bias *	4215 mV
4.3 V bias *	4234 mV
4.3 V bias *	4224 mV
5.0 V bias	4920 mV
CB 2 status:	
State	Online
Temperature Intake	20 degrees C / 68 degrees F
Temperature Exhaust	24 degrees C / 75 degrees F
Power 1	
1.2 V	1202 mV
1.5 V	1508 mV
1.8 V	1804 mV

2.5 V	2520 mV
3.3 V	3300 mV
3.3 V bias	3300 mV
12.0 V	12200 mV

show chassis environment cb (PTX5000 Packet Transport Switch)

```

user@host> show chassis environment cb
CB 0 status:
  State                               Online Master
  Intake Temperature                  38 degrees C / 100 degrees F
  Exhaust A Temperature               45 degrees C / 113 degrees F
  Exhaust B Temperature               42 degrees C / 107 degrees F
  Power 1
    1.2 V                             1200 mV
    1.25 V                            1250 mV
    2.5 V                             2500 mV
    3.3 V                             3300 mV
  Power 2
    1.0 V                             1000 mV
    3.3 V bias                        3293 mV
    3.9 V                             3921 mV
  Bus Revision                        132
  FPGA Revision                       27
CB 1 status:
  State                               Online Standby
  Intake Temperature                  34 degrees C / 93 degrees F
  Exhaust A Temperature               39 degrees C / 102 degrees F
  Exhaust B Temperature               36 degrees C / 96 degrees F
  Power 1
    1.2 V                             1199 mV
    1.25 V                            1250 mV
    2.5 V                             2499 mV
    3.3 V                             3299 mV
  Power 2
    1.0 V                             1000 mV
    3.3 V bias                        3312 mV
    3.9 V                             3961 mV
  Bus Revision                        132
  FPGA Revision                       28

```

show chassis environment cb (QFabric System)

```

user@switch> show chassis environment cb interconnect-device IC-123 0
CB 0 status:
  State                               Online Master
  Left Intake Temperature              33 degrees C / 91 degrees F
  Right Intake Temperature             33 degrees C / 91 degrees F
  Left Exhaust Temperature             36 degrees C / 96 degrees F
  Right Exhaust Temperature            35 degrees C / 95 degrees F
  Power                               OK
    VDD 3V3                           3294 mV
    VDD 2V5                           2436 mV
    VDD 1V8                           1746 mV
    VDD 1V5                           1460 mV
    VDD 1V25                          1210 mV
    VDD 1V2                           1164 mV
    CPU CORE 1V2                      1120 mV
    VDD 1V0                           968 mV
    VDD 5V0                           5088 mV
    CPU MP BIAS 4V3                   4050 mV

```

BIAS 3V3	3180 mV
VTT 0V9	866 mV

show chassis environment fpc

List of Syntax	Syntax on page 579 Syntax (TX Matrix and TX Matrix Plus Routers) on page 579 Syntax (MX Series Routers) on page 579 Syntax (MX2010 3D Universal Edge Routers) on page 579 Syntax (MX2020 3D Universal Edge Routers) on page 579 Syntax (QFX Series) on page 579
Syntax	show chassis environment fpc <slot>
Syntax (TX Matrix and TX Matrix Plus Routers)	show chassis environment fpc <lcc number> <slot>
Syntax (MX Series Routers)	show chassis environment fpc <slot> <all-members> <local> <member member-id>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment fpc <slot>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment fpc <slot>
Syntax (QFX Series)	show chassis environment fpc <fpc-slot> interconnect-device <i>name</i>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series switches only) Display environmental information about Flexible PIC Concentrators (FPCs).
Options	none —Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

slot or *fpc-slot*—(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Similarly, on a TX Matrix Plus router, if you specify the number of the router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
user@host> show chassis environment fpc 9
```

- M120 router—Replace **slot** with a value from 0 through 5.
- MX240 router—Replace **slot** with a value from 0 through 2.
- MX480 router—Replace **slot** with a value from 0 through 5.
- MX960 router—Replace **slot** with a value from 0 through 11.
- MX2010 router—Replace **slot** with a value from 0 through 9.
- MX2020 router—Replace **slot** with a value from 0 through 19.
- Other routers—Replace **slot** with a value from 0 through 7.
- EX Series switches:

- EX3200 switches and EX4200 standalone switches—Replace **slot** with 0.
- EX4200 switches in a Virtual Chassis configuration—Replace **slot** with a value from 0 through 9 (switch's member ID).
- EX6210 switches—Replace **slot** with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotating Engine (SRE) module), or 6 through 9 (line card only).
- EX8208 switches—Replace **slot** with a value from 0 through 7 (line card).
- EX8216 switches—Replace **slot** with a value from 0 through 15 (line card).
- QFX3500 switches —Replace **fpc-slot** with 0 through 15.
- PTX5000 Packet Transport Switch—Replace **fpc-slot** with 0 through 7.

Required Privilege Level view

- Related Documentation**
- [request chassis fpc on page 363](#)
 - [show chassis fpc on page 641](#)
 - *show chassis fpc-feb-connectivity*
 - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
 - *MX960 Flexible PIC Concentrator Description*

- List of Sample Output**
- [show chassis environment fpc \(M120 Router\) on page 583](#)
 - [show chassis environment fpc \(M160 Router\) on page 584](#)
 - [show chassis environment fpc \(M320 Router\) on page 584](#)
 - [show chassis environment fpc \(MX2020 Router\) on page 585](#)
 - [show chassis environment fpc \(MX2010 Router\) on page 588](#)
 - [show chassis environment fpc \(MX240 Router\) on page 590](#)
 - [show chassis environment fpc \(MX480 Router\) on page 591](#)
 - [show chassis environment fpc \(MX960 Router\) on page 592](#)
 - [show chassis environment fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 593](#)
 - [show chassis environment fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 594](#)
 - [show chassis environment fpc \(T320, T640, and T1600 Routers\) on page 594](#)
 - [show chassis environment fpc \(T4000 Router\) on page 595](#)
 - [show chassis environment fpc lcc \(TX Matrix Router\) on page 600](#)
 - [show chassis environment fpc lcc \(TX Matrix Plus Router\) on page 601](#)
 - [show chassis environment fpc \(QFX Series\) on page 601](#)
 - [show chassis environment fpc interconnect-device \(QFabric Systems\) on page 602](#)
 - [show chassis environment fpc 0 \(PTX5000 Packet Transport Switch\) on page 602](#)
 - [show chassis environment FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 603](#)

Output Fields Table 56 on page 582 lists the output fields for the **show chassis environment fpc** command. Output fields are listed in the approximate order in which they appear.

Table 56: show chassis environment fpc Output Fields

Field Name	Field Description
State	<p>Status of the FPC:</p> <ul style="list-style-type: none"> • Unknown—FPC is not detected by the router. • Empty—No FPC is present. • Present—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online. • Ready—FPC is in intermediate or transition state. • Announce online—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative. • Online—FPC is online and running. • Offline—FPC is powered down. • Diagnostics—FPC is set to operate in diagnostics mode.
Temperature	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
PMB Temperature	(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC).
Temperature Intake	(M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing into the chassis.
Temperature Top	(T Series routers only) Temperature of the air flowing past the top of the FPC.
Temperature Exhaust	<p>(M120 and M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing out of the chassis.</p> <p>The PTX Series Packet Transport Switches, and the MX2010 and MX2020 routers include exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</p>
Temperature Bottom	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
TL <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
TQ <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
Temperature MMBO	(T640 router only) Temperature of the air flowing past the type 3 FPC.
Temperature MMB1	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
Power	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
CMB Revision or BUS revision	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

Sample Output

show chassis environment fpc (M120 Router)

```

user@host> show chassis environment fpc
FPC 2 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               31 degrees C / 87 degrees F
  Power A-Board
    1.2 V                             1202 mV
    1.5 V                             1508 mV
    1.8 V                             1798 mV
    2.5 V                             2507 mV
    3.3 V                             3351 mV
    5.0 V                             4995 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 3 status:
  State                               Online
  Temperature Exhaust A               31 degrees C / 87 degrees F
  Temperature Exhaust B               33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                             1211 mV
    1.5 V                             1501 mV
    1.8 V                             1798 mV
    2.5 V                             2471 mV
    3.3 V                             3293 mV
    5.0 V                             4930 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  Power B-Board
    1.2 V                             1214 mV
    1.5 V                             1501 mV
    2.5 V                             2471 mV
    3.3 V                             3300 mV
    5.0 V                             4943 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 4 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               30 degrees C / 86 degrees F
  Power A-Board
    1.2 V                             1195 mV
    1.5 V                             1504 mV
    1.8 V                             1801 mV
    2.5 V                             2504 mV
    3.3 V                             3293 mV
    5.0 V                             4917 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1202 mV
    1.5 V Rocket IO                   1492 mV
  I2C Slave Revision                 12

```

show chassis environment fpc (M160 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature                         42 degrees C / 107 degrees F
  Power:
    1.5 V                             1500 mV
    2.5 V                             2509 mV
    3.3 V                             3308 mV
    5.0 V                             4991 mV
    5.0 V bias                         4952 mV
    8.0 V bias                         8307 mV
  CMB Revision                        12
FPC 1 status:
  State                               Online
  Temperature                         45 degrees C / 113 degrees F
  Power:
    1.5 V                             1498 mV
    2.5 V                             2501 mV
    3.3 V                             3319 mV
    5.0 V                             5020 mV
    5.0 V bias                         5025 mV
    8.0 V bias                         8307 mV
  CMB Revision                        12
```

show chassis environment fpc (M320 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature Intake                  27 degrees C / 80 degrees F
  Temperature Exhaust                 38 degrees C / 100 degrees F
  Temperature MMB1                    31 degrees C / 87 degrees F
  Power:
    1.5 V                             1487 mV
    1.5 V *                           1494 mV
    1.8 V                             1821 mV
    2.5 V                             2533 mV
    3.3 V                             3323 mV
    5.0 V                             5028 mV
    3.3 V bias                        3296 mV
    5.0 V bias                        4984 mV
  CMB Revision                        16
FPC 1 status:
  State                               Online
  Temperature Intake                  27 degrees C / 80 degrees F
  Temperature Exhaust                 37 degrees C / 98 degrees F
  Temperature MMB1                    32 degrees C / 89 degrees F
  Power:
    1.5 V                             1504 mV
    1.5 V *                           1499 mV
    1.8 V                             1820 mV
    2.5 V                             2529 mV
    3.3 V                             3328 mV
    5.0 V                             5013 mV
    3.3 V bias                        3294 mV
    5.0 V bias                        4984 mV
  CMB Revision                        16
FPC 2 status:
  State                               Online
```

```

Temperature Intake          28 degrees C / 82 degrees F
Temperature Exhaust         38 degrees C / 100 degrees F
Temperature MMB1            32 degrees C / 89 degrees F
Power:
  1.5 V                     1498 mV
  1.5 V *                   1487 mV
  1.8 V                     1816 mV
  2.5 V                     2531 mV
  3.3 V                     3324 mV
  5.0 V                     5025 mV
  3.3 V bias                3277 mV
  5.0 V bias                5013 mV
CMB Revision                17
FPC 3 status:
...
```

show chassis environment fpc (MX2020 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  41 degrees C / 105 degrees F
Temperature Exhaust A               48 degrees C / 118 degrees F
Temperature Exhaust B               60 degrees C / 140 degrees F
Temperature LU 0 TSen               56 degrees C / 132 degrees F
Temperature LU 0 Chip               59 degrees C / 138 degrees F
Temperature LU 1 TSen               56 degrees C / 132 degrees F
Temperature LU 1 Chip               61 degrees C / 141 degrees F
Temperature LU 2 TSen               56 degrees C / 132 degrees F
Temperature LU 2 Chip               52 degrees C / 125 degrees F
Temperature LU 3 TSen               56 degrees C / 132 degrees F
Temperature LU 3 Chip               52 degrees C / 125 degrees F
Temperature MQ 0 TSen               49 degrees C / 120 degrees F
Temperature MQ 0 Chip               49 degrees C / 120 degrees F
Temperature MQ 1 TSen               49 degrees C / 120 degrees F
Temperature MQ 1 Chip               52 degrees C / 125 degrees F
Temperature MQ 2 TSen               49 degrees C / 120 degrees F
Temperature MQ 2 Chip               45 degrees C / 113 degrees F
Temperature MQ 3 TSen               49 degrees C / 120 degrees F
Temperature MQ 3 Chip               46 degrees C / 114 degrees F
Power
  AS-BIAS3V3-z12105                3299 mV
  AS-VDD1V8-z12006                 1807 mV
  AS-VDD2V5-z12006                 2512 mV
  AS-AVDD1V0-z12004                 997 mV
  AS-PCIE_1V0-z12004                 996 mV
  AS-VDD3V3-z12004                 3294 mV
  AS-VDD_1V5A-z12004                1501 mV
  AS-VDD_1V5B-z12004                1498 mV
  AS-LU0_1V0-z12004                 998 mV
  AS-LU1_1V0-z12004                1002 mV
  AS-MQ0_1V0-z12004                 999 mV
  AS-MQ1_1V0-z12004                 994 mV
  AS-LU2_1V0-z12004                1000 mV
  AS-LU3_1V0-z12004                 998 mV
  AS-MQ2_1V0-z12004                1002 mV
  AS-MQ3_1V0-z12004                 999 mV
  AS-PMB_1V1-z12006                1096 mV
I2C Slave Revision                68
FPC 1 status:
State                               Online
```

Temperature Intake	39 degrees C / 102 degrees F
Temperature Exhaust A	48 degrees C / 118 degrees F
Temperature Exhaust B	55 degrees C / 131 degrees F
Temperature LU 0 TSen	52 degrees C / 125 degrees F
Temperature LU 0 Chip	54 degrees C / 129 degrees F
Temperature LU 1 TSen	52 degrees C / 125 degrees F
Temperature LU 1 Chip	56 degrees C / 132 degrees F
Temperature LU 2 TSen	52 degrees C / 125 degrees F
Temperature LU 2 Chip	49 degrees C / 120 degrees F
Temperature LU 3 TSen	52 degrees C / 125 degrees F
Temperature LU 3 Chip	50 degrees C / 122 degrees F
Temperature MQ 0 TSen	48 degrees C / 118 degrees F
Temperature MQ 0 Chip	48 degrees C / 118 degrees F
Temperature MQ 1 TSen	48 degrees C / 118 degrees F
Temperature MQ 1 Chip	51 degrees C / 123 degrees F
Temperature MQ 2 TSen	48 degrees C / 118 degrees F
Temperature MQ 2 Chip	45 degrees C / 113 degrees F
Temperature MQ 3 TSen	48 degrees C / 118 degrees F
Temperature MQ 3 Chip	45 degrees C / 113 degrees F
Power	
AS-BIAS3V3-z12105	3291 mV
AS-VDD1V8-z12006	1786 mV
AS-VDD2V5-z12006	2496 mV
AS-AVDD1V0-z12004	1000 mV
AS-PCIE_1V0-z12004	1000 mV
AS-VDD3V3-z12004	3294 mV
AS-VDD_1V5A-z12004	1500 mV
AS-VDD_1V5B-z12004	1498 mV
AS-LU0_1V0-z12004	1003 mV
AS-LU1_1V0-z12004	1000 mV
AS-MQ0_1V0-z12004	1000 mV
AS-MQ1_1V0-z12004	995 mV
AS-LU2_1V0-z12004	1002 mV
AS-LU3_1V0-z12004	997 mV
AS-MQ2_1V0-z12004	1000 mV
AS-MQ3_1V0-z12004	998 mV
AS-PMB_1V1-z12006	1096 mV
I2C Slave Revision	68
FPC 2 status:	
State	Online
Temperature Intake	39 degrees C / 102 degrees F
Temperature Exhaust A	48 degrees C / 118 degrees F
Temperature Exhaust B	58 degrees C / 136 degrees F
Temperature LU 0 TSen	55 degrees C / 131 degrees F
Temperature LU 0 Chip	57 degrees C / 134 degrees F
Temperature LU 1 TSen	55 degrees C / 131 degrees F
Temperature LU 1 Chip	63 degrees C / 145 degrees F
Temperature LU 2 TSen	55 degrees C / 131 degrees F
Temperature LU 2 Chip	51 degrees C / 123 degrees F
Temperature LU 3 TSen	55 degrees C / 131 degrees F
Temperature LU 3 Chip	52 degrees C / 125 degrees F
Temperature MQ 0 TSen	48 degrees C / 118 degrees F
Temperature MQ 0 Chip	50 degrees C / 122 degrees F
Temperature MQ 1 TSen	48 degrees C / 118 degrees F
Temperature MQ 1 Chip	52 degrees C / 125 degrees F
Temperature MQ 2 TSen	48 degrees C / 118 degrees F
Temperature MQ 2 Chip	47 degrees C / 116 degrees F
Temperature MQ 3 TSen	48 degrees C / 118 degrees F
Temperature MQ 3 Chip	47 degrees C / 116 degrees F
Power	
AS-BIAS3V3-z12105	3299 mV

```

AS-VDD1V8-z12006      1805 mV
AS-VDD2V5-z12006      2510 mV
AS-AVDD1V0-z12004      999 mV
AS-PCIE_1V0-z12004      998 mV
AS-VDD3V3-z12004      3296 mV
AS-VDD_1V5A-z12004     1492 mV
AS-VDD_1V5B-z12004     1497 mV
AS-LU0_1V0-z12004      997 mV
AS-LU1_1V0-z12004     1000 mV
AS-MQ0_1V0-z12004      998 mV
AS-MQ1_1V0-z12004     1001 mV
AS-LU2_1V0-z12004      996 mV
AS-LU3_1V0-z12004      995 mV
AS-MQ2_1V0-z12004      998 mV
AS-MQ3_1V0-z12004      997 mV
AS-PMB_1V1-z12006     1100 mV
I2C Slave Revision      68
FPC 3 status:
State                   Online
Temperature Intake       41 degrees C / 105 degrees F
Temperature Exhaust A    48 degrees C / 118 degrees F
Temperature Exhaust B    58 degrees C / 136 degrees F
Temperature LU 0 TSen     56 degrees C / 132 degrees F
Temperature LU 0 Chip     59 degrees C / 138 degrees F
Temperature LU 1 TSen     56 degrees C / 132 degrees F
Temperature LU 1 Chip     61 degrees C / 141 degrees F
Temperature LU 2 TSen     56 degrees C / 132 degrees F
Temperature LU 2 Chip     51 degrees C / 123 degrees F
Temperature LU 3 TSen     56 degrees C / 132 degrees F
Temperature LU 3 Chip     53 degrees C / 127 degrees F
Temperature MQ 0 TSen     50 degrees C / 122 degrees F
Temperature MQ 0 Chip     51 degrees C / 123 degrees F
Temperature MQ 1 TSen     50 degrees C / 122 degrees F
Temperature MQ 1 Chip     55 degrees C / 131 degrees F
Temperature MQ 2 TSen     50 degrees C / 122 degrees F
Temperature MQ 2 Chip     47 degrees C / 116 degrees F
Temperature MQ 3 TSen     50 degrees C / 122 degrees F
Temperature MQ 3 Chip     50 degrees C / 122 degrees F
Power
AS-BIAS3V3-z12105       3305 mV
AS-VDD1V8-z12006       1810 mV
AS-VDD2V5-z12006       2508 mV
AS-AVDD1V0-z12004       999 mV
AS-PCIE_1V0-z12004      1001 mV
AS-VDD3V3-z12004       3294 mV
AS-VDD_1V5A-z12004     1500 mV
AS-VDD_1V5B-z12004     1498 mV
AS-LU0_1V0-z12004       998 mV
AS-LU1_1V0-z12004       998 mV
AS-MQ0_1V0-z12004       999 mV
AS-MQ1_1V0-z12004       998 mV
AS-LU2_1V0-z12004      1000 mV
AS-LU3_1V0-z12004      1001 mV
AS-MQ2_1V0-z12004       996 mV
AS-MQ3_1V0-z12004       998 mV
AS-PMB_1V1-z12006     1098 mV
I2C Slave Revision      68
FPC 4 status:
...
```

show chassis environment fpc (MX2010 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A               42 degrees C / 107 degrees F
Temperature Exhaust B               51 degrees C / 123 degrees F
Temperature LU 0 TSen               49 degrees C / 120 degrees F
Temperature LU 0 Chip               50 degrees C / 122 degrees F
Temperature LU 1 TSen               49 degrees C / 120 degrees F
Temperature LU 1 Chip               54 degrees C / 129 degrees F
Temperature LU 2 TSen               49 degrees C / 120 degrees F
Temperature LU 2 Chip               45 degrees C / 113 degrees F
Temperature LU 3 TSen               49 degrees C / 120 degrees F
Temperature LU 3 Chip               46 degrees C / 114 degrees F
Temperature MQ 0 TSen               40 degrees C / 104 degrees F
Temperature MQ 0 Chip               41 degrees C / 105 degrees F
Temperature MQ 1 TSen               40 degrees C / 104 degrees F
Temperature MQ 1 Chip               44 degrees C / 111 degrees F
Temperature MQ 2 TSen               40 degrees C / 104 degrees F
Temperature MQ 2 Chip               38 degrees C / 100 degrees F
Temperature MQ 3 TSen               40 degrees C / 104 degrees F
Temperature MQ 3 Chip               41 degrees C / 105 degrees F
Power
AS-BIAS3V3-z12105                  3300 mV
AS-VDD1V8-z12006                   1805 mV
AS-VDD2V5-z12006                   2505 mV
AS-AVDD1V0-z12004                   998 mV
AS-PCIE_1V0-z12004                   999 mV
AS-VDD3V3-z12004                   3303 mV
AS-VDD_1V5A-z12004                 1497 mV
AS-VDD_1V5B-z12004                 1497 mV
AS-LU0_1V0-z12004                   998 mV
AS-LU1_1V0-z12004                  1003 mV
AS-MQ0_1V0-z12004                   998 mV
AS-MQ1_1V0-z12004                   998 mV
AS-LU2_1V0-z12004                   997 mV
AS-LU3_1V0-z12004                  1001 mV
AS-MQ2_1V0-z12004                   996 mV
AS-MQ3_1V0-z12004                   994 mV
AS-PMB_1V1-z12006                  1097 mV
I2C Slave Revision                 68
FPC 1 status:
State                               Online
Temperature Intake                  34 degrees C / 93 degrees F
Temperature Exhaust A               46 degrees C / 114 degrees F
Temperature Exhaust B               54 degrees C / 129 degrees F
Temperature LU 0 TSen               45 degrees C / 113 degrees F
Temperature LU 0 Chip               55 degrees C / 131 degrees F
Temperature LU 1 TSen               45 degrees C / 113 degrees F
Temperature LU 1 Chip               44 degrees C / 111 degrees F
Temperature LU 2 TSen               45 degrees C / 113 degrees F
Temperature LU 2 Chip               50 degrees C / 122 degrees F
Temperature LU 3 TSen               45 degrees C / 113 degrees F
Temperature LU 3 Chip               58 degrees C / 136 degrees F
Temperature XM 0 TSen               45 degrees C / 113 degrees F
Temperature XM 0 Chip               51 degrees C / 123 degrees F
Temperature XF 0 TSen               45 degrees C / 113 degrees F
Temperature XF 0 Chip               63 degrees C / 145 degrees F
Temperature PLX Switch TSen         45 degrees C / 113 degrees F

```



```

Temperature PLX Switch Chip47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105      3300 mV
MPC-VDD3V3-z16100      3294 mV
MPC-VDD2V5-z16100      2505 mV
MPC-VDD1V8-z12004      1796 mV
MPC-AVDD1V0-z12004      991 mV
MPC-VDD1V2-z16100      1196 mV
MPC-VDD1V5A-z12004      1491 mV
MPC-VDD1V5B-z12004      1492 mV
MPC-XF_0V9-z12004      996 mV
MPC-PCIE_1V0-z16100     1003 mV
MPC-LU0_1V0-z12004      996 mV
MPC-LU1_1V0-z12004      996 mV
MPC-LU2_1V0-z12004      998 mV
MPC-LU3_1V0-z12004      994 mV
MPC-12VA-BMR453         12031 mV
MPC-12VB-BMR453         12003 mV
MPC-PMB_1V1-z12006      1104 mV
MPC-PMB_1V2-z12106      1194 mV
MPC-XM_0V9-vt273m       911 mV
I2C Slave Revision      110
FPC 8 status:
State                    Online
Temperature Intake        32 degrees C / 89 degrees F
Temperature Exhaust A     44 degrees C / 111 degrees F
Temperature Exhaust B     37 degrees C / 98 degrees F
Temperature LU 0 TCAM TSen 41 degrees C / 105 degrees F
Temperature LU 0 TCAM Chip 49 degrees C / 120 degrees F
Temperature LU 0 TSen      41 degrees C / 105 degrees F
Temperature LU 0 Chip      52 degrees C / 125 degrees F
Temperature MQ 0 TSen      41 degrees C / 105 degrees F
Temperature MQ 0 Chip      47 degrees C / 116 degrees F
Temperature LU 1 TCAM TSen 39 degrees C / 102 degrees F
Temperature LU 1 TCAM Chip 42 degrees C / 107 degrees F
Temperature LU 1 TSen      39 degrees C / 102 degrees F
Temperature LU 1 Chip      46 degrees C / 114 degrees F
Temperature MQ 1 TSen      39 degrees C / 102 degrees F
Temperature MQ 1 Chip      45 degrees C / 113 degrees F
Power
MPC-BIAS3V3-z12105      3296 mV
MPC-VDD3V3-z12006      3298 mV
MPC-VDD2V5-z12006      2505 mV
MPC-TCAM_1V0-z12004      997 mV
MPC-AVDD1V0-z12006      1007 mV
MPC-VDD1V8-z12006      1803 mV
MPC-PCIE_1V0-z12006      1004 mV
MPC-LU0_1V0-z12004      1000 mV
MPC-MQ0_1V0-z12004      999 mV
MPC-VDD_1V5-z12004      1498 mV
MPC-PMB_1V1-z12006      1102 mV
MPC-9VA-BMR453          9009 mV
MPC-9VB-BMR453          8960 mV
MPC-PMB_1V2-z12105      1202 mV
MPC-LU1_1V0-z12004      1005 mV
MPC-MQ1_1V0-z12004      1000 mV
I2C Slave Revision      70
FPC 9 status:
State                    Online
Temperature Intake        34 degrees C / 93 degrees F
Temperature Exhaust A     41 degrees C / 105 degrees F

```

Temperature Exhaust B	54 degrees C / 129 degrees F
Temperature LU 0 TSen	51 degrees C / 123 degrees F
Temperature LU 0 Chip	52 degrees C / 125 degrees F
Temperature LU 1 TSen	51 degrees C / 123 degrees F
Temperature LU 1 Chip	55 degrees C / 131 degrees F
Temperature LU 2 TSen	51 degrees C / 123 degrees F
Temperature LU 2 Chip	47 degrees C / 116 degrees F
Temperature LU 3 TSen	51 degrees C / 123 degrees F
Temperature LU 3 Chip	47 degrees C / 116 degrees F
Temperature MQ 0 TSen	40 degrees C / 104 degrees F
Temperature MQ 0 Chip	42 degrees C / 107 degrees F
Temperature MQ 1 TSen	40 degrees C / 104 degrees F
Temperature MQ 1 Chip	44 degrees C / 111 degrees F
Temperature MQ 2 TSen	40 degrees C / 104 degrees F
Temperature MQ 2 Chip	38 degrees C / 100 degrees F
Temperature MQ 3 TSen	40 degrees C / 104 degrees F
Temperature MQ 3 Chip	40 degrees C / 104 degrees F
Power	
AS-BIAS3V3-z12105	3302 mV
AS-VDD1V8-z12006	1808 mV
AS-VDD2V5-z12006	2513 mV
AS-AVDD1V0-z12004	997 mV
AS-PCIE_1V0-z12004	999 mV
AS-VDD3V3-z12004	3294 mV
AS-VDD_1V5A-z12004	1503 mV
AS-VDD_1V5B-z12004	1502 mV
AS-LU0_1V0-z12004	996 mV
AS-LU1_1V0-z12004	999 mV
AS-MQ0_1V0-z12004	997 mV
AS-MQ1_1V0-z12004	999 mV
AS-LU2_1V0-z12004	997 mV
AS-LU3_1V0-z12004	998 mV
AS-MQ2_1V0-z12004	1000 mV
AS-MQ3_1V0-z12004	1000 mV
AS-PMB_1V1-z12006	1102 mV
I2C Slave Revision	68

show chassis environment fpc (MX240 Router)

```
user@host> show chassis environment fpc
```

```
FPC 1 status:
```

State	Online
Temperature Intake	34 degrees C / 93 degrees F
Temperature Exhaust A	39 degrees C / 102 degrees F
Temperature Exhaust B	53 degrees C / 127 degrees F
Temperature I3 0 TSensor	51 degrees C / 123 degrees F
Temperature I3 0 Chip	54 degrees C / 129 degrees F
Temperature I3 1 TSensor	50 degrees C / 122 degrees F
Temperature I3 1 Chip	53 degrees C / 127 degrees F
Temperature I3 2 TSensor	48 degrees C / 118 degrees F
Temperature I3 2 Chip	51 degrees C / 123 degrees F
Temperature I3 3 TSensor	45 degrees C / 113 degrees F
Temperature I3 3 Chip	48 degrees C / 118 degrees F
Temperature IA 0 TSensor	45 degrees C / 113 degrees F
Temperature IA 0 Chip	45 degrees C / 113 degrees F
Temperature IA 1 TSensor	45 degrees C / 113 degrees F
Temperature IA 1 Chip	49 degrees C / 120 degrees F
Power	
1.5 V	1492 mV
2.5 V	2507 mV
3.3 V	3306 mV

```

1.8 V PFE 0          1801 mV
1.8 V PFE 1          1804 mV
1.8 V PFE 2          1798 mV
1.8 V PFE 3          1798 mV
1.2 V PFE 0          1169 mV
1.2 V PFE 1          1189 mV
1.2 V PFE 2          1182 mV
1.2 V PFE 3          1176 mV
I2C Slave Revision   42
FPC 2 status:
State                Online
Temperature Intake    33 degrees C / 91 degrees F
Temperature Exhaust A 41 degrees C / 105 degrees F
Temperature Exhaust B 53 degrees C / 127 degrees F
Temperature I3 0 TSensor 53 degrees C / 127 degrees F
Temperature I3 0 Chip  58 degrees C / 136 degrees F
Temperature I3 1 TSensor 52 degrees C / 125 degrees F
Temperature I3 1 Chip  56 degrees C / 132 degrees F
Temperature I3 2 TSensor 50 degrees C / 122 degrees F
Temperature I3 2 Chip  52 degrees C / 125 degrees F
Temperature I3 3 TSensor 46 degrees C / 114 degrees F
Temperature I3 3 Chip  49 degrees C / 120 degrees F
Temperature IA 0 TSensor 51 degrees C / 123 degrees F
Temperature IA 0 Chip  49 degrees C / 120 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip  53 degrees C / 127 degrees F
Power
1.5 V                1492 mV
2.5 V                2445 mV
3.3 V                3293 mV
1.8 V PFE 0          1827 mV
1.8 V PFE 1          1775 mV
1.8 V PFE 2          1788 mV
1.8 V PFE 3          1798 mV
1.2 V PFE 0          1250 mV
1.2 V PFE 1          1234 mV
1.2 V PFE 2          1231 mV
1.2 V PFE 3          1192 mV
I2C Slave Revision   42

```

show chassis environment fpc (MX480 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                Online
Temperature Intake    36 degrees C / 96 degrees F
Temperature Exhaust A 41 degrees C / 105 degrees F
Temperature Exhaust B 55 degrees C / 131 degrees F
Temperature I3 0 TSensor 55 degrees C / 131 degrees F
Temperature I3 0 Chip  57 degrees C / 134 degrees F
Temperature I3 1 TSensor 53 degrees C / 127 degrees F
Temperature I3 1 Chip  53 degrees C / 127 degrees F
Temperature I3 2 TSensor 52 degrees C / 125 degrees F
Temperature I3 2 Chip  49 degrees C / 120 degrees F
Temperature I3 3 TSensor 47 degrees C / 116 degrees F
Temperature I3 3 Chip  47 degrees C / 116 degrees F
Temperature IA 0 TSensor 54 degrees C / 129 degrees F
Temperature IA 0 Chip  58 degrees C / 136 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip  53 degrees C / 127 degrees F
Power

```

1.5 V	1479 mV
2.5 V	2542 mV
3.3 V	3319 mV
1.8 V PFE 0	1811 mV
1.8 V PFE 1	1804 mV
1.8 V PFE 2	1804 mV
1.8 V PFE 3	1814 mV
1.2 V PFE 0	1192 mV
1.2 V PFE 1	1202 mV
1.2 V PFE 2	1205 mV
1.2 V PFE 3	1189 mV
I2C Slave Revision	40

show chassis environment fpc (MX960 Router)

```
user@host> show chassis environment fpc
```

```
FPC 5 status:
```

State	Online
Temperature Intake	27 degrees C / 80 degrees F
Temperature Exhaust A	34 degrees C / 93 degrees F
Temperature Exhaust B	40 degrees C / 104 degrees F
Temperature I3 0 TSensor	39 degrees C / 102 degrees F
Temperature I3 0 Chip	41 degrees C / 105 degrees F
Temperature I3 1 TSensor	38 degrees C / 100 degrees F
Temperature I3 1 Chip	37 degrees C / 98 degrees F
Temperature I3 2 TSensor	37 degrees C / 98 degrees F
Temperature I3 2 Chip	34 degrees C / 93 degrees F
Temperature I3 3 TSensor	32 degrees C / 89 degrees F
Temperature I3 3 Chip	33 degrees C / 91 degrees F
Temperature IA 0 TSensor	39 degrees C / 102 degrees F
Temperature IA 0 Chip	44 degrees C / 111 degrees F
Temperature IA 1 TSensor	36 degrees C / 96 degrees F
Temperature IA 1 Chip	44 degrees C / 111 degrees F
Power	
1.5 V	1479 mV
2.5 V	2523 mV
3.3 V	3254 mV
1.8 V PFE 0	1798 mV
1.8 V PFE 1	1798 mV
1.8 V PFE 2	1807 mV
1.8 V PFE 3	1791 mV
1.2 V PFE 0	1173 mV
1.2 V PFE 1	1179 mV
1.2 V PFE 2	1179 mV
1.2 V PFE 3	1185 mV
I2C Slave Revision	6

```
FPC 6 status:
```

State	Online
Temperature Intake	25 degrees C / 77 degrees F
Temperature Exhaust A	38 degrees C / 100 degrees F
Temperature Exhaust B	38 degrees C / 100 degrees F
Temperature I3 0 TSensor	40 degrees C / 104 degrees F
Temperature I3 0 Chip	40 degrees C / 104 degrees F
Temperature I3 1 TSensor	40 degrees C / 104 degrees F
Temperature I3 1 Chip	38 degrees C / 100 degrees F
Temperature I3 2 TSensor	37 degrees C / 98 degrees F
Temperature I3 2 Chip	32 degrees C / 89 degrees F
Temperature I3 3 TSensor	34 degrees C / 93 degrees F
Temperature I3 3 Chip	33 degrees C / 91 degrees F
Temperature IA 0 TSensor	45 degrees C / 113 degrees F
Temperature IA 0 Chip	47 degrees C / 116 degrees F

```

Temperature IA 1 TSensor 37 degrees C / 98 degrees F
Temperature IA 1 Chip    42 degrees C / 107 degrees F
Power
  1.5 V                  1485 mV
  2.5 V                  2510 mV
  3.3 V                  3332 mV
  1.8 V PFE 0            1801 mV
  1.8 V PFE 1            1814 mV
  1.8 V PFE 2            1804 mV
  1.8 V PFE 3            1820 mV
  1.2 V PFE 0            1192 mV
  1.2 V PFE 1            1189 mV
  1.2 V PFE 2            1202 mV
  1.2 V PFE 3            1156 mV
I2C Slave Revision      40

```

show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  32 degrees C / 89 degrees F
Temperature Exhaust A               39 degrees C / 102 degrees F
Temperature Exhaust B               37 degrees C / 98 degrees F
Temperature QX 0 TSen               44 degrees C / 111 degrees F
Temperature QX 0 Chip               48 degrees C / 118 degrees F
Temperature LU 0 TCAM TSen          44 degrees C / 111 degrees F
Temperature LU 0 TCAM Chip          47 degrees C / 116 degrees F
Temperature LU 0 TSen               44 degrees C / 111 degrees F
Temperature LU 0 Chip               48 degrees C / 118 degrees F
Temperature MQ 0 TSen               44 degrees C / 111 degrees F
Temperature MQ 0 Chip               47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105                 3297 mV
MPC-VDD3V3-z12105                 3306 mV
MPC-VDD2V5-z12105                 2498 mV
MPC-TCAM_1V0-z12004                999 mV
MPC-AVDD1V0-z12006                 999 mV
MPC-VDD1V8-z12006                 1796 mV
MPC-PCIE_1V0-z12006                1002 mV
MPC-LU0_1V0-z12004                 997 mV
MPC-MQ0_1V0-z12004                 995 mV
MPC-VDD_1V5-z12004                1496 mV
MPC-PMB_1V1-z12006                1094 mV
MPC-9VA-BMR453                    9054 mV
MPC-9VB-BMR453                    9037 mV
MPC-PMB_1V2-z12106                1191 mV
MPC-QXM0_1V0-z12006               1000 mV
I2C Slave Revision                 66
FPC 1 status:
State                               Online
Temperature Intake                  35 degrees C / 95 degrees F
Temperature Exhaust A               50 degrees C / 122 degrees F
Temperature Exhaust B               56 degrees C / 132 degrees F
Temperature LU 0 TSen               46 degrees C / 114 degrees F
Temperature LU 0 Chip               59 degrees C / 138 degrees F
Temperature LU 1 TSen               46 degrees C / 114 degrees F
Temperature LU 1 Chip               45 degrees C / 113 degrees F
Temperature LU 2 TSen               46 degrees C / 114 degrees F
Temperature LU 2 Chip               60 degrees C / 140 degrees F
Temperature LU 3 TSen               46 degrees C / 114 degrees F

```

```

Temperature LU 3 Chip      71 degrees C / 159 degrees F
Temperature XM 0 TSen      46 degrees C / 114 degrees F
Temperature XM 0 Chip      -18 degrees C / 0 degrees F
Temperature XF 0 TSen      46 degrees C / 114 degrees F
Temperature XF 0 Chip      76 degrees C / 168 degrees F
Power
MPC-BIAS3V3-z12105        3292 mV
MPC-VDD3V3-z16100         3303 mV
MPC-VDD2V5-z16100         2501 mV
MPC-VDD1V8-z12004         1801 mV
MPC-AVDD1V0-z12006         996 mV
MPC-VDD1V2-z16100         1199 mV
MPC-VDD1V5A-z12004        1493 mV
MPC-VDD1V5B-z12004        1498 mV
MPC-XF_0V9-z12006         996 mV
MPC-PCIE_1V0-z16100       1000 mV
MPC-LU0_1V0-z12004         994 mV
MPC-LU1_1V0-z12004         994 mV
MPC-LU2_1V0-z12004         992 mV
MPC-LU3_1V0-z12004         993 mV
MPC-12VA-BMR453           12003 mV
MPC-12VB-BMR453           12043 mV
MPC-PMB_1V1-z12006        1091 mV
MPC-PMB_1V2-z12106        1196 mV
MPC-XM_0V9-vt273m         899 mV
I2C Slave Revision        106

```

show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis environment fpc 1
FPC 1 status:
State                Online
Temperature Intake    36 degrees C / 96 degrees F
Temperature Exhaust A 39 degrees C / 102 degrees F
Temperature LU TSen    52 degrees C / 125 degrees F
Temperature LU Chip    54 degrees C / 129 degrees F
Temperature XM TSen    52 degrees C / 125 degrees F
Temperature XM Chip    60 degrees C / 140 degrees F
Temperature PCIE TSen  52 degrees C / 125 degrees F
Temperature PCIE Chip  69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106    3302 mV
MPC-VDD3V3-z16100     3325 mV
MPC-AVDD1V0-z16100    1007 mV
MPC-PCIE_1V0-z16100    904 mV
MPC-LU0_1V0-z12004     996 mV
MPC-VDD_1V5-z12004    1498 mV
MPC-12VA-BMR453       11733 mV
MPC-12VB-BMR453       11728 mV
MPC-XM_0V9-vt273m     900 mV
I2C Slave Revision    81

```

show chassis environment fpc (T320, T640, and T1600 Routers)

```

user@host> show chassis environment fpc
FPC 0 status:
State                Online
Temperature Top       42 degrees C / 107 degrees F
Temperature Bottom    36 degrees C / 96 degrees F
Temperature MMB1       39 degrees C / 102 degrees F
Power:

```

```

1.8 V          1959 mV
2.5 V          2495 mV
3.3 V          3344 mV
5.0 V          5047 mV
1.8 V bias     1787 mV
3.3 V bias     3291 mV
5.0 V bias     4998 mV
8.0 V bias     7343 mV
BUS Revision   40
FPC 1 status:
State          Online
Temperature Top 42 degrees C / 107 degrees F
Temperature Bottom 39 degrees C / 102 degrees F
Temperature MMB1 40 degrees C / 104 degrees F
Power:
1.8 V          1956 mV
2.5 V          2498 mV
3.3 V          3340 mV
5.0 V          5023 mV
1.8 V bias     1782 mV
3.3 V bias     3277 mV
5.0 V bias     4989 mV
8.0 V bias     7289 mV
BUS Revision   40
FPC 2 status:
State          Online
Temperature Top 43 degrees C / 109 degrees F
Temperature Bottom 39 degrees C / 102 degrees F
Temperature MMB1 41 degrees C / 105 degrees F
Power:
1.8 V          1963 mV
2.5 V          2503 mV
3.3 V          3340 mV
5.0 V          5042 mV
1.8 V bias     1797 mV
3.3 V bias     3311 mV
5.0 V bias     5013 mV
8.0 V bias     7221 mV
BUS Revision   40

```

show chassis environment fpc (T4000 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State          Online
Fan Intake     34 degrees C / 93 degrees F
Fan Exhaust    48 degrees C / 118 degrees F
PMB            47 degrees C / 116 degrees F
LMB0           50 degrees C / 122 degrees F
LMB1           41 degrees C / 105 degrees F
LMB2           35 degrees C / 95 degrees F
PFE1 LU2       46 degrees C / 114 degrees F
PFE1 LU0       41 degrees C / 105 degrees F
PFE0 LU0       57 degrees C / 134 degrees F
XF1            47 degrees C / 116 degrees F
XF0            52 degrees C / 125 degrees F
XM1            41 degrees C / 105 degrees F
XM0            50 degrees C / 122 degrees F
PFE0 LU1       56 degrees C / 132 degrees F
PFE0 LU2       45 degrees C / 113 degrees F
PFE1 LU1       37 degrees C / 98 degrees F

```

Power 1	
1.0 V	991 mV
1.2 V bias	1195 mV
1.8 V	1788 mV
2.5 V	2483 mV
3.3 V	3289 mV
3.3 V bias	3299 mV
12.0 V A	10608 mV
12.0 V B	10637 mV
Power 2	
0.9 V	881 mV
0.9 V PFE0	916 mV
0.9 V PFE1	903 mV
1.0 V PFE0	1012 mV
1.0 V PFE1	1002 mV
1.1 V	1095 mV
1.5 V_0	1494 mV
1.5 V_1	1479 mV
Power 3	
1.0 V PFE0	1000 mV
1.0 V PFE1	1002 mV
1.0 V PFE0 *	995 mV
1.0 V PFE1 *	995 mV
1.8 V PFE 0	1788 mV
1.8 V PFE 1	1789 mV
2.5 V	2482 mV
12.0 V	11614 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1003 mV
1.0 V PFE1 LU2	1004 mV
1.0 V PFE0 LU0 *	995 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	996 mV
12.0 V	11643 mV
12.0 V C	11711 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2488 mV
LMB0 VDD1V8	1788 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	1002 mV
LMB0 PFE0 LU0 VDD1V0	1000 mV
LMB0 VDD12V0	10752 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1792 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1002 mV
LMB1 VDD12V0	10800 mV
LMB2 VDD2V5	2472 mV
LMB2 VDD1V8	1792 mV
LMB2 VDD1V5	1486 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	998 mV
LMB2 VDD12V0	10704 mV
PMB 1.05v	1049 mV
PMB 1.5v	1500 mV
PMB 2.5v	2500 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 3 status:	

State	Online
Fan Intake	37 degrees C / 98 degrees F
Fan Exhaust	51 degrees C / 123 degrees F
PMB	43 degrees C / 109 degrees F
LMB0	57 degrees C / 134 degrees F
LMB1	54 degrees C / 129 degrees F
LMB2	38 degrees C / 100 degrees F
PFE1 LU2	63 degrees C / 145 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	69 degrees C / 156 degrees F
XF1	62 degrees C / 143 degrees F
XF0	63 degrees C / 145 degrees F
XM1	43 degrees C / 109 degrees F
XM0	67 degrees C / 152 degrees F
PFE0 LU1	63 degrees C / 145 degrees F
PFE0 LU2	66 degrees C / 150 degrees F
PFE1 LU1	41 degrees C / 105 degrees F
Power 1	
1.0 V	1002 mV
1.2 V bias	1201 mV
1.8 V	1785 mV
2.5 V	2485 mV
3.3 V	3288 mV
3.3 V bias	3285 mV
12.0 V A	10412 mV
12.0 V B	10515 mV
Power 2	
0.9 V	882 mV
0.9 V PFE0	920 mV
0.9 V PFE1	905 mV
1.0 V PFE0	1015 mV
1.0 V PFE1	1001 mV
1.1 V	1094 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
0.92 V PFE1	998 mV
1.0 V PFE0	997 mV
1.0 V PFE0 *	992 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1780 mV
1.8 V PFE 1	1797 mV
2.5 V	2492 mV
12.0 V	11604 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1004 mV
1.0 V PFE1 LU2	1003 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	1001 mV
1.0 V PFE1 LU2 *	1003 mV
12.0 V	11653 mV
12.0 V C	11672 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2512 mV
LMB0 VDD1V8	1790 mV
LMB0 VDD1V5	1500 mV
LMB0 PFE0 LU0 AVDD1V0	1004 mV
LMB0 PFE0 LU0 VDD1V0	1002 mV
LMB0 VDD12V0	10608 mV
LMB1 VDD2V5	2472 mV

LMB1 VDD1V8	1788 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	1000 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10672 mV
LMB2 VDD2V5	2488 mV
LMB2 VDD1V8	1798 mV
LMB2 VDD1V5	1494 mV
LMB2 PFE1 LU1 AVDD1V0	1000 mV
LMB2 PFE1 LU1 VDD1V0	1004 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1500 mV
PMB 2.5v	2499 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 5 status:	
State	Online
Temperature Top	39 degrees C / 102 degrees F
Temperature Bottom	38 degrees C / 100 degrees F
Power	
1.8 V	1804 mV
1.8 V bias	1802 mV
3.3 V	3294 mV
3.3 V bias	3277 mV
5.0 V bias	5008 mV
5.0 V TOP	5067 mV
8.0 V bias	6642 mV
Power (Base/PMB/MMB)	
1.2 V	1202 mV
1.5 V	1504 mV
5.0 V BOT	5079 mV
12.0 V TOP Base	11848 mV
12.0 V BOT Base	11780 mV
1.1 V PMB	1111 mV
1.2 V PMB	1189 mV
1.5 V PMB	1494 mV
1.8 V PMB	1819 mV
2.5 V PMB	2503 mV
3.3 V PMB	3294 mV
5.0 V PMB	5035 mV
12.0 V PMB	11788 mV
0.75 MMB TOP	766 mV
1.5 V MMB TOP	1484 mV
1.8 V MMB TOP	1772 mV
2.5 V MMB TOP	2485 mV
1.2 V MMB TOP	1137 mV
5.0 V MMB TOP	4946 mV
12.0 V MMB TOP	11772 mV
3.3 V MMB TOP	3289 mV
0.75 MMB BOT	759 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1792 mV
2.5 V MMB BOT	2490 mV
1.2 V MMB BOT	1145 mV
5.0 V MMB BOT	4922 mV
12.0 V MMB BOT	11625 mV
3.3 V MMB BOT	3282 mV
APS 00	2495 mV
APS 01	3308 mV
APS 02	3301 mV

```

5.0 V PIC 0          4967 mV
APS 10              2512 mV
APS 11              3316 mV
APS 12              3304 mV
5.0 V PIC 1          5081 mV
Bus Revision         49
FPC 6 status:
State                Online
Fan Intake           34 degrees C / 93 degrees F
Fan Exhaust          49 degrees C / 120 degrees F
PMB                  40 degrees C / 104 degrees F
LMB0                  60 degrees C / 140 degrees F
LMB1                  58 degrees C / 136 degrees F
LMB2                  40 degrees C / 104 degrees F
PFE1 LU2             69 degrees C / 156 degrees F
PFE1 LU0             45 degrees C / 113 degrees F
PFE0 LU0             71 degrees C / 159 degrees F
XF1                  58 degrees C / 136 degrees F
XF0                  65 degrees C / 149 degrees F
XM1                  40 degrees C / 104 degrees F
XM0                  66 degrees C / 150 degrees F
PFE0 LU1             69 degrees C / 156 degrees F
PFE0 LU2             68 degrees C / 154 degrees F
PFE1 LU1             42 degrees C / 107 degrees F
Power 1
1.0 V                998 mV
1.2 V bias           1191 mV
1.8 V                1781 mV
2.5 V                2487 mV
3.3 V                3302 mV
3.3 V bias           3300 mV
12.0 V A             10388 mV
12.0 V B             10388 mV
Power 2
0.9 V                902 mV
0.9 V PFE0           921 mV
0.9 V PFE1           907 mV
1.0 V PFE0           996 mV
1.0 V PFE1           974 mV
1.1 V                1095 mV
1.5 V_0              1495 mV
1.5 V_1              1478 mV
Power 3
1.0 V PFE0           997 mV
1.0 V PFE1           998 mV
1.0 V PFE0 *         993 mV
1.0 V PFE1 *         991 mV
1.8 V PFE 0          1796 mV
1.8 V PFE 1          1789 mV
2.5 V                2465 mV
12.0 V               11609 mV
Power 4
1.0 V PFE0 LU0       1003 mV
1.0 V PFE1 LU0       1006 mV
1.0 V PFE1 LU2       1002 mV
1.0 V PFE0 LU0 *     1000 mV
1.0 V PFE1 LU0 *     998 mV
1.0 V PFE1 LU2 *     998 mV
12.0 V               11638 mV
12.0 V C             11702 mV
Power (Base/PMB/MMB)

```

LMB0 VDD2V5	2484 mV
LMB0 VDD1V8	1780 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	998 mV
LMB0 PFE0 LU0 VDD1V0	1004 mV
LMB0 VDD12V0	10528 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1776 mV
LMB1 VDD1V5	1474 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10544 mV
LMB2 VDD2V5	2476 mV
LMB2 VDD1V8	1790 mV
LMB2 VDD1V5	1492 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	1010 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1499 mV
PMB 2.5v	2500 mV
PMB 3.3v	3300 mV
Bus Revision	80

show chassis environment fpc lcc (TX Matrix Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

FPC 1 status:

State	Online
Temperature Top	30 degrees C / 86 degrees F
Temperature Bottom	25 degrees C / 77 degrees F
Temperature MMB0	Absent
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1813 mV
2.5 V	2504 mV
3.3 V	3338 mV
5.0 V	5037 mV
1.8 V bias	1797 mV
3.3 V bias	3301 mV
5.0 V bias	5013 mV
8.0 V bias	7345 mV
BUS Revision	40

FPC 2 status:

State	Online
Temperature Top	37 degrees C / 98 degrees F
Temperature Bottom	26 degrees C / 78 degrees F
Temperature MMB0	32 degrees C / 89 degrees F
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1791 mV
2.5 V	2517 mV
3.3 V	3308 mV
5.0 V	5052 mV
1.8 V bias	1797 mV
3.3 V bias	3289 mV
5.0 V bias	4991 mV
8.0 V bias	7477 mV
BUS Revision	40

show chassis environment fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

```
-----
FPC 1 status:
```

State	Online
Temperature Top	46 degrees C / 114 degrees F
Temperature Bottom	47 degrees C / 116 degrees F
Power	
1.8 V	1788 mV
1.8 V bias	1787 mV
3.3 V	3321 mV
3.3 V bias	3306 mV
5.0 V bias	5018 mV
5.0 V TOP	5037 mV
8.0 V bias	7223 mV
Power (Base/PMB/MMB)	
1.2 V	1205 mV
1.5 V	1503 mV
5.0 V BOT	5084 mV
12.0 V TOP Base	11775 mV
12.0 V BOT Base	11794 mV
1.1 V PMB	1108 mV
1.2 V PMB	1196 mV
1.5 V PMB	1499 mV
1.8 V PMB	1811 mV
2.5 V PMB	2515 mV
3.3 V PMB	3318 mV
5.0 V PMB	5030 mV
12.0 V PMB	11832 mV
0.75 MMB TOP	752 mV
1.5 V MMB TOP	1489 mV
1.8 V MMB TOP	1782 mV
2.5 V MMB TOP	2498 mV
1.2 V MMB TOP	1155 mV
5.0 V MMB TOP	4902 mV
12.0 V MMB TOP	11721 mV
3.3 V MMB TOP	3316 mV
0.75 MMB BOT	754 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1758 mV
2.5 V MMB BOT	2488 mV
1.2 V MMB BOT	1157 mV
5.0 V MMB BOT	4962 mV
12.0 V MMB BOT	11691 mV
3.3 V MMB BOT	3308 mV
APS 00	1484 mV
APS 01	2503 mV
APS 02	3313 mV
5.0 V PIC 0	5025 mV
APS 10	1501 mV
APS 11	2466 mV
APS 12	3311 mV
5.0 V PIC 1	5081 mV
Bus Revision	49

show chassis environment fpc (QFX Series)

```
user@switch> show chassis environment fpc 0
```

```

FPC 0 status:
State          Online
Temperature    42 degrees C / 107 degrees F

```

show chassis environment fpc interconnect-device (QFabric Systems)

```

user@switch> show chassis environment fpc interconnect-device interconnect1 0
FC 0 FPC 0 status:
State          Online
Left Intake Temperature    24 degrees C / 75 degrees F
Right Intake Temperature   24 degrees C / 75 degrees F
Left Exhaust Temperature   27 degrees C / 80 degrees F
Right Exhaust Temperature  27 degrees C / 80 degrees F
Power
  BIAS 3V3          3330 mV
  VDD 3V3           3300 mV
  VDD 2V5           2502 mV
  VDD 1V5           1496 mV
  VDD 1V2           1194 mV
  VDD 1V0           1000 mV
  SW0 VDD 1V0       1020 mV
  SW0 CVDD 1V025    1032 mV
  SW1 VDD 1V0       1022 mV
  SW1 CVDD 1V025    1030 mV
  VDD 12V0 DIV3_33  3414 mV

```

show chassis environment fpc 0 (PTX5000 Packet Transport Switch)

```

user@switch> show chassis environment fpc 0
FPC 0 status:
State          Online
PMB Temperature    35 degrees C / 95 degrees F
Intake Temperature  33 degrees C / 91 degrees F
Exhaust A Temperature  51 degrees C / 123 degrees F
Exhaust B Temperature  43 degrees C / 109 degrees F
TL0 Temperature     48 degrees C / 118 degrees F
TQ0 Temperature     53 degrees C / 127 degrees F
TL1 Temperature     56 degrees C / 132 degrees F
TQ1 Temperature     58 degrees C / 136 degrees F
TL2 Temperature     55 degrees C / 131 degrees F
TQ2 Temperature     57 degrees C / 134 degrees F
TL3 Temperature     59 degrees C / 138 degrees F
TQ3 Temperature     59 degrees C / 138 degrees F
Power
  PMB 1.05v         1049 mV
  PMB 1.5v          1500 mV
  PMB 2.5v          2500 mV
  PMB 3.3v          3299 mV
  PFE0 1.5v         1500 mV
  PFE0 1.0v         999 mV
  TQ0 0.9v          900 mV
  TL0 0.9v          900 mV
  PFE1 1.5v         1499 mV
  PFE1 1.0v         999 mV
  TQ1 0.9v          899 mV
  TL1 0.9v          900 mV
  PFE2 1.5v         1500 mV
  PFE2 1.0v         1000 mV
  TQ2 0.9v          900 mV
  TL2 0.9v          900 mV
  PFE3 1.5v         1499 mV

```

PFE3	1.0v	1000 mV
TQ3	0.9v	900 mV
TL3	0.9v	900 mV
Bias	3.3v	3327 mV
FPC	3.3v	3300 mV
FPC	2.5v	2500 mV
SAM	0.9v	900 mV
A	12.0v	2014 mV
B	12.0v	2030 mV

show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis environment fpc 1
FPC 1 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A               39 degrees C / 102 degrees F
Temperature LU TSen                  52 degrees C / 125 degrees F
Temperature LU Chip                  54 degrees C / 129 degrees F
Temperature XM TSen                  52 degrees C / 125 degrees F
Temperature XM Chip                  60 degrees C / 140 degrees F
Temperature PCIE TSen                52 degrees C / 125 degrees F
Temperature PCIE Chip                69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106                  3302 mV
MPC-VDD3V3-z16100                    3325 mV
MPC-AVDD1V0-z16100                   1007 mV
MPC-PCIE_1V0-z16100                   904 mV
MPC-LU0_1V0-z12004                    996 mV
MPC-VDD_1V5-z12004                   1498 mV
MPC-12VA-BMR453                      11733 mV
MPC-12VB-BMR453                      11728 mV
MPC-XM_0V9-vt273m                    900 mV
I2C Slave Revision                   81

```

show chassis environment pem

List of Syntax	Syntax on page 604 Syntax (ACX4000 Router) on page 604 Syntax (TX Matrix Routers) on page 604 Syntax (TX Matrix Plus Routers) on page 604 Syntax (MX Series Router) on page 604 Syntax (QFX Series) on page 604
Syntax	show chassis environment pem <slot>
Syntax (ACX4000 Router)	show chassis environment pem
Syntax (TX Matrix Routers)	show chassis environment pem <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment pem <lcc number sfc number> <slot>
Syntax (MX Series Router)	show chassis environment pem <slot> <all-members> <local> <member member-id>
Syntax (QFX Series)	show chassis environment pem <slot (interconnect-device name slot) (node-device name)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS 11.3 for the QFX Series.
Description	Display Power Entry Module (PEM) environmental status information.



NOTE: The new high-capacity (4100W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

Options **none**—Display environmental information about both PEMs. For the TX Matrix router, display environmental information about the PEMs, the TX Matrix router, and its attached T640 routers. For the TX Matrix Plus router, display environmental information about the PEMs, the TX Matrix Plus router, and its attached routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the PEMs in all the member routers of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information about the PEM in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the PEM in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Node device.

scc—(TX Matrix routers only) (Optional) Display environmental information about the PEM in the TX Matrix router (switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display environmental information about the PEM in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about an individual PEM. Replace *slot* with 0 or 1.

Required Privilege Level

view

Related Documentation

- [show chassis hardware on page 667](#)

List of Sample Output

[show chassis environment pem \(M40e Router\) on page 607](#)
[show chassis environment pem \(M120 Router\) on page 607](#)
[show chassis environment pem \(M160 Router\) on page 607](#)
[show chassis environment pem \(M320 Router\) on page 607](#)
[show chassis environment pem \(MX240 Router\) on page 608](#)
[show chassis environment pem \(MX480 Router\) on page 608](#)

[show chassis environment pem \(MX960 Router\) on page 608](#)
[show chassis environment pem \(T320 Router\) on page 608](#)
[show chassis environment pem \(T640 Router\) on page 608](#)
[show chassis environment pem \(T4000 Router\) on page 609](#)
[show chassis environment pem \(T640/T1600/T4000 Routers With Six-Input DC Power Supply\) on page 609](#)
[show chassis environment pem lcc \(TX Matrix Routing Matrix\) on page 610](#)
[show chassis environment pem scc \(TX Matrix Routing Matrix\) on page 610](#)
[show chassis environment pem sfc \(TX Matrix Plus Routing Matrix\) on page 610](#)
[show chassis environment pem lcc \(TX Matrix Plus Routing Matrix\) on page 610](#)
[show chassis environment pem node-device \(QFabric System\) on page 611](#)
[show chassis environment pem \(QFX Series\) on page 611](#)
[show chassis environment pem interconnect-device \(QFabric System\) on page 612](#)

Output Fields [Table 57 on page 606](#) lists the output fields for the **show chassis environment pem** command. Output fields are listed in the approximate order in which they appear.

Table 57: show chassis environment pem Output Fields

Field Name	Field Description
PEM slot status	Number of the PEM slot.
State	Status of the PEM.
Temperature	Temperature of the air flowing past the PEM.
AC Input	Status of the AC input for the specified component
AC Output	Status of the AC output for the specified component.
DC input	Status of the DC input for the specified component.
DC output	Status of the DC output for the specified component.
Load	(Not available on M40e or M160 routers) Information about the load on supply, in percentage of rated current being used.
Voltage	(M120, M160, M320, T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about voltage supplied to the PEM.
Current	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM current.
Power	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM power.
SCG/CB/SIB	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) SONET Clock Generator/Control Board/Switch Interface Board.
FAN	(T640, T1600, and T4000 routers with six-input DC power supply only) Information about the DC output to the fan.

Sample Output

show chassis environment pem (M40e Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  AC input              OK
  DC output             OK
```

show chassis environment pem (M120 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       52864 mV
    48.0 V fan supply  41655 mV
    3.3 V              3399 mV
PEM 1 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       54537 mV
    48.0 V fan supply  42910 mV
    3.3 V              3506 mV
```

show chassis environment pem (M160 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC input              OK
  DC output             OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       54833 mV
    48.0 V fan supply  50549 mV
    8.0 V bias         8239 mV
    5.0 V bias         5006 mV
```

show chassis environment pem (M320 Router)

```
user@host> show chassis environment pem
PEM 2 status:
  State                Online
  Temperature           OK
  DC input              OK
  Load                Less than 40 percent
    48.0 V input       51853 mV
    48.0 V fan supply  48877 mV
    8.0 V bias         8449 mV
```

```
5.0 V bias          4998 mV
PEM 3 status:
  State              Online
  Temperature         OK
  DC input            OK
  Load              Less than 40 percent
    48.0 V input      51717 mV
    48.0 V fan supply 49076 mV
    8.0 V bias        8442 mV
    5.0 V bias        4998 mV
```

show chassis environment pem (MX240 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State              Online
  Temperature         OK
  DC Output:         OK
PEM 1 status:
  State              Online
  Temperature         OK
  DC Output:         OK
```

show chassis environment pem (MX480 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State              Online
  Temperature         OK
  DC Input:          OK
  DC Output:         OK
  Voltage:
PEM 1 status:
  State              Online
  Temperature         OK
  DC Input:          OK
  DC Output:         OK
  Voltage:
```

show chassis environment pem (MX960 Router)

```
user@host> show chassis environment pem
PEM 2 status:
  State              Present
PEM 3 status:
  State              Online
  Temperature         OK
  DC Output:         OK
```

show chassis environment pem (T320 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State              Online
  Temperature         OK
  DC input:          OK
```

show chassis environment pem (T640 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State              Online
```

```

Temperature                22 degrees C / 71 degrees F
AC input: OK
DC output:
      Voltage    Current    Power    Load
      FPC 0      56875      606      34      4
      FPC 1      57016      525      29      3
      FPC 2         0         0         0      0
      FPC 3         0         0         0      0
      FPC 4         0         0         0      0
      FPC 5         0         0         0      0
      FPC 6      57158      1581      90     12
      FPC 7         0         0         0      0
SCG/CB/SIB      56750      1125      63      5

```

show chassis environment pem (T4000 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature          33 degrees C / 91 degrees F
  DC Input:            OK
      Voltage(V)    Current(A)    Power(W)    Load(%)
      INPUT 0      54.625      9.812      535      22
      INPUT 1      54.625     10.250      559      23
      INPUT 2      55.125      0.125        6       0
      INPUT 3      54.500     10.062      548      22
      INPUT 4      54.750      9.375      513      21
      INPUT 5      54.750     10.187      557      23
  DC Output          Voltage(V)    Current(A)    Power(W)    Load(%)
      FPC 0      55.750     10.125      564      37
      FPC 1      51.625      0.000        0       0
      FPC 2      52.000      0.000        0       0
      FPC 3      55.062     10.437      574      38
      FPC 4      52.125      0.000        0       0
      FPC 5      55.000      9.375      515      34
      FPC 6      55.187      9.687      534      35
      FPC 7      51.437      0.000        0       0
      SCG/CB/SIB  55.375     15.750      872      35
      FAN        54.562     14.750      804      42

```

show chassis environment pem (T640/T1600/T4000 Routers With Six-Input DC Power Supply)

```

user@host> show chassis environment pem
PEM 1 status:
  State                Online
  Temperature          36 degrees C / 96 degrees F
  DC Input:            OK
      Voltage(V)    Current(A)    Power(W)    Load(%)
      INPUT 0         0.000      0.000        0       0
      INPUT 1      54.875      3.812      209      27
      INPUT 2      55.375      3.937      218      29
      INPUT 3      54.625      3.750      204      27
      INPUT 4      55.125      3.375      186      24
      INPUT 5      55.125      3.375      186      24
  DC Output          Voltage(V)    Current(A)    Power(W)    Load(%)
      FPC 0      52.312      0.000        0       0
      FPC 1      52.687      0.000        0       0
      FPC 2      52.812      0.000        0       0
      FPC 3      55.812      7.062      394      52
      FPC 4      52.625      0.000        0       0
      FPC 5      52.625      0.000        0       0
      FPC 6      52.750      0.000        0       0

```

FPC 7	52.750	0.000	0	0
SCG/CB/SIB	55.937	11.937	667	55
FAN	55.812	4.937	275	36

show chassis environment pem lcc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem 0 lcc 0
lcc0-re0:
```

```
-----
PEM 0 status:
State                               Present
Temperature                         27 degrees C / 80 degrees F
DC input:                           Check
DC output:                           Voltage Current      Power      Load
FPC 0                               0          0          0          0
FPC 1                               0          0          0          0
FPC 2                               0          0          0          0
FPC 3                               0          0          0          0
FPC 4                               0          0          0          0
FPC 5                               0          0          0          0
FPC 6                               0          0          0          0
FPC 7                               0          0          0          0
SCG/CB/SIB                          0          0          0          0
```

show chassis environment pem scc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem scc
scc-re0:
```

```
-----
PEM 1 status:
State                               Online
Temperature                         24 degrees C / 75 degrees F
DC input:                           OK
DC output:                           Voltage Current      Power      Load
SIB 0                               0          0          0          0
SIB 1                               0          0          0          0
SIB 2                               0          0          0          0
SIB 3                               56550       0          0          0
SIB 4                               55958      6912      386        51
```

show chassis environment pem sfc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment pem sfc 0
sfc0-re0:
```

```
-----
PEM 0 status:
State                               Online
Temperature                         35 degrees C / 95 degrees F
DC Input:                           OK
DC Output                           Voltage Current      Power      Load
Channel 0                          53820   14140      761        59
Channel 1                          53550   12720      681        53
Channel 2                          53840   12930      696        54
Channel 3                          53690   14990      804        63
Channel 4                          53620   15070      808        63
Channel 5                          53900   14820      798        62
Channel 6                          54120   5020       271        21
```

show chassis environment pem lcc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment lcc 0
```

```
1cc0-re1:
```

```
-----
PEM 0 status:
  State                Online
  Temperature          38 degrees C / 100 degrees F
  DC Input:           OK
  DC Output            Voltage    Current      Power      Load
    FPC 0              0          0            0          0
    FPC 1              0          0            0          0
    FPC 2              0          0            0          0
    FPC 3              0          0            0          0
    FPC 4             56408       7575          427        56
    FPC 5              0          0            0          0
    FPC 6             56266       7956          447        59
    FPC 7             56283       6100          343        45
    SCG/CB/SIB         55916       8950          500        41
PEM 1 status:
  State                Present
  Temperature          35 degrees C / 95 degrees F
  DC Input:           Check
  DC Output            Voltage    Current      Power      Load
    FPC 0              0          0            0          0
    FPC 1              0          0            0          0
    FPC 2              0          0            0          0
    FPC 3              0          0            0          0
    FPC 4              0          0            0          0
    FPC 5              0          0            0          0
    FPC 6              0          0            0          0
    FPC 7              0          0            0          0
    SCG/CB/SIB         0          0            0          0
```

show chassis environment pem node-device (QFabric System)

```
user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State                Check
  Airflow              Front to Back
  Temperature          OK
  AC Input:            OK
  DC Output            Voltage(V) Current(A) Power(W) Load(%)
                     12          10      120     18
FPC 0 PEM 1 status:
  State                Online
  Airflow              Back to Front
  Temperature          OK
  AC Input:            OK
  DC Output            Voltage(V) Current(A) Power(W) Load(%)
                     11          10      110     17
```

show chassis environment pem (QFX Series)

```
user@switch> show chassis environment pem
FPC 0 PEM 1 status:
  State                Online
  Airflow              Front to Back
  Temperature          OK
  AC Input:            OK
  DC Output            Voltage(V) Current(A) Power(W) Load(%)
                     12          17      204     31
```

show chassis environment pem interconnect-device (QFabric System)

```
user@switch> show chassis environment pem interconnect-device IC11
IC1 PEM 1 status:
  State                Online
  Airflow              Front to Back
  Temperature          OK
  AC Input:            OK
  DC Output            Voltage(V) Current(A) Power(W) Load(%)
                      12          18       216     33
```


show chassis environment routing-engine

List of Syntax	Syntax on page 613 Syntax (TX Matrix Routers) on page 613 Syntax (TX Matrix Plus Routers) on page 613 Syntax (MX2010 3D Universal Edge Routers) on page 613 Syntax (MX Series Routers) on page 613 Syntax (MX2020 3D Universal Edge Routers) on page 613 Syntax (QFX Series) on page 613
Syntax	show chassis environment routing-engine <slot>
Syntax (TX Matrix Routers)	show chassis environment routing-engine <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment routing-engine <lcc number sfc number> <slot>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment routing-engine <slot>
Syntax (MX Series Routers)	show chassis environment routing-engine <slot> <all-members> <local> <member member-id>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment routing-engine <slot>
Syntax (QFX Series)	show chassis environment routing-engine interconnect-device <i>name</i>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for the T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display Routing Engine environmental status information.
Options	none —Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix

router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all Routing Engines on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with the value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX2020 routers, and T Series routers, replace **slot** with 0 or 1. On M5, M7i, M10, and M40 routers and on the J Series router, replace **slot** with 0. On EX3200 and EX4200 standalone switches, replace **slot** with 0. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace **slot** with 0 or 1. On the QFX3500 switch, there is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Switches, replace **slot** with 0 or 1.

Required Privilege Level view

- Related Documentation**
- [request chassis routing-engine master on page 367](#)
 - [show chassis routing-engine on page 831](#)

- List of Sample Output**
- [show chassis environment routing-engine \(Nonredundant\) on page 615](#)
 - [show chassis environment routing-engine \(Redundant\) on page 615](#)
 - [show chassis environment routing-engine \(MX2010 Router\) on page 616](#)
 - [show chassis environment routing-engine \(MX2020 Router\) on page 616](#)
 - [show chassis environment routing-engine \(TX Matrix Plus Router\) on page 616](#)
 - [show chassis environment routing-engine \(T4000 Core Router\) on page 616](#)
 - [show chassis environment routing-engine \(QFX Series\) on page 616](#)
 - [show chassis environment routing-engine interconnect-device \(QFabric System\) on page 617](#)
 - [show chassis environment routing-engine \(PTX5000 Packet Transport Switch\) on page 617](#)

- Output Fields**
- Table 58 on page 615 lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 58: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine <i>slot</i> status	Number of the Routing Engine slot: 0 or 1.
State	Status of the Routing Engine: <ul style="list-style-type: none"> • Online Master—Routing Engine is online, operating as Master. • Online Standby—Routing Engine is online, operating as Standby. • Offline—Routing Engine is offline.
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

Sample Output

show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          27 degrees C / 80 degrees
```

show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State                Online Master
  Temperature          26 degrees C / 78 degrees F
Route Engine 1 status:
  State                Online Standby
  Temperature          26 degrees C / 78 degrees F
```

show chassis environment routing-engine (MX2010 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      37 degrees C / 98 degrees F
  CPU Temperature  37 degrees C / 98 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F
```

show chassis environment routing-engine (MX2020 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      44 degrees C / 111 degrees F
  CPU Temperature  43 degrees C / 109 degrees F
```

show chassis environment routing-engine (TX Matrix Plus Router)

```
user@host> show chassis environment routing-engine
sfc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      26 degrees C / 78 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      28 degrees C / 82 degrees F

lcc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      30 degrees C / 86 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      29 degrees C / 84 degrees F
```

show chassis environment routing-engine (T4000 Core Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  50 degrees C / 122 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  46 degrees C / 114 degrees F
```

show chassis environment routing-engine (QFX Series)

```
user@switch> show chassis environment routing-engine
```

```
Routing Engine 0 status:
  State           Online Master
  Temperature      42 degrees C / 107 degrees F
```

show chassis environment routing-engine interconnect-device (QFabric System)

```
user@switch> show chassis environment routing-engine interconnect-device interconnect1
routing-engine interconnect-device interconnect1
Routing Engine 0 status:
  State           Online Standby
  Temperature      52 degrees C / 125 degrees F
Routing Engine 1 status:
  State           Online Master
  Temperature      57 degrees C / 134 degrees F
```

show chassis environment routing-engine (PTX5000 Packet Transport Switch)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      55 degrees C / 131 degrees F
  CPU Temperature  66 degrees C / 150 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      52 degrees C / 125 degrees F
  CPU Temperature  64 degrees C / 147 degrees F
```

show chassis fan

List of Syntax	Syntax on page 618 Syntax (ACX4000 Series Router) on page 618 Syntax (MX Series Router) on page 618 Syntax (T Series Routers) on page 618 Syntax (MX2010 3D Universal Edge Router) on page 618 Syntax (MX2020 3D Universal Edge Router) on page 618 Syntax (QFabric Systems) on page 618 Syntax (TX Matrix Router) on page 618 Syntax (TX Matrix Plus Router) on page 618
Syntax	show chassis fan
Syntax (ACX4000 Series Router)	show chassis fan
Syntax (MX Series Router)	show chassis fan <all-members> <local> <member <i>member-id</i> >
Syntax (T Series Routers)	show chassis fan
Syntax (MX2010 3D Universal Edge Router)	show chassis fan
Syntax (MX2020 3D Universal Edge Router)	show chassis fan
Syntax (QFabric Systems)	show chassis fan <interconnect-device <i>name</i> >
Syntax (TX Matrix Router)	show chassis fan <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis fan <lcc <i>number</i> sfc <i>number</i> >
Release Information	Command introduced in Junos OS Release 10.0 on MX Series 3D Universal Edge Routers, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus routers. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 12.3 for PTX5000 Packet Transport Switches. Command introduced in Junos OS Release 12.1 for T4000 routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for ACX Series Routers.

Description	(T Series routers, TX Matrix routers, TX Matrix Plus routers, M120 routers, M320 routers, MX2010 routers, MX2020 routers, MX Series 3D Universal Edge Routers, QFX3008-I Interconnect devices, EX Series switches, and PTX Series Packet Transport Switches only) Show information about the fan tray and fans.
Options	<p>all-members—(MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.</p> <p>local—(MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.</p> <p>member <i>member-id</i>—(MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace <i>member-id</i> variable with a value 0 or 1.</p> <p>interconnect-device <i>name</i>—(QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified router (line-card chassis) that is connected to a TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. <p>scc —(TX Matrix routers only) (Optional) Display information about the fan tray and fans for the TX Matrix router (switch-card chassis).</p> <p>sfc <i>number</i>—(TX Matrix Plus routers only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (switch-fabric chassis). Replace <i>number</i> variable with 0.</p>
Required Privilege Level	view
List of Sample Output	<p>show chassis fan on page 620</p> <p>show chassis fan (QFabric Systems) on page 621</p> <p>show chassis fan (EX Series Switches) on page 622</p> <p>show chassis fan (T320 Router) on page 622</p>

[show chassis fan \(T640 Router\) on page 623](#)
[show chassis fan \(T1600 Router\) on page 623](#)
[show chassis fan \(T4000 Core Router\) on page 624](#)
[show chassis fan \(TX Matrix Router\) on page 624](#)
[show chassis fan \(TX Matrix Plus Router\) on page 625](#)
[show chassis fan \(TX Matrix Plus Router with 3D SIBs\) on page 626](#)
[show chassis fan \(PTX5000 Packet Transport Switch\) on page 628](#)
[show chassis fan \(MX2010 Router\) on page 629](#)
[show chassis fan \(MX2020 Router\) on page 629](#)
[show chassis fan \(ACX4000 Router\) on page 630](#)

Output Fields Table 59 on page 620 lists the output fields for the **show chassis fan** command. Output fields are listed in the approximate order in which they appear.

Table 59: show chassis fan Output Fields

Field Name	Field Description
Item	Fan item identifier.
Status	Status of the fan: <ul style="list-style-type: none"> • OK—Fan is running properly and within the normal range. • Check—Fan is in Check state because of some fault or alarm condition.
RPM	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
% RPM	(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Switches only) Percentage of the fan speed being used.
Measurement	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements: <ul style="list-style-type: none"> • Spinning at high speed • Spinning at intermediate speed • Spinning at normal speed • Spinning at low speed (except EX Series switches) (MX2010 routers, MX2020 routers, and PTX Series Packet Transport Switches only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.

Sample Output

show chassis fan

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
------	--------	-----	-------------

Top Tray Fan 1	OK	3790	Spinning at normal speed
Top Tray Fan 2	OK	3769	Spinning at normal speed
Top Tray Fan 3	OK	3769	Spinning at normal speed
Top Tray Fan 4	OK	3790	Spinning at normal speed
Top Tray Fan 5	OK	3790	Spinning at normal speed
Top Tray Fan 6	OK	3769	Spinning at normal speed
Top Tray Fan 7	OK	3790	Spinning at normal speed
Top Tray Fan 8	OK	3769	Spinning at normal speed
Top Tray Fan 9	OK	3769	Spinning at normal speed
Top Tray Fan 10	OK	3790	Spinning at normal speed
Top Tray Fan 11	OK	3790	Spinning at normal speed
Top Tray Fan 12	OK	3769	Spinning at normal speed
Bottom Tray Fan 1	OK	2880	Spinning at normal speed
Bottom Tray Fan 2	OK	2912	Spinning at normal speed
Bottom Tray Fan 3	OK	2928	Spinning at normal speed
Bottom Tray Fan 4	OK	2896	Spinning at normal speed
Bottom Tray Fan 5	OK	2896	Spinning at normal speed
Bottom Tray Fan 6	OK	2928	Spinning at normal speed

show chassis fan (QFabric Systems)

```
user@host> show chassis fan interconnect-device interconnect1
```

Item	Status	RPM	Measurement
TFT 0 Fan 0	OK	2849	Spinning at normal speed
TFT 0 Fan 1	OK	2821	Spinning at normal speed
TFT 0 Fan 2	OK	2735	Spinning at normal speed
TFT 0 Fan 3	OK	2815	Spinning at normal speed
TFT 0 Fan 4	OK	2828	Spinning at normal speed
TFT 0 Fan 5	OK	2863	Spinning at normal speed
BFT 1 Fan 0	OK	2941	Spinning at normal speed
BFT 1 Fan 1	OK	3008	Spinning at normal speed
BFT 1 Fan 2	OK	3073	Spinning at normal speed
BFT 1 Fan 3	OK	2925	Spinning at normal speed
BFT 1 Fan 4	OK	2863	Spinning at normal speed
BFT 1 Fan 5	OK	2933	Spinning at normal speed
SFT 0 Fan 0 Rotor 0	OK	15472	Spinning at normal speed
SFT 0 Fan 0 Rotor 1	OK	14477	Spinning at normal speed
SFT 0 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 0 Fan 1 Rotor 1	OK	14210	Spinning at normal speed
SFT 0 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 0 Fan 2 Rotor 1	OK	14248	Spinning at normal speed
SFT 0 Fan 3 Rotor 0	OK	16463	Spinning at normal speed
SFT 0 Fan 3 Rotor 1	OK	14099	Spinning at normal speed
SFT 1 Fan 0 Rotor 0	OK	15083	Spinning at normal speed
SFT 1 Fan 0 Rotor 1	OK	13533	Spinning at normal speed
SFT 1 Fan 1 Rotor 0	OK	16071	Spinning at normal speed
SFT 1 Fan 1 Rotor 1	OK	14400	Spinning at normal speed
SFT 1 Fan 2 Rotor 0	OK	15517	Spinning at normal speed
SFT 1 Fan 2 Rotor 1	OK	14210	Spinning at normal speed
SFT 1 Fan 3 Rotor 0	OK	16413	Spinning at normal speed
SFT 1 Fan 3 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 2 Fan 0 Rotor 1	OK	14634	Spinning at normal speed
SFT 2 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 2 Fan 1 Rotor 1	OK	14285	Spinning at normal speed
SFT 2 Fan 2 Rotor 0	OK	15835	Spinning at normal speed
SFT 2 Fan 2 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 3 Rotor 0	OK	15789	Spinning at normal speed
SFT 2 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 0 Rotor 0	OK	16314	Spinning at normal speed

SFT 3 Fan 0 Rotor 1	OK	14876	Spinning at normal speed
SFT 3 Fan 1 Rotor 0	OK	15835	Spinning at normal speed
SFT 3 Fan 1 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 2 Rotor 0	OK	16265	Spinning at normal speed
SFT 3 Fan 2 Rotor 1	OK	14594	Spinning at normal speed
SFT 3 Fan 3 Rotor 0	OK	16071	Spinning at normal speed
SFT 3 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 4 Fan 0 Rotor 0	OK	15652	Spinning at normal speed
SFT 4 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 4 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 4 Fan 1 Rotor 1	OK	14555	Spinning at normal speed
SFT 4 Fan 2 Rotor 0	OK	16023	Spinning at normal speed
SFT 4 Fan 2 Rotor 1	OK	14361	Spinning at normal speed
SFT 4 Fan 3 Rotor 0	OK	16216	Spinning at normal speed
SFT 4 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 5 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 5 Fan 0 Rotor 1	OK	14173	Spinning at normal speed
SFT 5 Fan 1 Rotor 0	OK	15472	Spinning at normal speed
SFT 5 Fan 1 Rotor 1	OK	13846	Spinning at normal speed
SFT 5 Fan 2 Rotor 0	OK	15340	Spinning at normal speed
SFT 5 Fan 2 Rotor 1	OK	13917	Spinning at normal speed
SFT 5 Fan 3 Rotor 0	OK	15835	Spinning at normal speed
SFT 5 Fan 3 Rotor 1	OK	13917	Spinning at normal speed
SFT 6 Fan 0 Rotor 0	OK	15743	Spinning at normal speed
SFT 6 Fan 0 Rotor 1	OK	14594	Spinning at normal speed
SFT 6 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 1 Rotor 1	OK	14634	Spinning at normal speed
SFT 6 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 2 Rotor 1	OK	14516	Spinning at normal speed
SFT 6 Fan 3 Rotor 0	OK	16666	Spinning at normal speed
SFT 6 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 0 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 1 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 1 Rotor 1	OK	14361	Spinning at normal speed
SFT 7 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 7 Fan 2 Rotor 1	OK	14555	Spinning at normal speed
SFT 7 Fan 3 Rotor 0	OK	15697	Spinning at normal speed
SFT 7 Fan 3 Rotor 1	OK	14361	Spinning at normal speed

show chassis fan (EX Series Switches)

user@host> show chassis fan

Item	Status	RPM	Measurement
Fan 1	OK	3477	Spinning at normal speed
Fan 2	OK	3477	Spinning at normal speed
Fan 3	OK	3479	Spinning at normal speed
Fan 4	OK	3508	Spinning at normal speed
Fan 5	OK	3517	Spinning at normal speed
Fan 6	OK	3531	Spinning at normal speed
Fan 7	OK	3439	Spinning at normal speed
Fan 8	OK	3424	Spinning at normal speed
Fan 9	OK	3413	Spinning at normal speed
Fan 10	OK	3439	Spinning at normal speed
Fan 11	OK	3446	Spinning at normal speed
Fan 12	OK	3432	Spinning at normal speed

show chassis fan (T320 Router)

user@host> show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	2850	Spinning at normal speed
Top Left Middle fan	OK	2820	Spinning at normal speed
Top Left Rear fan	OK	2970	Spinning at normal speed
Top Right Front fan	OK	2790	Spinning at normal speed
Top Right Middle fan	OK	2640	Spinning at normal speed
Top Right Rear fan	OK	2790	Spinning at normal speed
Bottom Left Front fan	OK	2520	Spinning at normal speed
Bottom Left Middle fan	OK	2610	Spinning at normal speed
Bottom Left Rear fan	OK	2550	Spinning at normal speed
Bottom Right Front fan	OK	2610	Spinning at normal speed
Bottom Right Middle fan	OK	2880	Spinning at normal speed
Bottom Right Rear fan	OK	2790	Spinning at normal speed
Rear Tray Top fan	OK	2130	Spinning at normal speed
Rear Tray Second fan	OK	2190	Spinning at normal speed
Rear Tray Middle fan	OK	2250	Spinning at normal speed
Rear Tray Fourth fan	OK	2220	Spinning at normal speed
Rear Tray Bottom fan	OK	2280	Spinning at normal speed

show chassis fan (T640 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3390	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5220	Spinning at normal speed
Rear Tray Second fan	OK	5220	Spinning at normal speed
Rear Tray Third fan	OK	5220	Spinning at normal speed
Rear Tray Fourth fan	OK	5220	Spinning at normal speed
Rear Tray Fifth fan	OK	5220	Spinning at normal speed
Rear Tray Sixth fan	OK	5220	Spinning at normal speed
Rear Tray Seventh fan	OK	5220	Spinning at normal speed
Rear Tray Bottom fan	OK	5220	Spinning at normal speed

show chassis fan (T1600 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed

Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5190	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	5190	Spinning at normal speed
Rear Tray Sixth fan	OK	5190	Spinning at normal speed
Rear Tray Seventh fan	OK	5190	Spinning at normal speed
Rear Tray Bottom fan	OK	5190	Spinning at normal speed

show chassis fan (T4000 Core Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	5190	Spinning at high speed
Top Left Middle fan	OK	5220	Spinning at high speed
Top Left Rear fan	OK	5190	Spinning at high speed
Top Right Front fan	OK	5160	Spinning at high speed
Top Right Middle fan	OK	5190	Spinning at high speed
Top Right Rear fan	OK	5160	Spinning at high speed
Bottom Left Front fan	OK	6030	Spinning at high speed
Bottom Left Middle fan	OK	6090	Spinning at high speed
Bottom Left Rear fan	OK	6090	Spinning at high speed
Bottom Right Front fan	OK	6030	Spinning at high speed
Bottom Right Middle fan	OK	6060	Spinning at high speed
Bottom Right Rear fan	OK	6060	Spinning at high speed
Rear Tray Top fan	OK	10000	Spinning at high speed
Rear Tray Second fan	OK	10000	Spinning at high speed
Rear Tray Third fan	OK	10000	Spinning at high speed
Rear Tray Fourth fan	OK	10000	Spinning at high speed
Rear Tray Fifth fan	OK	10000	Spinning at high speed
Rear Tray Sixth fan	OK	10000	Spinning at high speed
Rear Tray Seventh fan	OK	10000	Spinning at high speed
Rear Tray Bottom fan	OK	10000	Spinning at high speed

show chassis fan (TX Matrix Router)

```
user@host> show chassis fan
scc-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3390	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3450	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed

```

Rear Tray Seventh fan    OK      3420    Spinning at normal speed
Rear Tray Bottom fan     OK      3420    Spinning at normal speed

```

```
lcc2-re0:
```

```

-----
Item              Status  RPM    Measurement
Top Left Front fan    OK      3420    Spinning at normal speed
Top Left Middle fan   OK      3420    Spinning at normal speed
Top Left Rear fan     OK      3450    Spinning at normal speed
Top Right Front fan   OK      3420    Spinning at normal speed
Top Right Middle fan  OK      3450    Spinning at normal speed
Top Right Rear fan    OK      3360    Spinning at normal speed
Bottom Left Front fan  OK      3420    Spinning at normal speed
Bottom Left Middle fan OK      3480    Spinning at normal speed
Bottom Left Rear fan  OK      3420    Spinning at normal speed
Bottom Right Front fan OK      3420    Spinning at normal speed
Bottom Right Middle fan OK     3390    Spinning at normal speed
Bottom Right Rear fan  OK      3420    Spinning at normal speed
Rear Tray Top fan     OK      3420    Spinning at normal speed
Rear Tray Second fan  OK      3420    Spinning at normal speed
Rear Tray Third fan   OK      3420    Spinning at normal speed
Rear Tray Fourth fan  OK      3420    Spinning at normal speed
Rear Tray Fifth fan   OK      3420    Spinning at normal speed
Rear Tray Sixth fan   OK      3420    Spinning at normal speed
Rear Tray Seventh fan  OK      3420    Spinning at normal speed
Rear Tray Bottom fan  OK      3420    Spinning at normal speed

```

show chassis fan (TX Matrix Plus Router)

```

user@host> show chassis fan
sfc0-re0:

```

```

-----
Item              Status  RPM    Measurement
Fan Tray 0 Fan 1   OK      4350    Spinning at normal speed
Fan Tray 0 Fan 2   OK      4380    Spinning at normal speed
Fan Tray 0 Fan 3   OK      4410    Spinning at normal speed
Fan Tray 0 Fan 4   OK      4380    Spinning at normal speed
Fan Tray 0 Fan 5   OK      4350    Spinning at normal speed
Fan Tray 0 Fan 6   OK      4380    Spinning at normal speed
Fan Tray 1 Fan 1   OK      4410    Spinning at normal speed
Fan Tray 1 Fan 2   OK      4380    Spinning at normal speed
Fan Tray 1 Fan 3   OK      4410    Spinning at normal speed
Fan Tray 1 Fan 4   OK      4380    Spinning at normal speed
Fan Tray 1 Fan 5   OK      4410    Spinning at normal speed
Fan Tray 1 Fan 6   OK      4410    Spinning at normal speed
Fan Tray 2 Fan 1   OK      4380    Spinning at normal speed
Fan Tray 2 Fan 2   OK      4380    Spinning at normal speed
Fan Tray 2 Fan 3   OK      4380    Spinning at normal speed
Fan Tray 2 Fan 4   OK      4410    Spinning at normal speed
Fan Tray 2 Fan 5   OK      4380    Spinning at normal speed
Fan Tray 2 Fan 6   OK      4410    Spinning at normal speed
Fan Tray 2 Fan 7   OK      4410    Spinning at normal speed
Fan Tray 2 Fan 8   OK      4380    Spinning at normal speed
Fan Tray 2 Fan 9   OK      4380    Spinning at normal speed
Fan Tray 3 Fan 1   OK      4350    Spinning at normal speed
Fan Tray 3 Fan 2   OK      4380    Spinning at normal speed
Fan Tray 3 Fan 3   OK      4410    Spinning at normal speed
Fan Tray 3 Fan 4   OK      4440    Spinning at normal speed
Fan Tray 3 Fan 5   OK      4380    Spinning at normal speed
Fan Tray 3 Fan 6   OK      4410    Spinning at normal speed
Fan Tray 3 Fan 7   OK      4410    Spinning at normal speed

```

Fan Tray 3 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 9	OK	4410	Spinning at normal speed

lcc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3450	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3420	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	7050	Spinning at normal speed
Rear Tray Second fan	OK	7050	Spinning at normal speed
Rear Tray Third fan	OK	7050	Spinning at normal speed
Rear Tray Fourth fan	OK	7050	Spinning at normal speed
Rear Tray Fifth fan	OK	7050	Spinning at normal speed
Rear Tray Sixth fan	OK	7050	Spinning at normal speed
Rear Tray Seventh fan	OK	7050	Spinning at normal speed
Rear Tray Bottom fan	OK	7050	Spinning at normal speed

show chassis fan (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4770	Spinning at normal speed

Fan Tray 1 Fan 5	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4770	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 9	Check	2010	

1cc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3390	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed

Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

1cc2-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed
Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

show chassis fan (PTX5000 Packet Transport Switch)

user@host> show chassis fan

user@host> show chassis fan

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	29%	2700 RPM
Fan Tray 0 Fan 2	OK	29%	2700 RPM
Fan Tray 0 Fan 3	OK	29%	2742 RPM
Fan Tray 0 Fan 4	OK	29%	2700 RPM
Fan Tray 0 Fan 5	OK	30%	2828 RPM
Fan Tray 0 Fan 6	OK	30%	2828 RPM
Fan Tray 0 Fan 7	OK	29%	2700 RPM
Fan Tray 0 Fan 8	OK	30%	2785 RPM
Fan Tray 0 Fan 9	OK	30%	2828 RPM
Fan Tray 0 Fan 10	OK	30%	2828 RPM
Fan Tray 0 Fan 11	OK	30%	2785 RPM
Fan Tray 0 Fan 12	OK	30%	2828 RPM
Fan Tray 0 Fan 13	OK	31%	2871 RPM
Fan Tray 0 Fan 14	OK	30%	2828 RPM

Fan Tray 1 Fan 1	OK	42%	3033 RPM
Fan Tray 1 Fan 2	OK	42%	3066 RPM
Fan Tray 1 Fan 3	OK	43%	3099 RPM
Fan Tray 1 Fan 4	OK	43%	3166 RPM
Fan Tray 1 Fan 5	OK	45%	3266 RPM
Fan Tray 1 Fan 6	OK	43%	3133 RPM
Fan Tray 2 Fan 1	OK	29%	2099 RPM
Fan Tray 2 Fan 2	OK	30%	2199 RPM
Fan Tray 2 Fan 3	OK	30%	2166 RPM
Fan Tray 2 Fan 4	OK	33%	2399 RPM
Fan Tray 2 Fan 5	OK	29%	2133 RPM
Fan Tray 2 Fan 6	OK	32%	2366 RPM

show chassis fan (MX2010 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	38%	3480 RPM
Fan Tray 0 Fan 3	OK	37%	3360 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	38%	3480 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	38%	3480 RPM
Fan Tray 1 Fan 2	OK	40%	3600 RPM
Fan Tray 1 Fan 3	OK	38%	3480 RPM
Fan Tray 1 Fan 4	OK	38%	3480 RPM
Fan Tray 1 Fan 5	OK	38%	3480 RPM
Fan Tray 1 Fan 6	OK	38%	3480 RPM
Fan Tray 2 Fan 1	OK	38%	3480 RPM
Fan Tray 2 Fan 2	OK	41%	3720 RPM
Fan Tray 2 Fan 3	OK	38%	3480 RPM
Fan Tray 2 Fan 4	OK	38%	3480 RPM
Fan Tray 2 Fan 5	OK	38%	3480 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	40%	3600 RPM
Fan Tray 3 Fan 3	OK	40%	3600 RPM
Fan Tray 3 Fan 4	OK	40%	3600 RPM
Fan Tray 3 Fan 5	OK	40%	3600 RPM
Fan Tray 3 Fan 6	OK	38%	3480 RPM

show chassis fan (MX2020 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	37%	3360 RPM
Fan Tray 0 Fan 3	OK	36%	3240 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	37%	3360 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	37%	3360 RPM
Fan Tray 1 Fan 2	OK	37%	3360 RPM
Fan Tray 1 Fan 3	OK	37%	3360 RPM
Fan Tray 1 Fan 4	OK	37%	3360 RPM
Fan Tray 1 Fan 5	OK	37%	3360 RPM
Fan Tray 1 Fan 6	OK	36%	3240 RPM
Fan Tray 2 Fan 1	OK	37%	3360 RPM
Fan Tray 2 Fan 2	OK	37%	3360 RPM
Fan Tray 2 Fan 3	OK	37%	3360 RPM

Fan Tray 2 Fan 4	OK	37%	3360 RPM
Fan Tray 2 Fan 5	OK	37%	3360 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	38%	3480 RPM
Fan Tray 3 Fan 3	OK	38%	3480 RPM
Fan Tray 3 Fan 4	OK	37%	3360 RPM
Fan Tray 3 Fan 5	OK	37%	3360 RPM
Fan Tray 3 Fan 6	OK	37%	3360 RPM

show chassis fan (ACX4000 Router)

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	4140	Spinning at normal speed
Fan 2	OK	4200	Spinning at normal speed

show chassis firmware

List of Syntax	Syntax on page 631 Syntax (TX Matrix Routers) on page 631 Syntax (TX Matrix Plus Routers) on page 631 Syntax (MX Series Routers) on page 631 Syntax (MX2010 3D Universal Edge Routers) on page 631 Syntax (MX2020 3D Universal Edge Routers) on page 631 Syntax (QFX Series) on page 631 Syntax (ACX Series Universal Access Routers) on page 631 Syntax (EX Series Switches) on page 631
Syntax	show chassis firmware
Syntax (TX Matrix Routers)	show chassis firmware <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis firmware <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis firmware <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis firmware
Syntax (MX2020 3D Universal Edge Routers)	show chassis firmware
Syntax (QFX Series)	show chassis firmware interconnect-device <i>name</i> node-device <i>name</i>
Syntax (ACX Series Universal Access Routers)	show chassis firmware
Syntax (EX Series Switches)	show chassis firmware <detail>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced for EX8200 switches in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
Command introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.

Description On routers and switches, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), Flexible PIC Concentrators (FPCs), and Routing Engines. On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On EX2200, EX3200, and EX4200 switches, and the QFX Series, display the version levels of the firmware running on the switch. On an EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On an EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

Options **none**—Display the version levels of the firmware running. For an EX4200 switch that is a member of a Virtual Chassis, display version levels for all members. For a TX Matrix router, display version levels for the firmware on the TX Matrix router and on all the T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, display version levels for the firmware on the TX Matrix Plus router and on all the routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Display the version levels of the firmware running for all members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems) (Optional) Display the version levels of the firmware running on the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display version levels for the firmware on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the version levels for the firmware on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the version levels of the firmware running for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the version levels of the firmware running for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device—(QFabric systems only) (Optional) Display the version levels of the firmware running on the Node device.

scc—(TX Matrix router only) (Optional) Display version levels for the firmware on the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display version levels for the firmware on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

detail—(EX3200, EX3300, EX4200, and EX4500 standalone and Virtual Chassis member switches only) (Optional) Display version levels of the firmware running on the switch for its programmable hardware components.

Required Privilege Level view

Related Documentation • *Upgrading the HSM Firmware*

List of Sample Output

- [show chassis firmware \(M10 Router\) on page 634](#)
- [show chassis firmware \(M20 Router\) on page 634](#)
- [show chassis firmware \(M40 Router\) on page 635](#)
- [show chassis firmware \(M120 Router\) on page 635](#)
- [show chassis firmware \(M160 Router\) on page 635](#)
- [show chassis firmware \(MX240 Router\) on page 635](#)
- [show chassis firmware \(MX480 Router\) on page 635](#)
- [show chassis firmware \(MX960 Router\) on page 635](#)
- [show chassis firmware \(MX2010 Router\) on page 636](#)
- [show chassis firmware \(MX2020 Router\) on page 636](#)
- [show chassis firmware \(MX240, MX480, MX960 Router with Application Services Modular Line Card\) on page 636](#)
- [show chassis firmware \(EX4200 Switch\) on page 636](#)
- [show chassis firmware \(EX8200 Switch\) on page 637](#)
- [show chassis firmware lcc \(TX Matrix Router\) on page 637](#)
- [show chassis firmware scc \(TX Matrix Router\) on page 637](#)
- [show chassis firmware \(TX Matrix Plus Router\) on page 637](#)
- [show chassis firmware lcc \(TX Matrix Plus Router\) on page 639](#)
- [show chassis firmware sfc \(TX Matrix Plus Router\) on page 639](#)
- [show chassis firmware \(QFX Series\) on page 640](#)
- [show chassis firmware interconnect-device \(QFabric System\) on page 640](#)
- [show chassis firmware \(ACX2000 Universal Access Router\) on page 640](#)
- [show chassis firmware detail \(EX3300 Switch\) on page 640](#)
- [show chassis firmware \(MX Routers with Media Services Blade \[MSB\]\) on page 640](#)

Output Fields Table 60 on page 634 lists the output fields for the **show chassis firmware** command. Output fields are listed in the approximate order in which they appear.

Table 60: show chassis firmware Output Fields

Field Name	Field Description
Part	(MX Series, MX2010, and MX2020 routers) Chassis part name.
Type	(MX Series, MX2010, and MX2020 routers) Type of firmware: On routers: ROM or O/S . On switches: uboot or loader .
Version	(MX Series, MX2010, and MX2020 routers) Version of firmware running on the chassis part.
FPC	(<i>detail</i> option only) Number of FPC. For a standalone switch, the value is 0. For a Virtual Chassis configuration, value in the range of 0-9; refers to the member ID assigned to the switch.
Boot	(<i>detail</i> option only) Version of the SYSPLD.
PoE	(<i>detail</i> option only) Version of the PoE firmware.
PFE-<number>	(<i>detail</i> option only) Version of the PFE used in the switch.
PHY-	(<i>detail</i> option only) Version of the physical layer device (PHY) used in the switch.
microcode	(<i>detail</i> option only) Microcode of the physical layer devices (PHY) used in the switch.
uboot	(<i>detail</i> option only) Version of the u-boot used in the switch.
loader	(<i>detail</i> option only) Version of the loader used in the switch.

Sample Output

show chassis firmware (M10 Router)

```

user@host> show chassis firmware
Part          Type      Version
Forwarding engine board ROM      Juniper ROM Monitor Version 4.1b2
O/S           Version 4.1I1 by tlim on 2000-04-24 11:27

```

show chassis firmware (M20 Router)

```

user@host> show chassis firmware
Part          Type      Version
System switch board ROM      Juniper ROM Monitor Version 3.4b26
O/S           Version 3.4I16 by smackie on 2000-02-29 2
FPC 1         ROM      Juniper ROM Monitor Version 3.0b1
O/S           Version 3.4I4 by smackie on 2000-02-25 21
FPC 2         ROM      Juniper ROM Monitor Version 3.0b1
O/S           Version 3.4I4 by smackie on 2000-02-25 21

```

show chassis firmware (M40 Router)

```

user@host> show chassis firmware
Part                Type      Version
System control board ROM       Juniper ROM Monitor Version 2.0i126Copyri
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5               ROM       Juniper ROM Monitor Version 2.0i49Copyrig
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:59

```

show chassis firmware (M120 Router)

```

user@host> show chassis firmware
FPC 2               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FEB 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1
FEB 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1

```

show chassis firmware (M160 Router)

```

user@host> show chassis firmware
Part                Type      Version
SFM 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
SFM 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
FPC 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 2               ROM       Juniper ROM Monitor Version 4.0b3
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56

```

show chassis firmware (MX240 Router)

```

user@host> show chassis firmware
Part                Type      Version
FPC 1               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0
FPC 2               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0

```

show chassis firmware (MX480 Router)

```

user@host> show chassis firmware
Part                Type      Version
FPC 1               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20070916.3 by builder on 2007-0

```

show chassis firmware (MX960 Router)

```

user@host> show chassis firmware
Part                Type      Version
FPC 4               ROM       Juniper ROM Monitor Version 8.0b8
                  O/S       Version 8.2I59 by artem on 2006-10-31 19:22

```

FPC 7	ROM	Juniper ROM Monitor Version 8.2b1
	O/S	Version 8.2-20061026.1 by builder on 2006-1

show chassis firmware (MX2010 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 0	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil
FPC 1	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil
FPC 8	ROM	Juniper ROM Monitor Version 10.1b2
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil
FPC 9	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil
SPMB 0	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20120718_ib_12_3_psd.0 by buil

show chassis firmware (MX2020 Router)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 0		
FPC 1		
FPC 2		
FPC 3		
FPC 4		
FPC 5		
FPC 6		
FPC 7		
FPC 8		
FPC 9		
FPC 10		
FPC 11		
FPC 12		
FPC 13		
FPC 14		
FPC 15		
FPC 16		
FPC 17		
FPC 18		
FPC 19		
SPMB 0	ROM	Juniper ROM Monitor Version 11.2b1
	O/S	Version 12.3I5 by psampath on 2012-11-04 03
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3I5 by psampath on 2012-11-04 03

show chassis firmware (MX240, MX480, MX960 Router with Application Services Modular Line Card)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.2I21 by manish on 2012-06-19 17:

show chassis firmware (EX4200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
------	------	---------

FPC 0	uboot loader	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42) FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 1	uboot loader	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42) FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 2	uboot loader	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42) FreeBSD/PowerPC U-Boot bootstrap loader 2.1

show chassis firmware (EX8200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3	U-Boot loader	U-Boot 1.1.6 (Dec 4 2009 - 13:17:34) 3.1.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 5	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7	U-Boot loader	U-Boot 1.1.6 (Feb 6 2009 - 05:31:46) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 1	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2

show chassis firmware lcc (TX Matrix Router)

```
user@host> show chassis firmware lcc 0
lcc0-re0:
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0
FPC 2	ROM	Juniper ROM Monitor Version 6.4b20
	O/S	Version 7.0-20040804.0 by builder on 2004-0
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

show chassis firmware scc (TX Matrix Router)

```
user@host> show chassis firmware scc
scc-re0:
```

Part	Type	Version
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

show chassis firmware (TX Matrix Plus Router)

```
user@host> show chassis firmware
sfc0-re0:
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		

Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		
Global FPC 29		
Global FPC 31		
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

lcc0-re1:

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

lcc1-re1:

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

lcc2-re1:

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 5	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 7.5b4
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

lcc3-re1:

Part	Type	Version
------	------	---------

FPC 0	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 1	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 2	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 4	ROM	Juniper ROM Monitor Version 7.5b4
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 5	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware lcc (TX Matrix Plus Router)

```
user@host> show chassis firmware lcc 0
lcc0-re1:
```

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware sfc (TX Matrix Plus Router)

```
user@host> show chassis firmware sfc 0
sfc0-re0:
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		
Global FPC 29		
Global FPC 31		
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

show chassis firmware (QFX Series)

```
user@switch> show chassis firmware
Part                               Type      Version
FPC 0
Routing Engine 0                  U-Boot    U-Boot 1.1.6 (Sep 15 2010 - 02:11:11) 1.0.5
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware interconnect-device (QFabric System)

```
user@switch> show chassis firmware interconnect-device interconnect1
Part                               Type      Version
Routing Engine 0                  U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1                  U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
loader                            FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware (ACX2000 Universal Access Router)

```
user@switch> show chassis firmware
Part                               Type      Version
FPC                               O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
FEB                               O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
```

show chassis firmware detail (EX3300 Switch)

```
user@switch> show chassis firmware detail
FPC 0
  Boot SYSPLD                      3
  PoE firmware                     4.1.6
  PFE-0                            3
  PFE-1                            3
  PHY
    microcode                      0x514
  Boot Firmware
    uboot                          U-Boot 1.1.6 (Aug 21 2011 - 01:45:26) 1.0.0
    loader                         FreeBSD/arm U-Boot loader 1.0
```

show chassis firmware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis firmware
Part                               Type      Version
FPC 1                             ROM        Juniper ROM Monitor Version 12.1b1
                                O/S        Version 12.2I21 by manish on 2012-06-19 17:
```

show chassis fpc

List of Syntax	Syntax on page 641 Syntax (EX Series Switches) on page 641 Syntax (T4000 Routers) on page 641 Syntax (TX Matrix and TX Matrix Plus Routers) on page 641 Syntax (MX Series Routers) on page 641 Syntax (MX2010 3D Universal Edge Routers) on page 641 Syntax (MX2020 3D Universal Edge Routers) on page 641 Syntax (QFX Series) on page 641 Syntax (PTX Series Packet Transport Switches) on page 641 Syntax (ACX Series Universal Access Routers) on page 642
Syntax	<pre>show chassis fpc <detail <slot>> <pic-status <slot>></pre>
Syntax (EX Series Switches)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <fpc-slot></pre>
Syntax (T4000 Routers)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <slot></pre>
Syntax (MX Series Routers)	<pre>show chassis fpc <detail <slot>> <pic-status <slot>> <all-members> <local> <member member-id></pre>
Syntax (MX2010 3D Universal Edge Routers)	<pre>show chassis fpc <slot> detail <detail <slot>> <pic-status <slot>> <fpc-slot></pre>
Syntax (MX2020 3D Universal Edge Routers)	<pre>show chassis fpc < slot> detail <detail <slot>> <pic-status <slot>> <fpc-slot></pre>
Syntax (QFX Series)	<pre>show chassis fpc <detail> <interconnect-device name <fpc-slot fpc-slot>> <node-device name></pre>
Syntax (PTX Series Packet Transport Switches)	<pre>show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>> <fpc-slot></pre>

Syntax (ACX Series Universal Access Routers) `show chassis fpc`
`<detail <fpc-slot>> | <pic-status <fpc-slot>>`
`<fpc-slot>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for QFX Series.
 Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.
 Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
 Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
 Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

Description Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options **none**—Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached routers in the routing matrix.



NOTE: In EX8200 switches, line cards initialize Packet Forwarding Engine during startup. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error in system log messages (`/var/log/messages`), and generates an alarm to inform the user.

See the following sample output:

```
user@host> show chassis fpc
                               Temp  CPU Utilization (%)  Memory
Utilization (%)
Slot State                    (C)  Total  Interrupt          DRAM (MB) Heap
Buffer
0  Empty
1  Empty
2  Empty
3  Empty
4  Empty
5  Offline                    ---Hard FPC error---
6  Empty
7  Online                     26      4          0          1024      0
32
```

The following sample output shows the alarm raised for the failed FPCs.

```
user@host > show chassis alarms
4 alarms currently active
Alarm time          Class  Description
2011-03-24 00:52:51 UTC Major  FPC 5 Hard errors
2011-03-24 00:52:31 UTC Major  Fan Tray Failure
2011-03-24 00:52:31 UTC Major  Fan Tray Failure
2011-03-24 00:51:26 UTC Minor  Loss of communication with Backup
RE
```



NOTE: On T4000 routers, when you include the enhanced-mode statement at the [edit chassis network-services] hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The show chassis alarm command output displays FPC misconfiguration (FPC *fpc-slot* misconfig) as the reason for the generation the alarms.

The following sample output shows the FPC status after the enhanced-mode statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
```

	Temp	CPU Utilization (%)		Memory	
Utilization (%)					
Slot State	(C)	Total	Interrupt	DRAM (MB)	Heap
Buffer					
0 offline		---FPC misconfiguration---			
1 offline		---FPC misconfiguration---			
2 offline		---FPC misconfiguration---			
3 Empty					
4 Empty					
5 Online	66	50	0	2816	29
27					

The following sample output shows FPC misconfiguration alarms.

```
user@host > show chassis alarms
3 alarms currently active
Alarm time          Class  Description
2011-03-24 00:52:51 PST Major  FPC 1 misconfig
2011-03-24 00:52:31 PST Major  FPC 2 misconfig
2011-03-24 00:52:31 PST Major  FPC 3 misconfig
```

detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see *fpc-slot* or *slot*).

all-members—(MX Series routers only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

fpc-slot—(Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus router only)—On a TX Matrix router, if you specify the number of the T640 router (line-card chassis) by using the **lcc *number*** option (the recommended method), replace *fpc-slot* with a value from 0 through 7. Otherwise, replace *fpc-slot* with a value from 0 through 31. Likewise, on a TX Matrix Plus router, if you specify the number of the specified router (line-card chassis)

by using the **lcc number** option (the recommended method), replace **fpc-slot** with a value from 0 through 7. Otherwise, replace **fpc-slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fpc detail 1 lcc 1
user@host> show chassis fpc detail 9
```

- M120 router—Replace **fpc-slot** with a value from 0 through 5.
- MX80 router—Replace **fpc-slot** with a value from 0 through 1.
- MX240 router—Replace **fpc-slot** with a value from 0 through 2.
- MX480 router—Replace **fpc-slot** with a value from 0 through 5.
- MX-960 router—Replace **fpc-slot** with a value from 0 through 11.
- MX2010 router—Replace **fpc-slot-number** with a value from 0 through 9.
- MX2020 router—Replace **fpc-slot-number** with a value from 0 through 19.
- Other routers—Replace **fpc-slot** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **fpc-slot** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **fpc-slot** with a value from 0 through 9.
 - EX6210 switches—Replace **fpc-slot** with a value from 0 through 9.
 - EX8208 switches—Replace **fpc-slot** with a value from 0 through 7.
 - EX8216 switches—Replace **fpc-slot** with a value from 0 through 15.
- QFX Series:
 - QFX3500 switches—Replace **fpc-slot** with 0.
 - QFabric systems—Replace **fpc-slot** with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Switches:
 - PTX5000 Packet Transport Switch—Replace **fpc-slot** with a value from 0 through 7.
- ACX Series Universal Access Routers:
 - ACX1000 and ACX2000 Universal Access Routers—Replace **fpc-slot** with 0.

local—(MX Series routers only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

node-device name—(QFabric systems only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

pic-status—(Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).



NOTE: On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string “Not Supported” for 10x10GE(LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
Slot 0   Online      E2-FPC Type 1
  PIC 0   Online      1x G/E SFP, 1000 BASE
  PIC 1   Online      Adaptive Services-II
  PIC 2   Online      1x G/E IQ, 1000 BASE
  PIC 3   Online      1x G/E IQ, 1000 BASE
Slot 1   Online      FPC Type 3-ES
  PIC 0   Present     UNUSED- Not Supported
Slot 2   Online      FPC Type 4-ES
  PIC 0   Offline     4x OC-192 SONET XFP
  PIC 1   Present     10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4   Offline     FPC Type 1-ES
Slot 5   Offline     FPC Type 2-ES
Slot 6   Online      E2-FPC Type 3
  PIC 0   Online      1x OC-192 SONET XFP
  PIC 1   Online      4x OC-48 SONET
  PIC 2   Online      4x OC-48 SONET
  PIC 3   Online      MultiServices 500
Slot 7   Online      FPC Type 4-ES
  PIC 0   Online      4x 10GE (LAN/WAN) XFP
  PIC 1   Online      4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (/var/log/messages) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr  5 08:47:36  router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
PIC 1 in FPC 2 (type 763, version 257) is not supported
```

If you see this issue, contact Juniper Networks Technical Assistance Center (JTAC) for a possible fix. For more information about this issue and a possible solution, see [PSN-2010-03-696](https://www.juniper.net/psn/2010-03-696).

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level view

Related Documentation

- [request chassis fpc on page 363](#)
- *show chassis fpc-feb-connectivity*
- *show chassis fabric fpcs*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*
- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*
- *enhanced-mode*

List of Sample Output

- [show chassis fpc \(EX6210 Switch\) on page 649](#)
- [show chassis fpc \(M10 Router\) on page 650](#)
- [show chassis fpc \(M20 Router\) on page 650](#)
- [show chassis fpc detail \(M Series Routers\) on page 650](#)
- [show chassis fpc detail \(MX80 Router\) on page 650](#)
- [show chassis fpc \(MX240 Router\) on page 650](#)
- [show chassis fpc \(MX480 Router\) on page 651](#)
- [show chassis fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 651](#)
- [show chassis fpc pic-status \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 651](#)
- [show chassis fpc \(MX960 Router\) on page 651](#)
- [show chassis fpc \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 652](#)
- [show chassis fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 652](#)
- [show chassis fpc \(MX2010 Routers\) on page 652](#)
- [show chassis fpc \(MX2020 Routers\) on page 652](#)
- [show chassis fpc detail \(MX Series Routers\) on page 653](#)
- [show chassis fpc \(Hardware Not Supported\) on page 653](#)
- [show chassis fpc detail \(Hardware Not Supported\) on page 653](#)
- [show chassis fpc pic-status on page 654](#)
- [show chassis fpc pic-status \(M Series Routers\) on page 654](#)
- [show chassis fpc pic-status \(M120 Router\) on page 654](#)
- [show chassis fpc pic-status \(MX240, MX480, and MX960 Routers with Application Services Modular Line Card\) on page 654](#)
- [show chassis fpc lcc \(TX Matrix Router\) on page 655](#)
- [show chassis fpc pic-status \(TX Matrix Router\) on page 655](#)
- [show chassis fpc pic-status lcc \(TX Matrix Router\) on page 655](#)
- [show chassis fpc \(TX Matrix Plus Router\) on page 656](#)

[show chassis fpc lcc \(TX Matrix Plus Router\) on page 656](#)
[show chassis fpc detail \(TX Matrix Plus Router\) on page 657](#)
[show chassis fpc pic-status \(TX Matrix Plus Router\) on page 659](#)
[show chassis fpc \(T1600 Router\) on page 660](#)
[show chassis fpc detail \(T1600 Router\) on page 660](#)
[show chassis fpc slot \(T1600 Router\) on page 661](#)
[show chassis fpc pic-status \(T1600 Router\) on page 661](#)
[show chassis fpc \(T4000 Router\) on page 661](#)
[show chassis fpc detail \(T4000 Router\) on page 661](#)
[show chassis fpc pic-status \(T4000 Router\) on page 662](#)
[show chassis fpc \(QFX Series\) on page 662](#)
[show chassis fpc detail \(QFX3500 Switches\) on page 662](#)
[show chassis fpc pic-status \(QFX3500 Switches\) on page 662](#)
[show chassis fpc interconnect-device \(QFabric System\) on page 663](#)
[show chassis fpc interconnect-device \(QFabric System\) on page 663](#)
[show chassis fpc interconnect-device detail \(QFabric System\) on page 663](#)
[show chassis fpc pic-status interconnect-device \(QFabric System\) on page 663](#)
[show chassis fpc pic-status node-device \(QFabric System\) on page 664](#)
[show chassis fpc \(PTX5000 Packet Transport Switch\) on page 664](#)
[show chassis fpc detail \(PTX5000 Packet Transport Switch\) on page 664](#)
[show chassis fpc pic-status \(PTX5000 Packet Transport Switch\) on page 665](#)
[show chassis fpc \(ACX2000 Universal Access Router\) on page 665](#)
[show chassis fpc 0 \(ACX2000 Universal Access Router\) on page 666](#)
[show chassis fpc detail \(ACX2000 Universal Access Router\) on page 666](#)
[show chassis fpc pic-status \(ACX2000 Universal Access Router\) on page 666](#)
[show chassis FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 666](#)
[show chassis FPC 1 detail \(MX Routers with Media Services Blade \[MSB\]\) on page 666](#)

Output Fields [Table 61 on page 648](#) lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 61: show chassis fpc Output Fields

Field Name	Field Description	Level of Output
Slot or Slot State	Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> • Dead—Held in reset because of errors. • Diag—Slot is being ignored while the FPC is running diagnostics. • Dormant—Held in reset. • Empty—No FPC is present. • Offline—(PTX Series Packet Transport Switches only) One of the following two states is displayed: <ul style="list-style-type: none"> • FPC offlined due to unreachable destinations • FPC Offlined due to degraded FPC action • Online—FPC is online and running. • Present—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either Hardware Not Supported or Hardware Not In Right Slot. The FPC is coming up but not yet online. • Probed—Probe is complete; awaiting restart of the Packet Forwarding Engine. • Probe-wait—Waiting to be probed. 	all levels
Logical slot	Slot number.	all levels
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels all levels
Temperature (PTX Series)	On PTX Series Packet Transport Switches, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: <ul style="list-style-type: none"> • Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC. • Temperature (Intake)—Temperature of the air flowing into the chassis. • Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B). • Temperature (TLn)—Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC. • Temperature (TQn)—Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC. 	detail
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.	all levels
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.	none specified

Table 61: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak). <i>NOTE:</i> On MX Series routers in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.	none specified
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
Total CPU DRAM	Amount of DRAM available to the FPC's CPU.	detail
Total RLDRAM	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	detail
Total DDR DRAM	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	detail
Total SRAM	Amount of static RAM (SRAM) used by the FPC's CPU.	detail
Total SDRAM	Total amount of memory used for storing packets and notifications.	detail
I/O Manager ASICs information	I/O Manager version number, manufacturer, and part number.	detail
Start time	Time when the Routing Engine detected that the FPC was running.	detail
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	detail
PIC type	(pic-status output only) Type of PIC.	none specified

Sample Output

show chassis fpc (EX6210 Switch)

```

user@switch> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Online	7	5 0	1024	0 32
2	Empty				
3	Empty				
4	Online	25	17 2	2048	0 30
5	Online	25	3 0	2048	0 24
6	Online	6	5 0	1024	0 32
7	Empty				
8	Empty				
9	Online	8	7 0	1024	0 32

show chassis fpc (M10 Router)

```

user@host> show chassis fpc
FPC status:

Slot State      Temp
              (C)
0  Online       27
1  Online       28

```

show chassis fpc (M20 Router)

```

user@host> show chassis fpc
FPC status:

Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB)  Heap      Buffer
0  Empty         0      0      0      0      0      0
1  Online        38      0      0      8      0      4
2  Online        35      0      0      8      0      3
3  Empty         0      0      0      0      0      0

```

show chassis fpc detail (M Series Routers)

```

user@host> show chassis fpc detail 1
Slot 1 information:
State                               Online
Temperature                         48 degrees C
Total CPU DRAM                     32 MB
Total SRAM                         4 MB
Total SDRAM                        256 MB
I/O Manager ASICs information      Version 2.0, Foundry IBM, Part number 0
I/O Manager ASICs information      Version 2.0, Foundry IBM, Part number 0
Start time                         2000-02-08 02:18:49 UTC
Uptime                             14 hours, 41 minutes, 41 seconds

```

show chassis fpc detail (MX80 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
State                               Online
Temperature                         47 degrees C / 116 degrees F
Total CPU DRAM                     1024 MB
Total SRAM                         331 MB
Total SDRAM                        1280 MB
Start time                         2010-02-08 12:25:33 PST
Uptime                             2 hours, 13 minutes, 19 seconds
Slot 1 information:
State                               Online
Temperature                         47 degrees C / 116 degrees F
Total CPU DRAM                     1024 MB
Total SRAM                         331 MB
Total SDRAM                        1280 MB
Start time                         2010-02-08 12:25:33 PST
Uptime                             2 hours, 13 minutes, 19 seconds

```

show chassis fpc (MX240 Router)

```

user@host> show chassis fpc

Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB)  Heap      Buffer
0  Empty

```

1	Online	34	6	0	1024	18	30
2	Online	33	9	0	1024	24	30

show chassis fpc (MX480 Router)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Online	36	9	0	1024	17	57
2	Empty						
3	Empty						
4	Empty						
5	Empty						

show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Online	33	4	0	2048	10	13
1	Online	36	7	0	2048	16	13
2	Online	29	6	0	1024	27	29
3	Online	33	0	0	0	0	0
4	Online	36	7	0	2048	19	13
5	Online	34	31	11	2048	14	13

show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis fpc pic-status
```

Slot 1	Online	MPC Type 3
PIC 2	Online	1X100GE CFP
Slot 2	Online	DPCE 40x 1GE R EQ
PIC 0	Online	10x 1GE(LAN) EQ
PIC 1	Online	10x 1GE(LAN) EQ
PIC 2	Online	10x 1GE(LAN) EQ
PIC 3	Online	10x 1GE(LAN) EQ
Slot 3	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 4	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 5	Online	MPC Type 2 3D EQ
PIC 0	Online	2x 10GE XFP
PIC 1	Online	2x 10GE XFP
PIC 2	Online	10x 1GE(LAN) SFP
PIC 3	Online	10x 1GE(LAN) SFP

show chassis fpc (MX960 Router)

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Empty						
2	Empty						
3	Online	25	19	0	1024	15	57
4	Empty						
5	Online	26	27	0	1024	15	57
6	Empty						
7	Empty						

```

8 Empty
9 Empty
10 Empty
11 Empty

```

show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host>show chassis fpc 1
      Temp CPU Utilization (%)  Memory    Utilization (%)
Slot State      (C) Total  Interrupt    DRAM (MB) Heap      Buffer
1 Online        34     5      0      3072     5      13

```

show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis fpc 1 detail
Slot 1 information:
State Online
Temperature 34
Total CPU DRAM 3072 MB
Total RLD RAM 259 MB
Total DDR DRAM 4864 MB
Start time: 2012-06-19 10:51:43 PDT
Uptime: 16 minutes, 48 seconds
Max Power Consumption 550 Watts

```

show chassis fpc (MX2010 Routers)

```

user@host show chassis fpc
      Temp CPU Utilization (%)  Memory    Utilization (%)
Slot State      (C) Total  Interrupt    DRAM (MB) Heap      Buffer
0 Online        34     9      0      2048     18      13
1 Online        32     9      0      2048     15      13
2 Empty
3 Empty
4 Empty
5 Empty
6 Empty
7 Empty
8 Online        31    13      0      2048     11      13
9 Online        33    10      0      2048     18      13

```

show chassis fpc (MX2020 Routers)

```

user@host show chassis fpc
      Temp CPU Utilization (%)  Memory    Utilization (%)
Slot State      (C) Total  Interrupt    DRAM (MB) Heap      Buffer
0 Online        10    12      0      2048     18      13
1 Online         8     9      0      2048     18      13
2 Online         7     9      0      2048     18      13
3 Online         8    10      0      2048     18      13
4 Online         9    10      0      2048     18      13
5 Online         8     9      0      2048     18      13
6 Online         8    10      0      2048     18      13
7 Online         9     9      0      2048     18      13
8 Online         9    10      0      2048     18      13
9 Online        10     9      0      2048     18      13
10 Online        16     8      0      2048     18      13
11 Online        11    10      0      2048     18      13
12 Online        10    10      0      2048     18      13
13 Online        11     9      0      2048     18      13
14 Online        12    10      0      2048     18      13
15 Online        13     9      0      2048     18      13

```


16	Online	13	9	0	2048	18	13
17	Online	12	9	0	2048	18	13
18	Online	12	8	0	2048	18	13
19	Online	14	10	0	2048	18	13

show chassis fpc detail (MX Series Routers)

```

user@host> show chassis fpc detail 2
Slot 0 information:
  State                Online
  Temperature          36 degrees C / 96 degrees F
  Total CPU DRAM       1024 MB
  Total RLDRAM         256 MB
  Total DDR DRAM       4096 MB
  Start time:          2009-08-11 21:20:30 PDT
  Uptime:              2 hours, 8 minutes, 50 seconds
  Max Power Consumption 335 Watts

```

show chassis fpc (Hardware Not Supported)

```

user@host> show chassis fpc
show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online						
1	Present						
2	Online		0	0	0	0	0
3	Present						
4	Empty						
5	Empty						
6	Online		0	0	0	0	0

show chassis fpc detail (Hardware Not Supported)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                Online
  Total CPU DRAM       ---- CPU less FPC ----
  Start time           2006-07-07 03:21:00 UTC
  Uptime               27 minutes, 51 seconds
Slot 1 information:
  State                Present
  Reason               --- Hardware Not In Right Slot ---
Slot 2 information:
  State                Online
  Total CPU DRAM       32 MB
  Start time           2006-07-07 03:20:59 UTC
  Uptime               27 minutes, 52 seconds
Slot 3 information:
  State                Present
  Reason               --- Hardware Not Supported ---
  Total CPU DRAM       0 MB
Slot 6 information:
  State                Online
  Total CPU DRAM       32 MB
  Start time           2006-07-07 03:21:01 UTC
  Uptime               27 minutes, 50 seconds

```

show chassis fpc pic-status

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1    1x OC-12 ATM, MM
  PIC 2    1x OC-12 ATM, MM
  PIC 3    1x OC-12 ATM, MM
Slot 1 Online
  PIC 0    1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0    1x OC-192 SONET, SMSR
```

show chassis fpc pic-status (M Series Routers)

```
user@host> show chassis fpc pic-status
Slot 1  Online      FPC Type 1
  PIC 0  Present    2x OC-3 ATM, MM- Hardware Error
  PIC 1  Online     4x OC-3 SONET, SMIR
Slot 2  Online      E-FPC Type 2
  PIC 0  Online     4x G/E, 1000 BASE-SX
  PIC 1  Online     2x G/E SFP, 1000 BASE
  PIC 3  Online     1x Tunnel
Slot 3  Online      E-FPC Type 1
  PIC 0  Online     1x G/E IQ, 1000 BASE
  PIC 2  Online     1x G/E SFP, 1000 BASE
Slot 4  Online      E-FPC Type 2
  PIC 0  Online     4x G/E SFP, 1000 BASE
  PIC 1  Online     4x G/E SFP, 1000 BASE
  PIC 2  Online     4x G/E SFP, 1000 BASE
  PIC 3  Online     4x G/E SFP, 1000 BASE
Slot 5  Online      FPC Type 2
...
```

show chassis fpc pic-status (M120 Router)

```
user@host> show chassis fpc pic-status
Slot 1  Online      M120 CFPC 10GE
  PIC 0  Online     1x 10GE(LAN/WAN) XFP
Slot 3  Online      M120 FPC Type 2 (proto)
  PIC 0  Online     2x G/E IQ, 1000 BASE
  PIC 1  Online     4x OC-3 SONET, SMIR
  PIC 2  Online     2x G/E IQ, 1000 BASE
  PIC 3  Online     8x 1GE(LAN), IQ2
Slot 4  Online      M120 FPC Type 3 (proto)
  PIC 0  Online     10x 1GE(LAN), 1000 BASE
Slot 5  Online      M120 FPC Type 1 (proto)
  PIC 0  Present    1x G/E, 1000 BASE-LX- Not Supported
  PIC 1  Online     1x CHOC3 IQ SONET, SMLR
  PIC 2  Online     4x CHDS3 IQ
  PIC 3  Online     1x G/E SFP, 1000 BASE
```

show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1** and **Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```
user@host> show chassis fpc pic-status
Slot 2  Online      MPC Type 1 3D Q
Slot 1  Online      AS-MCC
```

```

PIC 0 Online AS-MSC
PIC 2 Online AS-MXC
Slot 4 Offline MPC 3D 16x 10GE
Slot 5 Offline AS-MCC

```

show chassis fpc lcc (TX Matrix Router)

```

user@host> show chassis fpc lcc 0
lcc0-re0:

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory Utilization (%) DRAM (MB) Heap Buffer
0	Empty				
1	Online	27	2	0	256 8 44
2	Online	27	3	0	256 15 44
3	Empty				
4	Empty				
5	Empty				
6	Empty				
7	Empty				

show chassis fpc pic-status (TX Matrix Router)

```

user@host> show chassis fpc pic-status
lcc0-re0:

```

Slot 0	Online	FPC Type 3
PIC 0	Online	1x OC-192 SM SR1
PIC 1	Online	1x OC-192 SM SR2
PIC 2	Online	1x OC-192 SM SR1
PIC 3	Online	1x Tunnel
Slot 1	Online	FPC Type 2
PIC 0	Online	1x OC-48 SONET, SMSR
PIC 1	Online	1x OC-48 SONET, SMSR

```
lcc1-re0:
```

```
lcc2-re0:
```

Slot 1	Online	FPC Type 3
PIC 0	Online	1x OC-192 SM SR1
Slot 5	Online	FPC Type 2
PIC 0	Online	1x OC-48 SONET, SMSR
PIC 1	Online	2x G/E, 1000 BASE-LX
PIC 2	Online	2x G/E, 1000 BASE-LX
PIC 3	Online	1x OC-48 SONET, SMSR

```
lcc3-re0:
```

show chassis fpc pic-status lcc (TX Matrix Router)

```

user@host> show chassis fpc pic-status lcc 0
lcc0-re0:

```

Slot 0	Online	FPC Type 3
PIC 0	Online	1x OC-192 SM SR2
Slot 1	Online	FPC Type 2
PIC 0	Online	2x OC-12 ATM2 IQ, MM
PIC 1	Online	1x OC-48 SONET, SMSR

PIC 2 Online 1x OC-48 SONET, SMSR
 PIC 3 Online 4x G/E, 1000 BASE-SX

show chassis fpc (TX Matrix Plus Router)

user@host> show chassis fpc
 lcc0-re0:

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Online	38	4	0	2048	3	24
2	Online	43	8	0	2048	6	24
3	Empty						
4	Online	43	6	0	2048	6	24
5	Empty						
6	Online	42	13	0	2048	6	24
7	Online	45	7	0	2048	3	24

lcc2-re0:

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Online	42	10	0	2048	6	24
1	Empty						
2	Online	42	11	0	2048	6	24
3	Online	40	5	0	2048	3	24
4	Online	33	26	0	1024	8	49
5	Empty						
6	Online	43	8	0	2048	6	24
7	Online	46	6	0	2048	3	24

lcc3-re0:

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Empty						
2	Online	39	30	0	2048	7	24
3	Empty						
4	Online	41	8	0	2048	6	24
5	Online	41	12	0	2048	6	24
6	Online	40	8	0	2048	6	24
7	Online	42	4	0	2048	3	24

show chassis fpc lcc (TX Matrix Plus Router)

user@host> show chassis fpc lcc 0
 lcc0-re0:

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Online	38	4	0	2048	3	24
2	Online	43	8	0	2048	6	24
3	Empty						
4	Online	43	6	0	2048	6	24
5	Empty						
6	Online	42	14	0	2048	6	24
7	Online	45	6	0	2048	3	24

show chassis fpc detail (TX Matrix Plus Router)

```
user@host> show chassis fpc details
```

```
lcc0-re0:
```

```
-----
```

```
Slot 1 information:
```

```
State                Online
Temperature           38 degrees C / 100 degrees F
Total CPU DRAM        2048 MB
Total SRAM            64 MB
Total SDRAM           1280 MB
Start time            2010-10-04 20:06:22 PDT
Uptime                1 hour, 32 minutes, 51 seconds
```

```
Slot 2 information:
```

```
State                Online
Temperature           43 degrees C / 109 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:06:37 PDT
Uptime                1 hour, 32 minutes, 36 seconds
```

```
Slot 4 information:
```

```
State                Online
Temperature           43 degrees C / 109 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:06:40 PDT
Uptime                1 hour, 32 minutes, 33 seconds
```

```
Slot 6 information:
```

```
State                Online
Temperature           42 degrees C / 107 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:06:42 PDT
Uptime                1 hour, 32 minutes, 31 seconds
```

```
Slot 7 information:
```

```
State                Online
Temperature           45 degrees C / 113 degrees F
Total CPU DRAM        2048 MB
Total SRAM            64 MB
Total SDRAM           1280 MB
Start time            2010-10-04 20:06:43 PDT
Uptime                1 hour, 32 minutes, 30 seconds
```

```
lcc2-re0:
```

```
-----
```

```
Slot 0 information:
```

```
State                Online
Temperature           42 degrees C / 107 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:06:35 PDT
Uptime                1 hour, 32 minutes, 38 seconds
```

```
Slot 2 information:
```

```
State                Online
Temperature           42 degrees C / 107 degrees F
Total CPU DRAM        2048 MB
```

```

Total SRAM                128 MB
Total SDRAM                2560 MB
Start time                2010-10-04 20:06:37 PDT
Uptime                    1 hour, 32 minutes, 36 seconds
Slot 3 information:
State                     Online
Temperature               40 degrees C / 104 degrees F
Total CPU DRAM            2048 MB
Total SRAM                64 MB
Total SDRAM               1280 MB
Start time                2010-10-04 20:06:28 PDT
Uptime                    1 hour, 32 minutes, 45 seconds
Slot 4 information:
State                     Online
Temperature               33 degrees C / 91 degrees F
Total CPU DRAM            1024 MB
Total SRAM                64 MB
Total SDRAM               1280 MB
Start time                2010-10-04 20:08:03 PDT
Uptime                    1 hour, 31 minutes, 10 seconds
Slot 6 information:
State                     Online
Temperature               43 degrees C / 109 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM               2560 MB
Start time                2010-10-04 20:06:44 PDT
Uptime                    1 hour, 32 minutes, 29 seconds
Slot 7 information:
State                     Online
Temperature               46 degrees C / 114 degrees F
Total CPU DRAM            2048 MB
Total SRAM                64 MB
Total SDRAM               1280 MB
Start time                2010-10-04 20:06:46 PDT
Uptime                    1 hour, 32 minutes, 27 seconds

```

lcc3-re0:

```

-----
Slot 2 information:
State                     Online
Temperature               38 degrees C / 100 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM               2560 MB
Start time                2010-10-04 20:17:31 PDT
Uptime                    1 hour, 21 minutes, 42 seconds
Slot 4 information:
State                     Online
Temperature               41 degrees C / 105 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM               2560 MB
Start time                2010-10-04 20:17:34 PDT
Uptime                    1 hour, 21 minutes, 39 seconds
Slot 5 information:
State                     Online
Temperature               41 degrees C / 105 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM               2560 MB

```

```

Start time                2010-10-04 20:17:36 PDT
Uptime                    1 hour, 21 minutes, 37 seconds
Slot 6 information:
State                     Online
Temperature               40 degrees C / 104 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM               2560 MB
Start time                2010-10-04 20:17:39 PDT
Uptime                    1 hour, 21 minutes, 34 seconds
Slot 7 information:
State                     Online
Temperature               42 degrees C / 107 degrees F
Total CPU DRAM            2048 MB
Total SRAM                64 MB
Total SDRAM               1280 MB
Start time                2010-10-04 20:17:41 PDT
Uptime                    1 hour, 21 minutes, 32 seconds

```

show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status
```

```
1cc0-re0:
```

```

-----
Slot 1  Online      FPC Type 2-ES
PIC 0   Online      8x 1GE(LAN), IQ2
Slot 2  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
Slot 4  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
Slot 6  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
PIC 1   Online      4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
PIC 0   Online      10x 1GE(LAN), 1000 BASE
PIC 2   Online      1x OC-192 SM SR2
PIC 3   Online      10x 1GE(LAN), 1000 BASE

```

```
1cc2-re0:
```

```

-----
Slot 0  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
Slot 2  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
PIC 1   Online      4x 10GE (LAN/WAN) XFP
Slot 3  Online      FPC Type 2-ES
PIC 0   Online      8x 1GE(LAN), IQ2
Slot 4  Online      FPC Type 4
PIC 0   Online      10x10GE(LAN/WAN) SFPP
Slot 6  Online      FPC Type 4-ES
PIC 0   Online      4x OC-192 SONET XFP
Slot 7  Online      FPC Type 3-ES
PIC 0   Online      10x 1GE(LAN), 1000 BASE
PIC 1   Offline     1x 10GE(LAN/WAN) IQ2E
PIC 2   Online      1x OC-192 SM SR2
PIC 3   Online      1x Tunnel

```

```
1cc3-re0:
```

```

-----
Slot 2  Online      FPC Type 4-ES

```

```

PIC 0 Online 10x10GE(LAN/WAN) SFPP
Slot 4 Online FPC Type 4-ES
PIC 0 Online 4x OC-192 SONET XFP
Slot 5 Online FPC Type 4-ES
PIC 0 Online 4x OC-192 SONET XFP
PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 6 Online FPC Type 4-ES
PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 7 Online FPC Type 3-ES
PIC 0 Online 10x 1GE(LAN), 1000 BASE
PIC 1 Online 8x 1GE(TYPE3), IQ2E
PIC 2 Online 4x OC-48 SONET

```

show chassis fpc (T1600 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Empty				
2	Online	49	3 0	2048	3 24
3	Online	46	6 0	2048	6 24
4	Empty				
5	Online	46	5 0	2048	3 24
6	Empty				
7	Online	44	8 0	1024	7 49

show chassis fpc detail (T1600 Router)

```

user@host> show chassis fpc detail

```

```

show chassis fpc detail
Slot 2 information:
  State Online
  Temperature 49 degrees C / 120 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM 64 MB
  Total SDRAM 1280 MB
  Start time 2010-10-04 21:12:52 PDT
  Uptime 32 minutes, 9 seconds
Slot 3 information:
  State Online
  Temperature 47 degrees C / 116 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM 128 MB
  Total SDRAM 2560 MB
  Start time 2010-10-04 21:13:06 PDT
  Uptime 31 minutes, 55 seconds
Slot 5 information:
  State Online
  Temperature 46 degrees C / 114 degrees F
  Total CPU DRAM 2048 MB
  Total SRAM 64 MB
  Total SDRAM 1280 MB
  Start time 2010-10-04 21:12:56 PDT
  Uptime 32 minutes, 5 seconds
Slot 7 information:
  State Online
  Temperature 44 degrees C / 111 degrees F
  Total CPU DRAM 1024 MB
  Total SRAM 64 MB

```



```

Total SDRAM                1280 MB
Start time                  2010-10-04 21:14:34 PDT
Uptime                      30 minutes, 27 seconds

```

show chassis fpc slot (T1600 Router)

```
user@host> show chassis fpc slot 2
```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
2	Online	49	3	0	2048	3	24

show chassis fpc pic-status (T1600 Router)

```
user@host> show chassis fpc pic-status
```

```

Slot 2  Online      FPC Type 1-ES
        PIC 0  Online  Load Type 1
        PIC 1  Online  4x 1GE(LAN), IQ2E
        PIC 3  Online  1x OC-12-3 SFP
Slot 3  Online      FPC Type 4-ES
        PIC 0  Online  4x 10GE (LAN/WAN) XFP
        PIC 1  Online  4x OC-192 SONET XFP
Slot 5  Online      FPC Type 2-ES
        PIC 0  Online  Load Type 2
        PIC 1  Online  8x 1GE(LAN), IQ2E
        PIC 2  Online  8x 1GE(LAN), IQ2E
        PIC 3  Online  1x OC-48-12-3 SFP
Slot 7  Online      FPC Type 4
        PIC 0  Online  4x 10GE (LAN/WAN) XFP

```

show chassis fpc (T4000 Router)

```
user@host> show chassis fpc
```

```

regress@stymphalian# run show chassis fpc

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Online	48	15	0	2816	21	27
1	Empty						
2	Empty						
3	Online	51	15	0	2816	21	27
4	Empty						
5	Online	39	8	0	2048	6	23
6	Online	49	15	0	2816	21	27
7	Empty						

show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
```

```
Slot 0 information:
```

```

State                Online
Temperature          48 degrees C / 118 degrees F
Total CPU DRAM       2816 MB
Total SRAM           1554 MB
Total SDRAM          10752 MB
Start time           2012-02-09 22:56:25 PST
Uptime               2 hours, 40 minutes, 52 seconds

```

```
Slot 3 information:
```

```

State                Online
Temperature          51 degrees C / 123 degrees F
Total CPU DRAM       2816 MB

```

```

Total SRAM                1554 MB
Total SDRAM                10752 MB
Start time                2012-02-09 22:56:22 PST
Uptime                    2 hours, 40 minutes, 55 seconds
Slot 5 information:
State                     Online
Temperature                39 degrees C / 102 degrees F
Total CPU DRAM            2048 MB
Total SRAM                128 MB
Total SDRAM                2560 MB
Start time                2012-02-09 22:51:27 PST
Uptime                    2 hours, 45 minutes, 50 seconds
Slot 6 information:
State                     Online
Temperature                49 degrees C / 120 degrees F
Total CPU DRAM            2816 MB
Total SRAM                1554 MB
Total SDRAM                10752 MB
Start time                2012-02-09 22:56:29 PST
Uptime                    2 hours, 40 minutes, 48 seconds

```

show chassis fpc pic-status (T4000 Router)

```

user@host> show chassis fpc pic-status
Slot 0  Online      FPC Type 5-3D
PIC 0   Online      12x10GE (LAN/WAN) SFPP
PIC 1   Online      12x10GE (LAN/WAN) SFPP
Slot 3  Online      FPC Type 5-3D
PIC 0   Online      1x100GE
PIC 1   Online      12x10GE (LAN/WAN) SFPP
Slot 5  Online      FPC Type 4-ES
PIC 0   Online      100GE
PIC 1   Online      100GE CFP
Slot 6  Online      FPC Type 5-3D
PIC 0   Online      12x10GE (LAN/WAN) SFPP
PIC 1   Online      12x10GE (LAN/WAN) SFPP

```

show chassis fpc (QFX Series)

```

user@switch> show chassis fpc
Temp CPU Utilization (%) Memory      Utilization (%)
Slot State              (C) Total Interrupt    DRAM (MB) Heap      Buffer
0  Online                26      2           0      2820      0      49

```

show chassis fpc detail (QFX3500 Switches)

```

user@switch> show chassis fpc detail
Slot 0 information:
State                     Online
Temperature                28 degrees C / 82 degrees F
Total CPU DRAM            2820 MB
Total SRAM                0 MB
Total SDRAM                0 MB
Start time                2010-09-20 01:34:13 PDT
Uptime                    3 days, 3 hours, 31 minutes, 48 seconds

```

show chassis fpc pic-status (QFX3500 Switches)

```

user@switch> show chassis fpc pic-status
Slot 0  Online      QFX 48x10G 4x40G Switch
PIC 0   Online      48x 10G-SFP+
PIC 1   Online      15x 10G-SFP+

```

show chassis fpc interconnect-device (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1
FPC status:

```

Slot	State	Temp (C)
0	Online	0
1	Online	0
2	Online	0
3	Online	0
4	Online	0
5	Online	0
6	Online	0
7	Online	0
8	Online	0
9	Online	0
10	Online	0
11	Online	0
12	Online	0
13	Online	0
14	Online	0
15	Online	0

show chassis fpc interconnect-device (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3
FPC status:

```

Slot	State	Temp (C)
3	Online	0

show chassis fpc interconnect-device detail (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3 detail
Slot 3 information:

```

State	Online
Temperature	0 degrees C / 32 degrees F
Start time	2011-08-18 10:45:04 PDT
Uptime	1 minute, 49 seconds

show chassis fpc pic-status interconnect-device (QFabric System)

```

user@switch> show chassis fpc pic-status interconnect-device interconnect1

```

Slot 0	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 1	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 2	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 3	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 4	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE
Slot 5	Online	QFX 16-port QSFP+ Front Card
PIC 0	Online	16x 40G-QSFP+
PIC 1	Online	16x 40G-GE

```

Slot 6  Online      QFX 16-port QSFP+ Front Card
PIC 0   Online      16x 40G-QSFP+
PIC 1   Online      16x 40G-GE
Slot 7  Online      QFX 16-port QSFP+ Front Card
PIC 0   Online      16x 40G-QSFP+
PIC 1   Online      16x 40G-GE
Slot 8  Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 9  Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 10 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 11 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 12 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 13 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 14 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE
Slot 15 Online      QFX Fabric Rear Card
PIC 0   Online      16x 40G-GE

```

show chassis fpc pic-status node-device (QFabric System)

```

user@switch> show chassis fpc pic-status node-device node1
Slot node1 Online      QFX 48x10G 4x40G Switch
PIC 0   Online      48x 10G-SFP+
PIC 1   Online      4x 40G-QSFP+

```

show chassis fpc (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Empty				
1	Empty				
2	Online	50	6 0	2816	5 27
3	Empty				
4	Empty				
5	Online	48	9 0	2816	5 27
6	Empty				
7	Online	49	8 0	2816	5 27

show chassis fpc detail (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 35 degrees C / 95 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 54 degrees C / 129 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 58 degrees C / 136 degrees F (TQ1)
Temperature 57 degrees C / 134 degrees F (TL2)
Temperature 58 degrees C / 136 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 61 degrees C / 141 degrees F (TQ3)

```

```

Total CPU DRAM          2816 MB
Total SRAM              0 MB
Total SDRAM            0 MB
Start time              2012-01-12 12:05:42 PST
Uptime                 3 hours, 14 minutes, 7 seconds

Slot 5 information:
State                  Online
Temperature            35 degrees C / 95 degrees F (PMB)
Temperature            34 degrees C / 93 degrees F (Intake)
Temperature            48 degrees C / 118 degrees F (Exhaust A)
Temperature            53 degrees C / 127 degrees F (Exhaust B)
Temperature            54 degrees C / 129 degrees F (TL0)
Temperature            52 degrees C / 125 degrees F (TQ0)
Temperature            69 degrees C / 156 degrees F (TL1)
Temperature            56 degrees C / 132 degrees F (TQ1)
Temperature            54 degrees C / 129 degrees F (TL2)
Temperature            56 degrees C / 132 degrees F (TQ2)
Temperature            59 degrees C / 138 degrees F (TL3)
Temperature            60 degrees C / 140 degrees F (TQ3)
Total CPU DRAM          2816 MB
Total SRAM              0 MB
Total SDRAM            0 MB
Start time              2012-01-12 12:05:43 PST
Uptime                 3 hours, 14 minutes, 6 seconds

Slot 7 information:
State                  Online
Temperature            35 degrees C / 95 degrees F (PMB)
Temperature            33 degrees C / 91 degrees F (Intake)
Temperature            50 degrees C / 122 degrees F (Exhaust A)
Temperature            55 degrees C / 131 degrees F (Exhaust B)
Temperature            56 degrees C / 132 degrees F (TL0)
Temperature            56 degrees C / 132 degrees F (TQ0)
Temperature            61 degrees C / 141 degrees F (TL1)
Temperature            57 degrees C / 134 degrees F (TQ1)
Temperature            55 degrees C / 131 degrees F (TL2)
Temperature            59 degrees C / 138 degrees F (TQ2)
Temperature            62 degrees C / 143 degrees F (TL3)
Temperature            62 degrees C / 143 degrees F (TQ3)
Total CPU DRAM          2816 MB
Total SRAM              0 MB
Total SDRAM            0 MB
Start time              2012-01-12 12:05:44 PST
Uptime                 3 hours, 14 minutes, 5 seconds

```

show chassis fpc pic-status (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc pic-status
Slot 2  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      24x 10GE(LAN) SFP+
Slot 5  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      2x 40GE CFP
Slot 7  Online      FPC
  PIC 0  Online      24x 10GE(LAN) SFP+
  PIC 1  Online      2x 40GE CFP

```

show chassis fpc (ACX2000 Universal Access Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Online	61	17 6	512	21 37

show chassis fpc 0 (ACX2000 Universal Access Router)

```
user@host> show chassis fpc 0
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
0	Online	61	17 6	512	21 37

show chassis fpc detail (ACX2000 Universal Access Router)

```
user@host> show chassis fpc detail
```

Slot 0 information:

State	Online
Temperature	61 degrees C / 141 degrees F
Total CPU DRAM	512 MB
Start time	2012-05-29 02:52:06 PDT
Uptime	27 minutes, 17 seconds

show chassis fpc pic-status (ACX2000 Universal Access Router)

```
user@host> show chassis fpc pic-status
```

Slot 0	Online	
PIC 0	Online	16x CHE1T1, RJ48
PIC 1	Online	8x 1GE(LAN) RJ45
PIC 2	Online	2x 1GE(LAN) SFP
PIC 3	Online	2x 10GE(LAN) SFP+

show chassis FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
1	Online	34	5 0	3072	5 13

show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1 detail
```

Slot 1 information:

State	Online
Temperature	34
Total CPU DRAM	3072 MB
Total RLDRAM	259 MB
Total DDR DRAM	4864 MB
Start time:	2012-06-19 10:51:43 PDT
Uptime:	16 minutes, 48 seconds
Max Power Consumption	550 Watts

show chassis hardware

List of Syntax	Syntax on page 667 Syntax (EX Series) on page 667 Syntax (T4000 Router) on page 667 Syntax (TX Matrix Router) on page 667 Syntax (TX Matrix Plus Router) on page 667 Syntax (MX Series Routers) on page 667 Syntax (MX2010 3D Universal Edge Routers) on page 667 Syntax (MX2020 3D Universal Edge Routers) on page 668 Syntax (QFX Series) on page 668 Syntax (PTX Series Packet Transport Switches) on page 668 Syntax (ACX Series Universal Access Routers) on page 668
Syntax	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (EX Series)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (T4000 Router)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (TX Matrix Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis hardware <detail extensive> <clei-models> <models> <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis hardware <clei-models> <detail extensive>

	<models>
Syntax (MX2020 3D Universal Edge Routers)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (QFX Series)	show chassis hardware <detail extensive> <clei-models> <interconnect-device <i>name</i> > <node-device <i>name</i> > <models>
Syntax (PTX Series Packet Transport Switches)	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (ACX Series Universal Access Routers)	show chassis hardware <detail extensive> <clei-models> <models>
Release Information	Command introduced before Junos OS Release 7.4. models option introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number. In the EX Series switch command output, FPC refers to the following: <ul style="list-style-type: none">• On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC <i>number</i> is always 0.• On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC <i>number</i> equals the member ID, from 0 through 9.• On EX8208 and EX8216 switches—Refers to a line card; FPC <i>number</i> equals the slot number for the line card. On a QFX3500 standalone switch, both the FPC and FPC <i>number</i> are always 0. On Type 5 FPC on T4000 routers, there are no top temperature sensor or bottom temperature sensor parameters. Instead, fan intake temperature sensor and fan exhaust temperature sensors parameters are displayed.

- Options** **none**—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached routers.
- clei-models**—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).
- detail**—(Optional) Include RAM and disk information in output.
- extensive**—(Optional) Display ID EEPROM information.
- all-members**—(MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.
- interconnect-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.
- lcc *number***—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.
- Replace *number* with the following values depending on the LCC configuration:
- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
 - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
 - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
 - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- local**—(MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.
- member *member-id***—(MX Series routers only) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.
- models**—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.
- node-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.
- scc**—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

Additional Information The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines:

Table 62: Routing Engines Displaying DIMM Information

Routing Engines	Routers
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 routers
RE-A-1800x2	M120 and M320 routers

Required Privilege Level view

Related Documentation

- [show chassis power](#)

List of Sample Output

- [show chassis hardware \(EX8216 Switch\) on page 675](#)
- [show chassis hardware clei-models \(EX8216 Switch\) on page 676](#)
- [show chassis hardware clei-models \(T1600 Router\) on page 677](#)
- [show chassis hardware detail \(EX4200 Switch\) on page 678](#)
- [show chassis hardware models \(EX4500 Switch\) on page 678](#)
- [show chassis hardware \(J6350 Router\) on page 678](#)
- [show chassis hardware \(J6300 Router\) on page 678](#)
- [show chassis hardware \(M7i Router\) on page 679](#)
- [show chassis hardware \(M10 Router\) on page 679](#)
- [show chassis hardware models \(M10 Router\) on page 680](#)
- [show chassis hardware \(M20 Router\) on page 680](#)
- [show chassis hardware models \(M20 Router\) on page 681](#)
- [show chassis hardware \(M40 Router\) on page 681](#)
- [show chassis hardware \(M40e Router\) on page 682](#)
- [show chassis hardware \(M120 Router\) on page 682](#)
- [show chassis hardware detail \(M120 Router\) on page 683](#)
- [show chassis hardware models \(M120 Router\) on page 684](#)
- [show chassis hardware \(M160 Router\) on page 685](#)
- [show chassis hardware models \(M160 Router\) on page 685](#)
- [show chassis hardware detail \(M160 Router\) on page 686](#)
- [show chassis hardware \(M320 Router\) on page 687](#)
- [show chassis hardware models \(M320 Router\) on page 688](#)
- [show chassis hardware \(MX5 Router\) on page 688](#)
- [show chassis hardware \(MX10 Router\) on page 689](#)
- [show chassis hardware \(MX40 Router\) on page 689](#)
- [show chassis hardware \(Fixed MX80 Router\) on page 690](#)
- [show chassis hardware \(Modular MX80 Router\) on page 691](#)
- [show chassis hardware \(MX240 Router\) on page 691](#)
- [show chassis hardware detail \(MX 240 Router with Routing Engine Displaying DIMM information\) on page 692](#)

[show chassis hardware \(MX240 Router with Enhanced MX SCB\) on page 692](#)
[show chassis hardware \(MX480 Router\) on page 693](#)
[show chassis hardware \(MX480 Router with Enhanced MX SCB\) on page 694](#)
[show chassis hardware \(MX960 Router\) on page 694](#)
[show chassis hardware \(MX960 Router with Bidirectional Optics\) on page 694](#)
[show chassis hardware \(MX960 Router with Enhanced MX SCB\) on page 695](#)
[show chassis hardware models \(MX960 Router with Enhanced MX SCB\) on page 697](#)
[show chassis hardware detail \(MX960 Router\) on page 698](#)
[show chassis hardware \(MX2010 Router\) on page 698](#)
[show chassis hardware detail \(MX2010 Router\) on page 700](#)
[show chassis hardware extensive \(MX2010 Router\) on page 705](#)
[show chassis hardware models \(MX2010 Router\) on page 710](#)
[show chassis hardware clei-models \(MX2010 Routers\) on page 711](#)
[show chassis hardware \(MX2020 Router\) on page 712](#)
[show chassis hardware detail \(MX2020 Router\) on page 720](#)
[show chassis hardware models \(MX2020 Router\) on page 729](#)
[show chassis hardware clei-models \(MX2020 Router\) on page 731](#)
[show chassis hardware \(MX Series routers with ATM MIC\) on page 732](#)
[show chassis hardware \(MX240, MX480, MX960 routers with Application Services Modular Line Card\) on page 732](#)
[show chassis hardware extensive \(MX240, MX480, MX960 routers with Application Services Modular Line Card\) on page 733](#)
[show chassis hardware \(T320 Router\) on page 734](#)
[show chassis hardware \(T640 Router\) on page 735](#)
[show chassis hardware models \(T640 Router\) on page 736](#)
[show chassis hardware extensive \(T640 Router\) on page 736](#)
[show chassis hardware \(T4000 Router\) on page 737](#)
[show chassis hardware \(T4000 Router with 16 GB line card chassis \(LCC\) Routing Engine\) on page 739](#)
[show chassis hardware \(T4000 Router with LSR FPC\) on page 740](#)
[show chassis hardware clei-models \(T4000 Router\) on page 740](#)
[show chassis hardware detail \(T4000 Router\) on page 740](#)
[show chassis hardware models \(T4000 Router\) on page 742](#)
[show chassis hardware lcc \(TX Matrix Router\) on page 743](#)
[show chassis hardware scc \(TX Matrix Router\) on page 743](#)
[show chassis hardware \(T1600 Router\) on page 744](#)
[show chassis hardware \(TX Matrix Plus Router\) on page 746](#)
[show chassis hardware sfc \(TX Matrix Plus Router\) on page 751](#)
[show chassis hardware extensive \(TX Matrix Plus Router\) on page 752](#)
[show chassis hardware clei-models \(TX Matrix Plus Router\) on page 753](#)
[show chassis hardware detail \(TX Matrix Plus Router\) on page 756](#)
[show chassis hardware models \(TX Matrix Plus Router\) on page 758](#)
[show chassis hardware \(TX Matrix Plus router with 3D SIBs\) on page 760](#)
[show chassis hardware clei-models \(TX Matrix Plus router with 3D SIBs\) on page 764](#)
[show chassis hardware detail \(TX Matrix Plus router with 3D SIBs\) on page 768](#)
[show chassis hardware lcc \(TX Matrix Plus router with 3D SIBs\) on page 771](#)
[show chassis hardware sfc \(TX Matrix Plus router with 3D SIBs\) on page 772](#)
[show chassis hardware \(16-Port 10-Gigabit Ethernet MPC with SFP+ Optics \[MX Series Routers\]\) on page 773](#)

[show chassis hardware \(MPC3E \[MX Series Routers\]\) on page 774](#)
[show chassis hardware \(QFX3500 Switches\) on page 775](#)
[show chassis hardware detail \(QFX3500 Switches\) on page 775](#)
[show chassis hardware models \(QFX3500 Switches\) on page 777](#)
[show chassis hardware clei-models \(QFX3500 Switches\) on page 777](#)
[show chassis hardware interconnect-device \(QFabric Systems\) on page 777](#)
[show chassis hardware node-device \(QFabric Systems\) on page 777](#)
[show chassis hardware \(PTX5000 Packet Transport Switch\) on page 777](#)
[show chassis hardware clei-models \(PTX5000 Packet Transport Switch\) on page 779](#)
[show chassis hardware detail \(PTX5000 Packet Transport Switch\) on page 779](#)
[show chassis hardware models \(PTX5000 Packet Transport Switch\) on page 781](#)
[show chassis hardware extensive \(PTX5000 Packet Transport Switch\) on page 781](#)
[show chassis hardware \(MX Routers with Media Services Blade \[MSB\]\) on page 782](#)
[show chassis hardware extensive \(MX Routers with Media Services Blade \[MSB\]\) on page 782](#)

Output Fields [Table 63 on page 673](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 63: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
Item	<p>Chassis component:</p> <ul style="list-style-type: none"> (EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>. (MX Series routers)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). (M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller. (M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB. (QFX Series)—Information about the chassis, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). (PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal). (MX2010 and MX2020 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays. 	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels
Serial number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels
Assb ID or Assembly ID	(extensive keyword only) Identification number that describes the FRU hardware.	extensive

Table 63: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
Assembly Version	(extensive keyword only) Version number of the FRU hardware.	extensive
Assembly Flags	(extensive keyword only) Flags.	extensive
FRU model number	(clei-models , extensive , and models keyword only) Model number of the FRU hardware component.	none specified
CLEI code	(clei-models and extensive keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
EEPROM Version	ID EEPROM version used by the hardware component: 0x00 (version 0), 0x01 (version 1), or 0x02 (version 2).	extensive
Description	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> Type of power supply. Type of PIC. If the PIC type is not supported on the current software release, the output states Hardware Not Supported. Type of FPC: FPC Type 1, FPC Type 2, FPC Type 3, FPC Type 4, or FPC TypeOC192. <p>On EX Series switches, a brief description of the FPC.</p> <p>On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> 2x FE—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM 4x FE—4-port Fast Ethernet ePIM 1x GE Copper—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port) 1x GE SFP—SFP Gigabit Ethernet ePIM (one fiber port) 4x GE Base PIC—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM) 2x Serial—Dual-port serial PIM 2x T1—Dual-port T1 PIM 2x E1—Dual-port E1 PIM 2x CTIE1—Dual-port channelized T1/E1 PIM 1x T3—T3 PIM (one port) 1x E3—E3 PIM (one port) 4x BRI S/T—4-port ISDN BRI S/T PIM 4x BRI U—4-port ISDN BRI U PIM 1x ADSL Annex A—ADSL 2/2+ Annex A PIM (one port, for POTS) 1x ADSL Annex B—ADSL 2/2+ Annex B PIM (one port, for ISDN) 2x SHDSL (ATM)—G SHDSL PIM (2-port two-wire module or 1-port four-wire module) 	All levels

Table 63: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • 1x TGM550—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog LINE ports, and two analog TRUNK ports) • 1x DS1 TIM510—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup) • 4x FXS, 4x FXO, TIM514—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog LINE ports and four analog TRUNK ports) • 4x BRI TIM521—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports) • Crypto Accelerator Module—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services • MPCM16x10GE—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: LX, SX, LH, or T. • LCD description for EX Series switches (except EX2200 switches). • MPC2—1-port MPC2 that supports two separate slots for MICs. • MPC3E—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs. • 100GBASE-LR4, pluggable CFP optics • Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy. • Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs). • LCD description for MX Series routers 	

Sample Output

show chassis hardware (EX8216 Switch)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 06                CY0109220035  EX8216
Midplane      REV 06    710-016845    BA0909120112  EX8216-MP
CB 0          REV 22    710-020771    AX0109197723  EX8216-RE320
CB 1          REV 22    710-020771    AX0109197726  EX8216-RE320
Routing Engine 1    BUILTIN    BUILTIN      RE-EX8216
FPC 3         REV 19    710-020683    BC0109083125  EX8200-48F
CPU           REV 13    710-020598    BF0109144549  EX8200-CPU
FPC 4         REV 17    710-020683    BC0108500127  EX8200-48F
CPU           REV 10    710-020598    BF0108460510  EX8200-CPU
PIC 0                     BUILTIN      BUILTIN      48x 100 Base-QFX/1000
Base-X
Xcvr 1        REV 01    740-011613    PE70V89       SFP-SX

```

Xcvr 11	REV 01	740-011613	PE70YCE	SFP-SX
Xcvr 12	REV 01	740-011613	PE70VSH	SFP-SX
Xcvr 13	REV 01	740-011613	E08C02063	SFP-SX
Xcvr 14	REV 01	740-011613	PE70VKU	SFP-SX
Xcvr 15	REV 01	740-011613	E08E03372	SFP-SX
Xcvr 21	REV 01	740-011613	PE70VAD	SFP-SX
Xcvr 22	REV 01	740-011613	E08E01228	SFP-SX
Xcvr 23	REV 01	740-011613	PE70VSL	SFP-SX
Xcvr 24	REV 01	740-011613	E08E03409	SFP-SX
Xcvr 25	REV 01	740-011613	PE70VL4	SFP-SX
Xcvr 26	REV 01	740-011613	PDQ4L2Z	SFP-SX
Xcvr 27	REV 01	740-011613	PE70WFK	SFP-SX
Xcvr 28	REV 01	740-011782	PBD2B5U	SFP-SX
Xcvr 29	REV 01	740-011613	PE70UQX	SFP-SX
Xcvr 30	REV 01	740-011613	PE70VL5	SFP-SX
Xcvr 31	REV 01	740-011613	PE70V0F	SFP-SX
Xcvr 32	REV 01	740-011613	E08C02052	SFP-SX
Xcvr 33	REV 01	740-011613	E08C02197	SFP-SX
Xcvr 34	REV 01	740-011613	PE70V0L	SFP-SX
Xcvr 35	REV 01	740-011613	E08E03390	SFP-SX
Xcvr 36	REV 01	740-011613	PDQ4VL9	SFP-SX
Xcvr 37	REV 01	740-011613	E08E03370	SFP-SX
Xcvr 38	REV 01	740-011613	E08E03362	SFP-SX
Xcvr 39	REV 01	740-011613	E08C02065	SFP-SX
Xcvr 40	REV 01	740-011613	E08E03405	SFP-SX
Xcvr 41	REV 01	740-011613	E08E03411	SFP-SX
Xcvr 43	REV 01	740-011613	E08C02171	SFP-SX
Xcvr 45	REV 01	740-011613	E08E03410	SFP-SX
FPC 13	REV 16	710-016837	BB0109051344	EX8200-8XS
CPU				
SIB 0	REV 10	710-021613	AY0109166244	EX8216-SF320
SIB 1	REV 10	710-021613	AY0109166357	EX8216-SF320
SIB 2	REV 10	710-021613	AY0109166362	EX8216-SF320
SIB 3	REV 10	710-021613	AY0109166338	EX8216-SF320
SIB 4	REV 10	710-021613	AY0109166350	EX8216-SF320
SIB 5	REV 10	710-021613	AY0109166365	EX8216-SF320
SIB 6	REV 10	710-021613	AY0109166361	EX8216-SF320
SIB 7	REV 10	710-021613	AY0109166399	EX8216-SF320
PSU 0	REV 17	740-021466	BG0709170003	EX8200-AC2K
PSU 1	REV 17	740-021466	BG0709170004	EX8200-AC2K
PSU 2	REV 17	740-021466	BG0709170020	EX8200-AC2K
PSU 3	REV 17	740-021466	BG0709170017	EX8200-AC2K
PSU 4	REV 17	740-021466	BG0709170008	EX8200-AC2K
PSU 5	REV 17	740-021466	BG0709170018	EX8200-AC2K
Top Fan Tray				
FTC 0	REV 4	760-022620	CX1209140212	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140212	EX8216-FT
Bottom Fan Tray				
FTC 0	REV 4	760-022620	CX1209140211	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140211	EX8216-FT
LCD 0	REV 04	710-025742	CE0109186919	EX8200 LCD

show chassis hardware clei-models (EX8216 Switch)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Midplane     REV 08   710-016845
PSU 0        REV 05   740-023002  COUPAEAEAA     EX8200-PWR-AC3KR
PSU 1        REV 05   740-023002  COUPAEAEAA     EX8200-PWR-AC3KR
PSU 2        REV 05   740-023002  COUPAEAEAA     EX8200-PWR-AC3KR

```



```

PSU 3          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 4          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 5          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray

```

show chassis hardware clei-models (T1600 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-005608		CHAS-BP-T640-S
FPM Display	REV 05	710-002897		CRAFT-T640-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 07	740-017906	IPUPAC7KTA	PWR-T1600-3-80-DC-S
PEM 1	Rev 18	740-002595		PWR-T-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 08	740-014082		RE-A-2000-4096-S
Routing Engine 1	REV 07	740-014082		RE-A-2000-4096-S
CB 0	REV 05	710-007655		CB-T-S
CB 1	REV 03	710-017707		CB-T-S
FPC 0	REV 07	710-013558		T640-FPC2-E2
PIC 0	REV 01	750-010618		PB-4GE-SFP
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 14	750-001901		PB-40C12-SON-SMIR
PIC 3	REV 07	750-001900		PB-10C48-SON-SMSR
FPC 1	REV 06	710-013553		T640-FPC1-E2
PIC 0	REV 08	750-001072		P-1GE-SX
PIC 1	REV 10	750-012266		PB-4GE-TYPE1-SFP-IQ2
PIC 2	REV 22	750-005634		PB-1CHOC12SMIR-QPP
FPC 2				
PIC 0	REV 16	750-007141		PC-10GE-SFP
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 05	750-004695		PC-TUNNEL
PIC 3	REV 17	750-009553		PC-40C48-SON-SFP
FPC 3	REV 01	710-010154		T640-FPC3-E
PIC 0	REV 07	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 25	750-007141		PC-10GE-SFP
PIC 2	REV 17	750-009553		PC-40C48-SON-SFP
PIC 3	REV 32	750-003700		PC-10C192-SON-VSR
FPC 4	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP
FPC 5	REV 02	710-013037		T1600-FPC4-ES
PIC 0	REV 16	750-012518		PD-40C192-SON-XFP
PIC 1	REV 01	750-010850		PD-10C768-SON-SR
FPC 6	REV 14	710-013037		T1600-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
PIC 1	REV 13	750-017405		PD-4XGE-XFP
FPC 7	REV 09	710-007529		T640-FPC3
PIC 0	REV 10	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 3	REV 15	750-009450		PC-10C192-SON-SR2
SIB 0	REV 07	710-013074		SIB-I-T1600-S
SIB 1	REV 07	710-013074		SIB-I-T1600-S
SIB 2	REV 07	710-013074		SIB-I-T1600-S
SIB 3	REV 07	710-013074		SIB-I-T1600-S
SIB 4	REV 07	710-013074		SIB-I-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FAN-REAR-TX-T640-S

show chassis hardware detail (EX4200 Switch)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               BM0208327733  EX4200-24T
Routing Engine 0 REV 11    750-021256  BM0208327733  EX4200-24T, 8 POE
Routing Engine 0                               BM0208327733  EX4200-24T, 8 POE
FPC 0          REV 11    750-021256  BM0208327733  EX4200-24T, 8 POE
  CPU                               BUILTIN       BUILTIN       FPC CPU
  PIC 0                               BUILTIN       BUILTIN       24x 10/100/1000 Base-T
  PIC 1          REV 03B  711-021270  AR0208162285  4x GE SFP
  BRD            REV 08    711-021264  AK0208328289  EX4200-24T, 8 POE
Power Supply 0  REV 03    740-020957  AT0508346354  PS 320W AC
Fan Tray

```

show chassis hardware models (EX4500 Switch)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Routing Engine 0 REV 01    750-035700  GG0210271867  EX4500-40F-FB-C
FPC 0          REV 01    750-035700  GG0210271867  EX4500-40F-FB-C
  PIC 0                               BUILTIN       BUILTIN       EX4500-40F-FB-C
Power Supply 1  REV 01    740-029654  H884FS00JC09  EX4500-PWR1-AC-FB

```

show chassis hardware (J6350 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1090E07ADB  JSR6350
Midplane          REV 03    710-014593  NP1265
System IO          REV 01    710-016210  NN9950        JX350 System IO
Crypto Module                               RE-J6350-3400  Crypto Acceleration
Routing Engine    REV 08    710-015273  NM6509        RE-J6350-3400
  ad0             248 MB   256MB   CKS           00102006C24A00000039 Compact
Flash
FPC 0
  PIC 0
FPC 1          REV 06    750-010355  AI07030023    FPC
  PIC 0
FPC 3          REV 06    750-011148  AJ06520151    FPC
  PIC 0
FPC 6          REV 06    750-013492  NC4170        FPC
  PIC 0
Power Supply 0

```

show chassis hardware (J6300 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN000164AB    J6300
Midplane          REV 02.04 710-010001  CORE99570
System IO          REV 02.00 710-010003  CORE100848    System IO board
Routing Engine    RevX2.6  750-010006  IWGS40735390  RE-J.3
FPC 0
  PIC 0
FPC 1          RevX2.0  750-011380  N3960005      FPC
  PIC 0

```

FPC 2	RevX2.0	750-011380	N3960002	FPC
PIC 0				1xADSL pic Annex B
FPC 3	REV 03	750-010354	N0780028	FPC
PIC 0				1x T3

show chassis hardware (M7i Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			31959	M7i
Midplane	REV 02	710-008761	CA0209	M7i Midplane
Power Supply 0	Rev 04	740-008537	PD10272	AC Power Supply
Routing Engine	REV 01	740-008846	1000396803	RE-5.0
CFEB	REV 02	750-009492	CA0166	Internet Processor IIv1
FPC 0				E-FPC
PIC 0	REV 04	750-003163	HJ6416	1x G/E, 1000 BASE-SX
PIC 1	REV 04	750-003163	HJ6423	1x G/E, 1000 BASE-SX
PIC 2	REV 04	750-003163	HJ6421	1x G/E, 1000 BASE-SX
PIC 3	REV 02	750-003163	HJ0425	1x G/E, 1000 BASE-SX
FPC 1				E-FPC
PIC 2	REV 01	750-009487	HM2275	ASP - Integrated
PIC 3	REV 01	750-009098	CA0142	2x F/E, 100 BASE-TX

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			B1157	M7i
Midplane	REV 05	710-008761	DM0840	M7i Midplane
Power Supply 0	Rev 08	740-008537	TE53755	AC Power Supply
Routing Engine	REV 07	740-011202	1000736567	RE-850
CFEB	REV 09	750-010463	DK6952	Internet Processor II
FPC 0				E-FPC
PIC 0	REV 12	750-012838	DL7993	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011614	PD94TDJ	SFP-LX10
Xcvr 1	REV 01	740-011615	PAD5EER	UNSUPPORTED
Xcvr 2	REV 01	740-011614	PD94THU	SFP-LX10
Xcvr 3		NON-JNPR	PDC2E7A	SFP-LX10
PIC 1	REV 03	750-023116	JT0203	4x CHSTM1 SDH CE SFP
Xcvr 0	REV 01	740-012434	AGT063832PS	SFP-SR
Xcvr 1	REV 01	740-012434	AGT063832LY	SFP-SR
Xcvr 3	REV 01	740-016064	C06J19018	SFP-LR
PIC 2	REV 15	750-014895	DM5757	MultiServices 100
PIC 3	REV 01	750-025390	JW9448	12x T1/E1 CE
FPC 1				E-FPC
PIC 2		BUILTIN	BUILTIN	1x Tunnel
PIC 3	REV 09	750-009099	DM0899	1x G/E, 1000 BASE
Xcvr 0	REV 01	740-012434	AGT07150HGJ	UNSUPPORTED
Fan Tray				Rear Fan Tray

show chassis hardware (M10 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			1122	M10
Midplane	REV 1.1	710-001950	S/N AC6626	
Power supply A	Rev 01	740-002497	S/N LC36095	AC
Power supply B	Rev 01	740-002497	S/N LC36100	AC
Display	REV 1.2	710-001995	S/N AC6656	
Host			18000005dfb3fb01	teknon
FEB	REV 01	710-001948	S/N AC6632	Internet Processor II
FPC 0				

PIC 0	REV 08	750-001072	S/N AB2485	1x G/E, 1000 BASE-SX
PIC 1	REV 01	750-000613	S/N AA1048	1x OC-12 SONET, SMIR
FPC 1				
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware models (M10 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-008920		CHAS-MP-M10i-S
Power Supply 0	Rev 06	740-008537		PWR-M10i-M7i-AC-S
Power Supply 1	Rev 06	740-008537		PWR-M10i-M7i-AC-S
HCM 0	REV 03	710-010580		HCM-M10i-S
HCM 1	REV 03	710-010580		HCM-M10i-S
Routing Engine 0	REV 09	740-009459		RE-400-256-S
CFEB 0	REV 05	750-010465		FEB-M10i-M7i-S
FPC 0				
PIC 0	REV 10	750-002971		PE-40C3-SON-MM
PIC 1	REV 11	750-002992		PE-4FE-TX
PIC 2	REV 03	750-002977		PE-20C3-ATM-MM
PIC 3	REV 08	750-005724		PE-20C3-ATM2-MM
FPC 1				
PIC 2	REV 12	750-008425		PE-AS
PIC 3	REV 13	750-005636		PE-4CHDS3-QPP
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware (M20 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			20033	M20
Backplane	REV 07	710-001517	S/N AA7940	
Power supply B	Rev 01	740-001465	S/N 000001	AC
Display	REV 02	710-001519	S/N AA9704	
Host 0			98000004f8f27501	teknor
SSB slot 0	REV 01	710-001951	S/N AD5905	Internet Processor II
SSRAM bank 0	REV 01	710-001385	S00480	2 MB
SSRAM bank 1	REV 01	710-001385	S00490	2 MB
SSRAM bank 2	REV 01	710-001385	S001:?	2 MB
SSRAM bank 3	REV 01	710-001385	S00483	2 MB
SSB slot 1	N/A	N/A	N/A	Backup
FPC 1	REV 01	710-001292	S/N AB7528	
SSRAM	REV 01	710-000077	S/N 304209	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 000603	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 000414	64 MB
PIC 0	REV 03	750-000612	S/N AB8433	2x OC-3 ATM, MM
PIC 1	REV 01	750-000616	S/N AA1168	1x OC-12 ATM, MM
PIC 2	REV 01	750-000613	S/N AA1008	1x OC-12 SONET, SMIR
PIC 3	REV 01	750-002501	S/N AD5810	4x E3
FPC 2	REV 01	710-001292	S/N AC0119	
SSRAM	REV 01	710-000077	S/N 503241	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 306835	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 306832	64 MB
Fan Tray 0				Front Upper Fan Tray
Fan Tray 1				Front Middle Fan Tray
Fan Tray 2				Front Bottom Fan Tray
Fan Tray 3				Rear Fan Tray

show chassis hardware models (M20 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Backplane     REV 03   710-002334
Power Supply A REV 06   740-001465
Display       REV 04   710-001519
Routing Engine 0 REV 06   740-003239
Routing Engine 1 REV 06   740-003239
SSB 0         REV 02   710-001951
SSB 1         N/A      N/A
FPC 0         REV 03   710-003308
  PIC 0       REV 08   750-002303
  PIC 1       REV 07   750-004745
  PIC 2       REV 03   750-002965
FPC 1         REV 03   710-003308
  PIC 0       REV 03   750-002914
Fan Tray 0
Fan Tray 1
Fan Tray 2
Fan Tray 3
CHAS-MP-M20-S
PWR-M20-AC-S
CRAFT-M20-S
RE-333-768-S
RE-333-768-S
SSB-E-M20
FPC-E
P-4FE-TX
P-2MCD53
PE-4CHDS3
FPC-E
P-20C3-ATM-MM
FANTRAY-F-M20-S
FANTRAY-F-M20-S
FANTRAY-F-M20-S
FANTRAY-R-M20-S

```

show chassis hardware (M40 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Backplane     REV 02   710-000073   S/N AA0053
Power supply A Rev 2     740-000235   S/N 000042    DC
Maxicab       REV X1   710-000229   S/N AA0139
Minicab       REV X1   710-000482   S/N AA0201
Display       REV 06   710-000150   S/N AA0905
Host
SCB
  SSRAM bank 0 REV 02   710-000077   S/N AA2267    1 MB
  SSRAM bank 1 REV 02   710-000077   S/N AA2270    1 MB
  SSRAM bank 2 REV 02   710-000077   S/N AA2269    1 MB
  SSRAM bank 3 REV 02   710-000077   S/N AA2268    1 MB
FPC 0
  SSRAM       REV 01   710-000175   S/N AA0048    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2332    64 MB
  SDRAM bank 1 REV X1   710-000099   S/N AA2337    64 MB
  PIC 0       REV 04   750-000613   S/N aa0343    1x OC-12 SONET, SMIR
  PIC 1       REV 04   750-000613   S/N AA0379    1x OC-12 SONET, SMIR
  PIC 2       REV 04   750-000613   S/N AA0377    1x OC-12 SONET, SMIR
  PIC 3       REV 04   750-000613   S/N AA0378    1x Tunnel
FPC 2
  SSRAM       REV 01   710-000175   S/N AA0042    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2331    64 MB
  SDRAM bank 1 REV 01   710-000099   S/N AA2330    64 MB
  PIC 0       REV X1   750-000603   S/N AA0143    4x OC-3 SONET, SMIR
  PIC 1       REV X1   750-000615   S/N AA0149    4x OC-3 SONET, MM
  PIC 2       REV X1   750-000611   S/N AA0148    4x OC-3 SONET, MM
  PIC 3       REV 04   750-000613   S/N AA0330    1x OC-12 SONET, SMIR
FPC 4
  SSRAM       REV 01   710-000175   S/N AA0050    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2329    64 MB
  SDRAM bank 1 REV 01   710-000099   S/N AA2328    64 MB
  PIC 0       REV 04   750-000613   S/N AA0320    1x OC-12 SONET, SMIR
  PIC 2       REV 05   750-000616   S/N AA1341    1x OC-12 ATM, MM

```

PIC 3	REV 08	750-001072	S/N AB2462	1x G/E, 1000 BASE-SX
FPC 5	REV 10	710-000175	S/N AA7663	
SSRAM	REV 01	710-000077	S/N 501590	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 300949	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 300868	64 MB
PIC 1	REV 01	750-001323	S/N AB1670	1x Tunnel

show chassis hardware (M40e Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis				m40e
Midplane	REV 01	710-005071	AX3671	
FPM CMB	REV 03	710-001642	AR9074	
FPM Display	REV 03	710-001647	AR7331	
CIP	REV 04	710-002649	BB4449	
PEM 0	Rev 01	740-003787	MC12364	Power Entry Module
PEM 1	Rev 01	740-003787	MC12383	Power Entry Module
PCG 0	REV 07	710-001568	AG1332	
PCG 1	REV 07	710-001568	AR3789	
Host 0			3e000007c8176601	Present
MCS 0	REV 11	710-001226	AN5813	
SFM 0 SPP	REV 07	710-001228	AG4676	
SFM 0 SPR	REV 05	710-002189	AE4735	Internet Processor II
SFM 1 SPP	REV 07	710-001228	AP1347	
SFM 1 SPR	REV 05	710-002189	BE0063	Internet Processor II
FPC 0	REV 01	710-011725	BE0669	M40e-EP-FPC Type 1
CPU	REV 01	710-004600	BD9504	
PIC 0	REV 03	750-003737	AY3991	4x G/E, 1000 BASE-SX
FPC 1	REV 01	710-005197	BD9842	M40e-FPC Type 2
CPU	REV 01	710-004600	BB4869	
PIC 0	REV 07	750-001900	AR8278	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005197	BD9824	M40e-FPC Type 2
CPU	REV 01	710-004600	BD9531	
PIC 0	REV 03	750-003737	AY3986	4x G/E, 1000 BASE-SX
FPC 4	REV 02	710-005078	BE0664	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9559	
PIC 0	REV 03	750-001894	AG7963	1x G/E, 1000 BASE-SX
PIC 2	REV 01	750-002575	AF2472	4x OC-3 SONET, SMIR
FPC 6	REV 02	710-005078	BE0652	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9607	
PIC 0	REV 02	750-002911	AN2286	4x F/E, 100 BASE-TX
PIC 2	REV 01	750-002577	AP6345	4x OC-3 SONET, MM

show chassis hardware (M120 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module
PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE

PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware detail (M120 Router)

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module
PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
ad0	248 MB	SILICONSYSTEMS INC	256M 126CT505S0763SC00110	Compact Flash
ad2	38154 MB	HTES41040G9SA00	MPBBTOX2HS2E3M	Hard Disk
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE

Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware models (M120 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:				
Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-013667		
FPM CIP	REV 02	710-011410		CRAFT-M120-S
PEM 0	Rev 05	740-011936		PWR-M120-AC-S
PEM 1	Rev 05	740-011936		PWR-M120-AC-S
Routing Engine 0	REV 03	740-014080		RE-A-1000-2048-S
CB 0	REV 03	710-011403		CB-M120-S
CB 1	REV 06	710-011403		CB-M120-S
FPC 1	REV 02	710-015908		M120-cFPC-1XGE-XFP
FPC 3				
PIC 0	REV 16	750-008155		PB-2GE-SFP-QPP
PIC 1	REV 09	750-007745		PC-40C3-SON-SMIR
PIC 2	REV 16	750-008155		PB-2GE-SFP-QPP
PIC 3	REV 07	750-011800		PB-8GE-TYPE2-SFP-IQ2
FPC 4				
PIC 0	REV 16	750-007141		PC-10GE-SFP
FPC 5				
PIC 1	REV 05	750-012052		PB-1CHOC3-SMIR-QPP
PIC 2	REV 01	750-013167		PE-4CHDS3-QPP
PIC 3	REV 01	750-010240		PB-1GE-SFP
Fan Tray 0				FFANTRAY-M120-S
Fan Tray 1				FFANTRAY-M120-S


```

Fan Tray 2
Fan Tray 3
RFANTRAY-M120-S
RFANTRAY-M120-S

```

show chassis hardware (M160 Router)

```

user@host> show chassis hardware
Item                Version  Part number  Serial number  Description
Chassis              REV 02   710-001245   101            M160
Midplane             REV 01   710-001642   S/N AB4107
FPM CMB              REV 01   710-001642   S/N AA2911
FPM Display          REV 01   710-001647   S/N AA2999
CIP                  REV 02   710-001593   S/N AA9563
PEM 0                Rev 01   740-001243   S/N KJ35769    DC
PEM 1                Rev 01   740-001243   S/N KJ35765    DC
PCG 0                REV 01   710-001568   S/N AA9794
PCG 1                REV 01   710-001568   S/N AA9804
Host 1               da000004f8d57001  teknor
MCS 1                REV 03   710-001226   S/N AA9777
SFM 0 SPP            REV 04   710-001228   S/N AA2975
SFM 0 SPR            REV 02   710-001224   S/N AA9838      Internet Processor I
SFM 1 SPP            REV 04   710-001228   S/N AA2860
SFM 1 SPR            REV 01   710-001224   S/N AB0139      Internet Processor I
FPC 0                REV 03   710-001255   S/N AA9806      FPC Type 1
CPU                  REV 02   710-001217   S/N AA9590
PIC 1                REV 05   750-000616   S/N AA1527      1x OC-12 ATM, MM
PIC 2                REV 05   750-000616   S/N AA1535      1x OC-12 ATM, MM
PIC 3                REV 01   750-000616   S/N AA1519      1x OC-12 ATM, MM
FPC 1                REV 02   710-001611   S/N AA9523      FPC Type 2
CPU                  REV 02   710-001217   S/N AA9571
PIC 0                REV 03   750-001900   S/N AA9626      1x STM-16 SDH, SMIR
PIC 1                REV 01   710-002381   S/N AD3633      2x G/E, 1000 BASE-SX
FPC 2                REV 03   710-001217   S/N AB3329      FPC Type OC192
CPU                  REV 01
PIC 0                REV 01
Fan Tray 0           Rear Bottom Blower
Fan Tray 1           Rear Top Blower
Fan Tray 2           Front Top Blower
Fan Tray 3           Front Fan Tray

```

show chassis hardware models (M160 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  CLEI code  FRU model number
Midplane             REV 03   710-009120
FPM Display          REV 02   710-009351
CIP                  REV 03   710-005926
PEM 2                Rev X4   740-009148
PEM 3                Rev X4   740-009148
Routing Engine 0     REV 02   740-008883
Routing Engine 1     REV 02   740-008883
FPC 0                REV 02   710-010419
PIC 0                REV 01   750-001323
PIC 1                REV 02   750-002987
PIC 2                REV 04   750-001894
PIC 3                REV 04   750-001896
FPC 1                REV 02   710-010419
PIC 0                REV 04   750-001894
PIC 1                REV 04   750-001894
PIC 3                REV 03   750-001894
FPC 2                REV 02   710-010419
CHAS-BP-M320-S
CRAFT-M320-S
CIP-M320-S
PWR-M-DC-S
PWR-M-DC-S
RE-1600-2048-S
RE-1600-2048-S
M320-FPC1
P-TUNNEL
PE-10C12-SON-SMIR
PB-1GE-SX
PB-10C12-SON-SMIR
M320-FPC1
PB-1GE-SX
PB-1GE-SX
PB-1GE-SX
M320-FPC1

```

PIC 0	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
FPC 3			
PIC 0	REV 03	750-001895	PB-10C12-SON-MM
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 04	750-003141	PB-1GE-SX-B
FPC 4	REV 02	710-010419	M320-FPC1
FPC 5	REV 02	710-010419	M320-FPC1
FPC 6	REV 02	710-010419	M320-FPC1
FPC 7			
PIC 0	REV 15	750-001901	PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900	PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900	PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737	PB-4GE-SX
SIB 0	REV 03	710-009184	SIB-M-S
SIB 1	REV 03	710-009184	SIB-M-S
SIB 2	REV 03	710-009184	SIB-M-S
SIB 3	REV 03	710-009184	SIB-M-S
Fan Tray 0			FFANTRAY-M320-S
Fan Tray 1			FFANTRAY-M320-S
Fan Tray 2			RFANTRAY-M320-S

show chassis hardware detail (M160 Router)

```

user@host> show chassis hardware detail
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 306456	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 306474	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 306388	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 306392	1 MB
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 302917	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 302662	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 302593	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 100160	1 MB
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
SSRAM	REV 01	710-000077	S/N 302836	1 MB
SDRAM 0	REV 01	710-001196	S00141	32 MB
SDRAM 1	REV 01	710-001196	S0010;	32 MB
SSRAM	REV 01	710-000077	S/N 302633	1 MB

SDRAM 0	REV 01	710-001196	S00143	32 MB
SDRAM 1	REV 01	710-001196	S00115	32 MB
SSRAM	REV 01	710-000077	S/N 302952	1 MB
SDRAM 0	REV 01	710-001196	S00135	32 MB
SDRAM 1	REV 01	710-001196	S001=3	32 MB
SSRAM	REV 01	710-000077	S/N 302892	1 MB
SDRAM 0	REV 01	710-001196	S00076	32 MB
SDRAM 1	REV 01	710-001196	S001=5	32 MB
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
SSRAM	REV 01	710-000077	S/N 306340	1 MB
SDRAM 0	REV 01	710-001196	S00012	32 MB
SDRAM 1	REV 01	710-001196	S0001?	32 MB
SSRAM	REV 01	710-000077	S/N 306454	1 MB
SDRAM 0	REV 01	710-001196	S00028	32 MB
SDRAM 1	REV 01	710-001196	S0002?	32 MB
SSRAM	REV 01	710-000077	S/N 306492	1 MB
SDRAM 0	REV 01	710-001196	S00015	32 MB
SDRAM 1	REV 01	710-001196	S00031	32 MB
SSRAM	REV 01	710-000077	S/N 306363	1 MB
SDRAM 0	REV 01	710-001196	S00013	32 MB
SDRAM 1	REV 01	710-001196	S00032	32 MB
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
... SSRAM	REV 01	710-000077	S/N 306466	1 MB

show chassis hardware (M320 Router)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			67245	M320
Midplane	REV 05	710-009120	RB1202	M320 Midplane
FPM GBUS	REV 04	710-005928	HZ5697	M320 Board
FPM Display	REV 05	710-009351	HR1464	M320 FPM Display
CIP	REV 04	710-005926	HT8672	M320 CIP
PEM 0	Rev 05	740-009148	QK34208	DC Power Entry Module
PEM 1	Rev 05	740-009148	QK34262	DC Power Entry Module
PEM 2	Rev 05	740-009148	QF10449	DC Power Entry Module
PEM 3	Rev 05	740-009148	QJ18257	DC Power Entry Module
Routing Engine 0	REV 06	740-008883	P11123901185	RE-4.0
CB 0	REV 07	710-009115	JB2382	M320 Control Board
FPC 0	REV 02	710-005017	CD9926	M320 FPC Type 2
CPU	REV 01	710-011659	CJ6940	M320 PCA SCPU
PIC 0	REV 07	750-001900	AT1594	1x OC-48 SONET, SMSR
PIC 1	REV 03	750-001850	HS2746	1x Tunnel
PIC 2	REV 05	750-010618	JE7117	4x G/E SFP, 1000 BASE
PIC 3	REV 06	750-001900	HE6083	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005017	CH0319	M320 FPC Type 1
CPU	REV 01	710-011659	CJ6942	M320 PCA SCPU
PIC 0	REV 05	750-003034	BD8705	4x OC-3 SONET, SMIR
FPC 5	REV 02	710-005017	CD9938	M320 FPC Type 2
CPU				
FPC 7	REV 02	710-005017	CD9934	M320 FPC Type 2
CPU				
SIB 0	REV 09	710-009184	JA6540	M320 SIB
SIB 1	REV 09	710-009184	HV9511	M320 SIB

SIB 2	REV 09	710-009184	HW2057	M320 SIB
SIB 3	REV 09	710-009184	JA6687	M320 SIB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (M320 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S
Routing Engine 1	REV 02	740-008883		RE-1600-2048-S
FPC 0	REV 02	710-010419		M320-FPC1
PIC 0	REV 01	750-001323		P-TUNNEL
PIC 1	REV 02	750-002987		PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-001896		PB-10C12-SON-SMIR
FPC 1	REV 02	710-010419		M320-FPC1
PIC 0	REV 04	750-001894		PB-1GE-SX
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 03	750-001894		PB-1GE-SX
FPC 2	REV 02	710-010419		M320-FPC1
PIC 0	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
FPC 3				
PIC 0	REV 03	750-001895		PB-10C12-SON-MM
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-003141		PB-1GE-SX-B
FPC 4	REV 02	710-010419		M320-FPC1
FPC 5	REV 02	710-010419		M320-FPC1
FPC 6	REV 02	710-010419		M320-FPC1
FPC 7				
PIC 0	REV 15	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900		PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737		PB-4GE-SX
SIB 0	REV 03	710-009184		SIB-M-S
SIB 1	REV 03	710-009184		SIB-M-S
SIB 2	REV 03	710-009184		SIB-M-S
SIB 3	REV 03	710-009184		SIB-M-S
Fan Tray 0				FFANTRAY-M320-S
Fan Tray 1				FFANTRAY-M320-S
Fan Tray 2				RFANTRAY-M320-S

show chassis hardware (MX5 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			E1368	MX5-T

Midplane	REV 01	711-038215	YF5288	MX5-T
PEM 0	Rev 04	740-028288	VA01215	AC Power Entry Module
PEM 1	Rev 04	740-028288	VA01218	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9136	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX9820	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUAQ3	SFP-SX
Xcvr 1	REV 01	740-031851	AM1045SUAPA	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUAN7	SFP-SX
Xcvr 3	REV 01	740-031851	AM1045SU91Q	SFP-SX
Xcvr 4	REV 01	740-031851	AM1045SUDDR	SFP-SX
Xcvr 9	REV 01	740-011613	AM0848SB6A1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUANO	SFP-SX
Xcvr 1	REV 01	740-011613	AS0812S0719	SFP-SX
Xcvr 2	REV 01	740-011613	AM0821SA121	SFP-SX
Xcvr 3	REV 01	740-011613	PF21K21	SFP-SX
Xcvr 4	REV 01	740-011613	AM0848SB69Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9POXV3	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8WJN	SFP-SX
Xcvr 7	REV 01	740-011613	PAM3G9Q	SFP-SX
Xcvr 8	REV 01	740-011613	AM0848SB4A6	SFP-SX
Xcvr 9	REV 01	740-011782	P9MOU37	SFP-SX
MIC 1	REV 20	750-028380	ZG2657	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Fan Tray				Fan Tray

show chassis hardware (MX10 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			E1372	MX10-T
Midplane	REV 01	711-038211	YF5285	MX10-T
PEM 0	Rev 04	740-028288	VB01678	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9053	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX9436	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1107SUFQW	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Fan Tray				Fan Tray

show chassis hardware (MX40 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			E1367	MX40-T
Midplane	REV 01	711-038211	YF5284	MX40-T
PEM 0	Rev 04	740-028288	VB01680	AC Power Entry Module
PEM 1	Rev 04	740-028288	VB01700	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9048	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
Xcvr 0	REV 01	740-014279	M7067UPP	XFP-10G-LR
Xcvr 1		NON-JNPR	K9J02UN	XFP-10G-LR
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX3504	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0812S8WTE	SFP-SX
Xcvr 1	REV 01	740-011613	PFA6KV2	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUDDM	SFP-SX
Xcvr 3	REV 01	740-011613	PD63C7M	SFP-SX
Xcvr 4	REV 01	740-011613	PD63DJY	SFP-SX
Xcvr 5	REV 02	740-011613	AA0950STLL9	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1YHC	SFP-SX
Xcvr 7	REV 01	740-011782	P9P0XXL	SFP-SX
Xcvr 8	REV 01	740-011613	PD63D95	SFP-SX
Xcvr 9	REV 01	740-031851	AM1045SU9B8	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	PF21L3Z	SFP-SX
Xcvr 1	REV 01	740-031851	AM1045SU7M9	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUAPT	SFP-SX
Xcvr 3	REV 01	740-011613	PFF2BZH	SFP-SX
Xcvr 4	REV 01	740-031851	AM1045SUDDN	SFP-SX
Xcvr 5	REV 01	740-031851	AM1039S00ZR	SFP-SX
Xcvr 6	REV 01	740-031851	AM1045SUD6Y	SFP-SX
Xcvr 8	REV 01	740-011613	PFM1QBS	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2E25	SFP-SX
MIC 1	REV 01	750-021130	KG4391	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-011571	C645XJ04G	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0		NON-JNPR	CA49BK0AE	XFP-10G-SR
Fan Tray				Fan Tray

show chassis hardware (Fixed MX80 Router)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				MX80-48T
Midplane	REV 01	711-031603	KF9250	MX80-48T
Routing Engine		BUILTIN	BUILTIN	Routing Engine
FEB 0		BUILTIN	BUILTIN	Forwarding Engine Board
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
Xcvr 0		NON-JNPR	M6439D41	XFP-10G-LR
Xcvr 1	REV 01	740-014279	6XE931N00202	XFP-10G-LR
Xcvr 2	REV 01	740-014289	C715XU05F	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C650XU0EP	XFP-10G-SR

FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 01	711-029399	JR6981	12x 1GE(LAN) RJ45
PIC 0		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 1		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
MIC 1	REV 01	BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 2		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 3		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
Fan Tray				Fan Tray

show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis
Midplane            REV 02    711-031594   JR7084         MX80
PEM 0               Rev 01    740-028288   000018         AC Power Entry Module
Routing Engine
FEB 0
QXM 0               REV 05    711-028408   JR7041         MPC QXM
FPC 0
MIC 0
PIC 0
FPC 1
MIC 0               REV 02    750-028380   JR6598         3D 2x 10GE XFP
PIC 0
Xcvr 0              REV 01    740-014289   T07M86365     XFP-10G-SR
PIC 1
Xcvr 0              REV 01    740-014289   T07M71094     XFP-10G-SR
MIC 1               REV 02    750-028380   JG8548         3D 2x 10GE XFP
PIC 2
Xcvr 0              REV 02    740-014289   T08L86302     XFP-10G-SR
PIC 3
Xcvr 0              REV 02    740-014289   C810XU0BA     XFP-10G-SR
Fan Tray
```

show chassis hardware (MX240 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis
Midplane            REV 01    710-021041   TR1502         MX240 Backplane
FPM Board           REV 01    710-017254   KD4017         Front Panel Display
PEM 0               Rev 02    740-017330   000332         PS 1.2-1.7kW; 100-240V
AC in
PEM 1               Rev 02    740-017330   000226         PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0    REV 06    740-013063   1000703522     RE-S-2000
Routing Engine 1    REV 06    740-015113   1000687625     RE-S-1300
CB 0
CB 1               REV 05    710-013385   JY4760         MX SCB
FPC 1               REV 01    750-021679   KC7340         DPCE 40x 1GE R
CPU                 REV 06    710-013713   KD4078         DPC PMB
PIC 0
Xcvr 0              REV 01    740-011613   P9F18ME        SFP-SX
PIC 1
PIC 2
PIC 3
```

FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware detail (MX 240 Router with Routing Engine Displaying DIMM information)

```
user@host> show chassis hardware detail
```

Item	Version	Part number	Serial number	Description
Chassis			JN11279B4AFC	MX240 Backplane
Midplane	REV 07	760-021404	TS2474	MX240 Backplane
FPM Board	REV 03	760-021392	XC2643	Front Panel Display
PEM 0	Rev 03	740-017343	QCS0908A068	DC Power Entry Module
Routing Engine 0	REV 01	740-031117	AARCH00	RE-S-1800x4
ad0 3764 MB	STEC M2+	CF 9.0.2	STIM2Q3209239145303	Removable Compact Flash
ad1 28626 MB	WDC SSD-F0030S-5000		C933Z036237215548S00	Compact Flash
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263E-F8S DIE REV-0	PCB REV-0		MFR ID-ce80
DIMM 1	VL31B5263E-F8S DIE REV-0	PCB REV-0		MFR ID-ce80
DIMM 2	VL31B5263E-F8S DIE REV-0	PCB REV-0		MFR ID-ce80
DIMM 3	SL31B5263E-F8S DIE REV-0	PCB REV-0		MFR ID-ce80
CB 0	REV 03	710-021523	XD7225	MX SCB
Fan Tray 0	REV 01	710-021113	WZ4986	MX240 Fan Tray

show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7EAFC	MX240
Midplane	REV 01	710-021041	TR1502	MX240 Backplane
FPM Board	REV 01	710-017254	KD4017	Front Panel Display
PEM 0	Rev 02	740-017330	000332	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	000226	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 06	740-013063	1000703522	RE-S-2000
Routing Engine 1	REV 06	740-015113	1000687625	RE-S-1300

CB 0	REV 02	710-031391	YE8494	Enhanced MX SCB
CB 1	REV 05	710-031391	YOP5764	Enhanced MX SCB
FPC 1	REV 01	750-021679	KC7340	DPCE 40x 1GE R
CPU	REV 06	710-013713	KD4078	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18ME	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7FAFB	MX480
Midplane	REV 04	710-017414	TR2071	MX480 Midplane
FPM Board	REV 02	710-017254	KB8459	Front Panel Display
PEM 0	Rev 02	740-017330	QCS07519029	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	QCS07519041	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 02	740-017330	QCS07519097	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 07	740-013063	1000733381	RE-S-2000
Routing Engine 1	REV 07	740-013063	1000733540	RE-S-2000
CB 0	REV 07	710-013385	KA8022	MX SCB
CB 1	REV 07	710-013385	KA8303	MX SCB
FPC 0	REV 09	750-020452	KA8660	DPCE 40x 1GE X EQ
CPU	REV 06	710-013713	KA8185	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Fan Tray				Left Fan Tray

show chassis hardware (MX480 Router with Enhanced MX SCB)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis              JN10C7F7FAFB  MX480
Midplane            REV 04   710-017414   TR2071         MX480 Midplane
FPM Board           REV 02   710-017254   KB8459         Front Panel Display
PEM 0               Rev 02   740-017330   QCS07519029    PS 1.2-1.7kW; 100-240V
AC in
PEM 1               Rev 02   740-017330   QCS07519041    PS 1.2-1.7kW; 100-240V
AC in
PEM 2               Rev 02   740-017330   QCS07519097    PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0    REV 07   740-013063   1000733381     RE-S-2000
Routing Engine 1    REV 07   740-013063   1000733540     RE-S-2000
CB 0                REV 07   710-013385   KA8022         Enhanced MX SCB
CB 1                REV 07   710-013385   KA8303         Enhanced MX SCB
FPC 0               REV 09   750-020452   KA8660         DPCE 40x 1GE X EQ
CPU                 REV 06   710-013713   KA8185         DPC PMB
PIC 0               BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 1               BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 2               BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 3               BUILTIN   BUILTIN      10x 1GE(LAN) EQ
Fan Tray
Left Fan Tray

```

show chassis hardware (MX960 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis              AA6082     MX960
Midplane            REV 01   710-013698   710-013698     MX960 Midplane
PIM                 Rev 01   740-013110   000008         Power Inlet Module
PEM 2
PEM 3               Rev 01   740-013682   000038         PS 1.7kW; 200-240VAC in
Routing Engine 0    REV 00   740-015113   1000617944     RE-S-1300
CB 0                REV 05   710-013725   JK6947         MX960 Test SCB
FPC 4               REV 01   710-013305   JM7617         MX960 Test DPC
CPU
PIC 0               BUILTIN   BUILTIN      1x 10GE(LAN/WAN)
PIC 1               BUILTIN   BUILTIN      10x 1GE
FPC 7               REV 01   710-013305   JL9634         MX960 Test DPC
CPU
PIC 0               BUILTIN   BUILTIN      1x 10GE(LAN/WAN)
Xcvr 0              NON-JNPR   MYBG65I82C    XFP-10G-SR
PIC 1               BUILTIN   BUILTIN      10x 1GE
Xcvr 1              REV 01   740-011782   P7N0368        SFP-SX
Xcvr 4              REV 01   740-011782   P8J1W27        SFP-SX
Xcvr 6              REV 01   740-011782   P8J1VSD        SFP-SX
Xcvr 9              REV 01   740-011782   P8J1W25        SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware (MX960 Router with Bidirectional Optics)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis              JN10BA5B9AFA  MX960
Midplane            REV 03   710-013698   TR0234         MX960 Backplane

```

FPM Board	REV 03	710-014974	JA0878	Front Panel Display
PDM	Rev 03	740-013110	QCS11135028	Power Distribution Module
PEM 0	Rev 03	740-013682	QCS11154036	PS 1.7kW; 200-240VAC in
PEM 1	Rev 03	740-013682	QCS11154010	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11154022	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000691458	RE-S-2000
CB 0	REV 07	710-013385	KA2190	MX SCB
CB 1	REV 07	710-013385	KA0837	MX SCB
FPC 3	REV 02	750-018122	KB3890	DPCE 40x 1GE R
CPU				
FPC 4	REV 01	750-018122	KB3889	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3976	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 1	REV 01	740-020426	4910549	SFP-1000BASE-BX40-D
Xcvr 2	REV 01	740-020426	4910551	SFP-1000BASE-BX40-D
Xcvr 5	REV 01	740-021340	77E245N00006	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-020425	4882821	SFP-1000BASE-BX40-U
Xcvr 8	REV 01	740-020425	4882820	SFP-1000BASE-BX40-U
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020465	77E555N00894	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020465	75E467X00818	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020465	75E467X00573	SFP-1000BASE-BX10-D
Xcvr 3	REV 01	740-020465	4888227	SFP-1000BASE-BX10-D
Xcvr 4	REV 01	740-020465	4888241	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021340	77E245N00005	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-021340	76E245X00487	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021341	5255889	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255887	SFP-1000BASE-BX10-U
Xcvr 9	REV 01	740-021340	77E245N00004	SFP-1000BASE-BX10-U
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020424	5007582	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020424	4888187	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020424	4656500	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021341	5255886	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021340	77E245N00003	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255888	SFP-1000BASE-BX10-U
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-017726	74S184H30341	SFP-EX
Xcvr 1	REV 01	740-017726	4814061	SFP-EX
Xcvr 5	REV 01	740-017726	6ZS184H31108	SFP-EX
Xcvr 9	REV 01	740-021340	76E245X00486	SFP-1000BASE-BX10-U
Fan Tray 0				
Fan Tray 1	REV 03	740-014971	TP0850	Fan Tray

show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1096805AFA	MX960
Midplane	REV 03	710-013698	TR0183	MX960 Backplane
Fan Extender	REV 02	710-018051	JY5227	Extended Cable Manager
FPM Board	REV 03	710-014974	JZ6876	Front Panel Display
PDM	Rev 03	740-013110	QCS11035023	Power Distribution Module
PEM 1	Rev 03	740-013682	QCS1109400L	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11094015	PS 1.7kW; 200-240VAC in
PEM 3	Rev 03	740-013682	QCS11094012	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000
CB 0	REV 11	750-031391	YZ6072	Enhanced MX SCB
CB 1	REV 11	750-031391	YZ6068	Enhanced MX SCB

CB 2	REV 11	750-031391	YZ6081	Enhanced MX SCB
FPC 0	REV 01	750-018122	KA5576	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3961	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18GF	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TL9	SFP-SX
Xcvr 7	REV 01	740-011782	P9POXXH	SFP-SX
Xcvr 9	REV 01	740-011782	P9M0TN1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PAJ4UHC	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PFF2CD0	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3ZUT	SFP-SX
Xcvr 2	REV 01	740-011613	PFF2DDV	SFP-SX
Xcvr 5	REV 01	740-011613	P8E2SST	SFP-SX
Xcvr 9	REV 01	740-011782	PB8329N	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-026192	1U0201084503342	SFP-100BASE-BX10-U
Xcvr 1	REV 01	740-026193	1U1201084503313	SFP-100BASE-BX10-D
Xcvr 2	REV 01	740-011613	PAJ4Y5B	SFP-SX
Xcvr 6	REV 01	740-011782	P9MOU3M	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0TLA	SFP-SX
FPC 1	REV 16	750-031089	YL0719	MPC Type 2 3D
CPU	REV 06	711-030884	YL1463	MPC PMB 2G
MIC 0	REV 07	750-028387	JR6500	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	733019A00154	XFP-10G-LR
Xcvr 1	REV 02	740-014289	T09F55034	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	913019B00791	XFP-10G-LR
Xcvr 1	REV 01	740-014289	98S803A90384	XFP-10G-SR
MIC 1	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014279	T10B36134	XFP-10G-LR
Xcvr 1	REV 01	740-014289	T07M86354	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 2	REV 08	710-014219	JY9654	DPCE 4x 10GE R
CPU	REV 06	710-013713	JZ6549	DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
Xcvr 0	REV 03	740-011571	C931BK028	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
FPC 3	REV 10	750-024199	XJ6692	MX FPC Type 3
CPU	REV 03	710-022351	XF5182	DPC PMB
PIC 0	REV 17	750-009553	RJ2945	4x OC-48 SONET
Xcvr 1	REV 01	740-011785	PCP3YLL	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMRY	SFP-SR
PIC 1	REV 32	750-003700	DP2113	1x OC-192 12xMM VSR
FPC 5	REV 25	750-028467	YM8256	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YL3029	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 1	REV 01	740-031980	AHNOX1Z	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
FPC 7	REV 02	750-031092	JR6658	MPC Type 1 3D Q
CPU	REV 01	711-030884	JZ9038	MPC PMB 2G
MIC 0	REV 08	750-028392	JZ8737	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PBE2C6Y	SFP-SX

Xcvr 2		NON-JNPR	U8105N8	SFP-SX
Xcvr 4	REV 01	740-011613	PFM18EF	SFP-SX
Xcvr 7	REV 01	740-011613	PFF2AM8	SFP-SX
Xcvr 8	REV 01	740-011613	PFF2CT6	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PB82VHH	SFP-SX
Xcvr 1	REV 01	740-011613	PFF2CSW	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2BY0	SFP-SX
QXM 0	REV 04	711-028408	JR6372	MPC QXM
FPC 8	REV 05	750-024387	JW9754	MX FPC Type 2
CPU	REV 03	710-022351	KF1651	DPC PMB
PIC 0	REV 08	750-014730	DM3664	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 01	740-016065	81S290N00077	SFP-SR
Xcvr 1		NON-JNPR	2191844	SFP-SR
Xcvr 2	REV 01	740-011618	PD81EE5	SFP-IR
PIC 1	REV 08	750-014637	DM3671	4x OC-12-3 SFP
Xcvr 0	REV 01	740-011785	PCK3UNK	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMPZ	SFP-SR
FPC 10	REV 04	710-013699	JY4654	DPCE 40x 1GE R
CPU	REV 05	710-013713	JS9717	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 5	REV 01	740-011782	PAR1L72	SFP-SX
Xcvr 6	REV 01	740-011782	P8N1YQ4	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011782	P8Q2AVL	SFP-SX
Xcvr 5	REV 01	740-011782	PAR1L7B	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1L2J	SFP-SX
Xcvr 8	REV 01	740-011782	P8N1YMY	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Fan Tray 0	REV 03	740-014971	TP0567	Fan Tray
Fan Tray 1	REV 03	740-014971	TP0702	Fan Tray

show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-013698	TR0183	CHAS-BP-MX960-S
Fan Extender	REV 02	710-018051	JY5227	ECM-MX960
FPM Board	REV 03	710-014974	JZ6876	CRAFT-MX960-S
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000-4096-S
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000-4096-S
CB 0	REV 11	750-031391	YZ6072	SCBE-MX-S
CB 1	REV 11	750-031391	YZ6068	SCBE-MX-S
CB 2	REV 11	750-031391	YZ6081	SCBE-MX-S
FPC 0	REV 01	750-018122	KA5576	DPCE-R-40GE-SFP
FPC 1	REV 16	750-031089	YL0719	MX-MPC2-3D
MIC 0	REV 07	750-028387	JR6500	MIC-3D-4XGE-XFP
MIC 1	REV 24	750-028387	YJ3950	MIC-3D-4XGE-XFP
FPC 2	REV 08	710-014219	JY9654	DPC-R-4XGE-XFP
FPC 3	REV 10	750-024199	XJ6692	MX-FPC3
PIC 0	REV 17	750-009553	RJ2945	PC-40C48-SON-SFP
PIC 1	REV 32	750-003700	DP2113	PC-10C192-SON-VSR
FPC 5	REV 25	750-028467	YM8256	MPC-3D-16XGE-SFP
FPC 7	REV 02	750-031092	JR6658	MX-MPC1-3D-Q
MIC 0	REV 08	750-028392	JZ8737	MIC-3D-20GE-SFP
FPC 8	REV 05	750-024387	JW9754	MX-FPC2
PIC 0	REV 08	750-014730	DM3664	PB-40C3-10C12-SON2-SFP
PIC 1	REV 08	750-014637	DM3671	PB-40C3-40C12-SON-SFP
FPC 10	REV 04	710-013699	JY4654	DPC-R-40GE-SFP

Fan Tray 0	REV 03	740-014971	TP0567	FFANTRAY-MX960-S
Fan Tray 1	REV 03	740-014971	TP0702	FFANTRAY-MX960-S

show chassis hardware detail (MX960 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis
Midplane          REV 01    710-013698   AA6082         MX960 Midplane
PIM               Rev 01    740-013110   000008         Power Inlet Module
PEM 2
PEM 3             Rev 01    740-013682   000038         PS 1.7kW; 200-240VAC in
Routing Engine 0  REV 00    740-015113   1000617944     RE-S-1300
  ad0             245 MB   SanDisk     SDCFB-256      111419E1805T1141 Compact Flash
  ad2             38154 MB FUJITSU     MHT2040BH      NR0WT5925N77    Hard Disk
CB 0              REV 05    710-013725   JK6947         MX960 Test SCB
FPC 4             REV 01    710-013305   JM7617         MX960 Test DPC
CPU
PIC 0
PIC 1
FPC 7             REV 01    710-013305   JL9634         MX960 Test DPC
CPU
PIC 0
  Xcvr 0
PIC 1
  Xcvr 1          REV 01    740-011782   P7N0368        SFP-SX
  Xcvr 4          REV 01    740-011782   P8J1W27        SFP-SX
  Xcvr 6          REV 01    740-011782   P8J1VSD        SFP-SX
  Xcvr 9          REV 01    740-011782   P8J1W25        SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware (MX2010 Router)

```

user@host > show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis
Midplane          REV 01    750-044636   ABAB8506       Lower Backplane
Midplane 1        REV 01    711-044557   ZY8296         Upper Backplane
PMP               REV 03    711-032426   ACAJ1388       Power Midplane
FPM Board         REV 06    711-032349   ZX8744         Front Panel Display
PSM 4             REV 0C    740-033727   VK00254        DC 52V Power Supply
Module
PSM 5             REV 0B    740-033727   VG00015        DC 52V Power Supply
Module
PSM 6             REV 0B    740-033727   VH00097        DC 52V Power Supply
Module
PSM 7             REV 0C    740-033727   VJ00151        DC 52V Power Supply
Module
PSM 8             REV 0C    740-033727   VJ00149        DC 52V Power Supply
Module
PDM 0             REV 0B    740-038109   WA00008        DC Power Dist Module
PDM 1             REV 0B    740-038109   WA00014        DC Power Dist Module
Routing Engine 0  REV 02    740-041821   9009094134     RE-S-1800x4
Routing Engine 1  REV 02    740-041821   9009094141     RE-S-1800x4
CB 0              REV 08    750-040257   CAAB3491       Control Board
CB 1              REV 08    750-040257   CAAB3489       Control Board
SPMB 0            REV 02    711-041855   CAAA6135       PMB Board
SPMB 1            REV 02    711-041855   CAAA6137       PMB Board

```

SFB 0	REV 06	711-032385	ZV1828	Switch Fabric Board
SFB 1	REV 07	711-032385	ZZ2568	Switch Fabric Board
SFB 2	REV 07	711-032385	ZZ2563	Switch Fabric Board
SFB 3	REV 07	711-032385	ZZ2564	Switch Fabric Board
SFB 4	REV 07	711-032385	ZZ2580	Switch Fabric Board
SFB 5	REV 07	711-032385	ZZ2579	Switch Fabric Board
SFB 6	REV 07	711-032385	CAAB4882	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4898	Switch Fabric Board
FPC 0	REV 33	750-028467	CAAB1919	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAB7174	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH02RE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH038C	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH0390	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMG0SUA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0579	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0SGP	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH04SV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04X3	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0135	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02NC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02XB	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH02PN	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH057Y	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0JHE	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02HT	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04V4	SFP+-10G-SR
FPC 1	REV 21	750-033205	ZG5027	MPC Type 3
CPU	REV 04	711-035209	YT4780	HMPD PMB 2G
MIC 0	REV 03	750-033307	ZV6299	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	083363A00410	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	083363A00334	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	113363A00125	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	083363A00953	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AHR013D	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ40JUR	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JKL	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ30ECK	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100864	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511100868	SFP+-10G-SR
MIC 1	REV 03	750-033307	ZV6268	10X10GE SFPP
PIC 2		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	AJC0JML	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403PC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ10N25	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JF4	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JSJ	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ403V7	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JN3	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ40JSU	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100468	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511101363	SFP+-10G-SR
FPC 8	REV 22	750-031089	ZT9746	MPC Type 2 3D
CPU	REV 06	711-030884	ZS1271	MPC PMB 2G
MIC 0	REV 26	750-028392	ABBS1150	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	PLG023C	SFP-SX

Xcvr 1	REV 01	740-031851	PLG09C6	SFP-SX
Xcvr 2	REV 02	740-011613	AM0950SF9L7	SFP-SX
Xcvr 3	REV 02	740-011613	AM1001SFN1H	SFP-SX
Xcvr 4	REV 02	740-011613	AM1001SFM9D	SFP-SX
Xcvr 5	REV 02	740-011613	AM1001SFLTJ	SFP-SX
Xcvr 6	REV 01	740-031851	AC1108S03L9	SFP-SX
Xcvr 7	REV 01	740-031851	AC1102S00NC	SFP-SX
Xcvr 8	REV 01	740-031851	AC1102S00MX	SFP-SX
Xcvr 9	REV 01	740-031851	AC1102S0085	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AC1102S00KU	SFP-SX
Xcvr 1	REV 01	740-031851	AC1102S00NG	SFP-SX
Xcvr 2	REV 01	740-031851	AC1102S00K3	SFP-SX
Xcvr 3	REV 01	740-031851	AC1102S008R	SFP-SX
Xcvr 4	REV 01	740-031851	AM1107SUFVJ	SFP-SX
Xcvr 5	REV 01	740-031851	AC1108S03LG	SFP-SX
MIC 1	REV 26	750-028387	ABBR9582	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T10A91703	XFP-10G-SR
Xcvr 1		NON-JNPR	T09L42604	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 9	REV 11	750-036284	ZL3591	MPC 3D 16x 10GE EM
CPU	REV 10	711-029089	ZL0513	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101825	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101821	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101682	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ13R6	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101828	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101716	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALP0TR1	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101741	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101829	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14E3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101826	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101817	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101735	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ159A	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2073	Adapter Card
ADC 1	REV 01	750-043596	ZV4117	Adapter Card
ADC 8	REV 01	750-043596	ZV4107	Adapter Card
ADC 9	REV 02	750-043596	ZW1555	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0015	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0019	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0020	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0021	172mm FanTray - 6 Fans

show chassis hardware detail (MX2010 Router)

user@host > show chassis hardware detail

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E233DAFK	MX2010
Midplane	REV 26	750-044636	ABAB9357	Lower Backplane
Midplane 1	REV 01	711-044557	ABAB8643	Upper Backplane
PMP	REV 04	711-032426	ACAJ1677	Power Midplane

FPM Board	REV 08	760-044634	ABBV9726	Front Panel Display
PSM 0 Module	REV 01	740-045050	1E02224000P	DC 52V Power Supply
PSM 1 Module	REV 01	740-045050	1E02224000M	DC 52V Power Supply
PSM 2 Module	REV 01	740-045050	1E022240010	DC 52V Power Supply
PSM 3 Module	REV 01	740-045050	1E02224000G	DC 52V Power Supply
PSM 4 Module	REV 01	740-045050	1E022240013	DC 52V Power Supply
PSM 5 Module	REV 01	740-045050	1E022240007	DC 52V Power Supply
PSM 6 Module	REV 01	740-045050	1E02224001C	DC 52V Power Supply
PSM 7 Module	REV 01	740-045050	1E02224001D	DC 52V Power Supply
PSM 8 Module	REV 01	740-045050	1E02224001B	DC 52V Power Supply
PDM 0	REV 01	740-045234	1E262250067	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009099704	RE-S-1800x4
ad0 3831 MB	UGB30SFA4000T1	SFA4000T1 00000651	Compact Flash	
ad1 30533 MB	UGB94BPH32H0S1-KCI	11000019592	Disk 1	
usb0 (addr 1)	EHCI root hub 0	Intel	uhub0	
usb0 (addr 2)	product 0x0020 32	vendor 0x8087	uhub1	
DIMM 0	SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 1	SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 2	SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 3	SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
Routing Engine 1	REV 02	740-041821	9009099706	RE-S-1800x4
ad0 3998 MB	Virtium - TuffDrive	VCF P1T0200262860208 114	Compact Flash	
ad1 30533 MB	UGB94ARF32H0S3-KC	UNIGEN-499551-000404	Disk 1	
CB 0	REV 13	750-040257	CAAF8436	Control Board
CB 1	REV 13	750-040257	CAAF8434	Control Board
SPMB 0	REV 02	711-041855	ABBV3825	PMB Board
SPMB 1	REV 02	711-041855	ABBV3833	PMB Board
SFB 0	REV 05	711-044466	ABBX5682	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5676	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5665	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBX5699	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5603	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBX5587	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBX5607	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5669	Switch Fabric Board
FPC 0	REV 09	750-037355	CAAF0924	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9842	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	19T511101656	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA04RU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00558	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M00202	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00328	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AMA088W	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10L04211	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101602	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10L04151	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00332	CFP-100G-SR10
FPC 1	REV 18	750-033205	ZE0128	MPC Type 3

CPU	REV 06	711-035209	ZG5431	HMPC PMB 2G
MIC 0	REV 15	750-033199	ZP6435	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	J11E46118	CFP-100G-LR4
MIC 1	REV 15	750-033199	ZP6442	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	UMN03T4	CFP-100G-LR4
FPC 2	REV 16	750-037358	CAAL1001	MPC Type 4-1
CPU	REV 08	711-035209	CAAK7927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00589	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00028	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00376	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00016	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00499	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00039	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E01239	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B10M00075	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00014	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA0638	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00063	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AMA0629	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00053	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00344	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00046	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062M	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00080	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00580	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00064	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	093363A01494	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00020	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	123363A00047	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00072	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01033	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00022	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00013	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01028	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00079	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01018	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00025	SFP+-10G-SR
FPC 3	REV 33	750-028467	CAAF5400	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7626	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00066	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00021	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00062	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00027	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00065	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00069	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00003	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00035	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00004	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00049	SFP+-10G-SR

Xcvr 3	REV 01	740-021308	973152A00055	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00010	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00001	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00073	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00012	SFP+-10G-SR
FPC 4	REV 21	750-033205	ZG5028	MPC Type 3
CPU	REV 05	711-035209	YX3911	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2036	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220708	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220735	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2028	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220727	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220715	QSFP+-40G-SR4
FPC 5	REV 11	750-037358	CAAE2196	MPC Type 4-1
CPU	REV 08	711-035209	CAAD9074	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062S	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA062P	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA052R	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0632	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00564	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00229	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00363	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00278	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04CC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A001W	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA04N2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA062U	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00491	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	183363A01511	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00565	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00405	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA07QX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA06MS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00318	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	193363A00402	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00174	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00388	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00377	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00234	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00550	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00364	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0630	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00509	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00459	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	113363A00191	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00352	SFP+-10G-SR
FPC 6	REV 33	750-028467	CAAF5552	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7601	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AD0927A0036	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A003M	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0927A003G	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0927A0031	SFP+-10G-SR

PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	193363A00331	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00325	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00417	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A02509	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75140	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11A04356	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01952	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01914	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75157	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	T09K75194	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01926	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01936	SFP+-10G-SR
FPC 7	REV 16	750-037358	CAAL1012	MPC Type 4-1
CPU	REV 08	711-035209	CAAJ3851	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04NK	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00260	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11E02192	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04CP	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JJK	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11F00238	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B10M00275	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00211	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B11D05577	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11G00586	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA08B7	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04Q0	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11D05840	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11E00467	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E00029	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101712	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00568	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00166	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10M00212	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11D05823	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01005	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	03DZ06A01003	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01009	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01004	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01017	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	03DZ06A01016	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01024	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	03DZ06A01008	SFP+-10G-SR
Xcvr 4	REV 01	740-030658	AD0946A02UH	SFP+-10G-USR
Xcvr 5	REV 01	740-021308	T09J67913	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AD0837ES09G	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01015	SFP+-10G-SR
FPC 8	REV 03	750-045372	CAAD3111	MPC Type 3
CPU	REV 08	711-035209	CAAD8033	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2032	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB230273	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB230254	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2021	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP

Xcvr 0	REV 01	740-032986	QB390962	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB390960	QSFP+-40G-SR4
FPC 9	REV 09	750-037355	CAAF1531	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00525	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00504	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00368	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JSS	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	123363A00042	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00023	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ802EM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11E02348	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
ADC 0	REV 13	750-043596	ABBX5532	Adapter Card
ADC 1	REV 13	750-043596	ABBX5550	Adapter Card
ADC 2	REV 13	750-043596	ABBX5571	Adapter Card
ADC 3	REV 13	750-043596	ABBX5568	Adapter Card
ADC 4	REV 13	750-043596	ABBX5556	Adapter Card
ADC 5	REV 13	750-043596	ABBX5553	Adapter Card
ADC 6	REV 13	750-043596	ABBX5541	Adapter Card
ADC 7	REV 13	750-043596	ABBX5578	Adapter Card
ADC 8	REV 13	750-043596	ABBX5560	Adapter Card
ADC 9	REV 07	750-043596	ABBV7188	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0127	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0068	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0072	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0070	172mm FanTray - 6 Fans

show chassis hardware extensive (MX2010 Router)

```

user@host > show chassis hardware extensive
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              JN11E233DAFK  MX2010
  Jedec Code: 0x7fb0      EEPROM Version: 0x02
                        S/N: JN11E233DAFK
  Assembly ID: 0x0557      Assembly Version: 00.00
  Date: 00-00-0000      Assembly Flags: 0x00
  ID: MX2010
  Board Information Record:
    Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
    Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
    Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
    Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Midplane            REV 26  750-044636  ABAB9357      Lower Backplane
    Jedec Code: 0x7fb0      EEPROM Version: 0x02
    P/N: 750-044636      S/N: S/N ABAB9357
    Assembly ID: 0x0b66      Assembly Version: 01.26
    Date: 08-28-2012      Assembly Flags: 0x00
    Version: REV 26      CLEI Code: PROTOXCLEI
    ID: Lower Backplane      FRU Model Number: PROTO-ASSEMBLY
  Board Information Record:

```

```

Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

Midplane 1      REV 01      711-044557      ABAB8643      Upper Backplane
Jedec Code:     0x7fb0      EEPROM Version: 0x01
P/N:            711-044557      S/N:           S/N ABAB8643
Assembly ID:    0x0b65      Assembly Version: 01.01
Date:           07-27-2012      Assembly Flags: 0x00
Version:        REV 01
ID: Upper Backplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

PMP            REV 04      711-032426      ACAJ1677      Power Midplane
Jedec Code:     0x7fb0      EEPROM Version: 0x01
P/N:            711-032426      S/N:           S/N ACAJ1677
Assembly ID:    0x045d      Assembly Version: 01.04
Date:           07-20-2012      Assembly Flags: 0x00
Version:        REV 04
ID: Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Board      REV 08      760-044634      ABBV9726      Front Panel Display
Jedec Code:     0x7fb0      EEPROM Version: 0x02
P/N:            760-044634      S/N:           S/N ABBV9726
Assembly ID:    0x0b64      Assembly Version: 01.08
Date:           09-10-2012      Assembly Flags: 0x00
Version:        REV 08      CLEI Code:     IPMYA4EJRA
ID: Front Panel Display      FRU Model Number: MX2010-CRAFT-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d

```

```

Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0          REV 01   740-045050   1E02224000P   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045050      S/N:           1E02224000P
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          12-06-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01   740-045050   1E02224000M   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045050      S/N:           1E02224000M
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          12-06-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0          REV 01   740-045234   1E262250067   DC Power Dist Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045234      S/N:           1E262250067
Assembly ID:   0x047b          Assembly Version: 01.01
Date:          06-28-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     IPUPAJSKAA
ID: DC Power Dist Module      FRU Model Number: MX2000-PDM-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02   740-041821   9009099704   RE-S-1800x4

```

```

Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-041821          S/N: 9009099704
Assembly ID: 0x09c0        Assembly Version: 01.02
Date: 03-15-2012          Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4          FRU Model Number: RE-S-1800X4-16G-S

Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000651 Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009099706 RE-S-1800x4
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-041821          S/N: 9009099706
Assembly ID: 0x09c0        Assembly Version: 01.02
Date: 02-23-2012          Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4          FRU Model Number: RE-S-1800X4-16G-S

Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0 REV 13 750-040257 CAAF8436 Control Board
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 750-040257          S/N: S/N CAAF8436
Assembly ID: 0x0b26        Assembly Version: 01.13
Date: 08-29-2012          Assembly Flags: 0x00
Version: REV 13          CLEI Code: PROTOXCLEI
ID: Control Board          FRU Model Number: PROTO-ASSEMBLY

Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00

```



```

Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff
...
SPMB 0          REV 02   711-041855   ABBV3825          PMB Board
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-041855      S/N:             S/N ABBV3825
Assembly ID:   0x0b29          Assembly Version: 01.02
Date:          08-14-2012      Assembly Flags:   0x00
Version:       REV 02
ID: PMB Board
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
...
SFB 0          REV 05   711-044466   ABBX5682          Switch Fabric Board
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           711-044466      S/N:             S/N ABBX5682
Assembly ID:   0x0b25          Assembly Version: 01.05
Date:          09-07-2012      Assembly Flags:   0x00
Version:       REV 05          CLEI Code:        PROTOXCLEI
ID: Switch Fabric Board      FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 34 36 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 48 00
...
FPC 0          REV 09   750-037355   CAAF0924          MPC Type 4-2
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037355      S/N:             S/N CAAF0924
Assembly ID:   0x0b4e          Assembly Version: 01.09
Date:          05-21-2012      Assembly Flags:   0x00
Version:       REV 09          CLEI Code:        PROTOXCLEI
ID: MPC Type 4-2              FRU Model Number: MPC4E-2CGE-8XGE
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 08   711-035209   CAAB9842          HMPMPC PMB 2G
Jedec Code:    0x7fb0          EEPROM Version:    0x01

```

```

P/N:          711-035209          S/N:          S/N CAAB9842
Assembly ID:  0x0b04             Assembly Version: 01.08
Date:         05-17-2012         Assembly Flags:  0x00
Version:      REV 08
ID: HMPC PMB 2G
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN          BUILTIN          4x10GE SFPP
Jedec Code:    0x0000             EEPROM Version: 0x00
P/N:           BUILTIN            S/N:           BUILTIN
Assembly ID:   0x0a53             Assembly Version: 00.00
Date:         00-00-0000         Assembly Flags:  0x00
ID: 4x10GE SFPP
Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
  Xcvr 0      REV 01      740-021308      19T511101656      SFP+-10G-SR
  Xcvr 1      REV 01      740-031980      AMA04RU          SFP+-10G-SR
  Xcvr 2      REV 01      740-031980      193363A00558      SFP+-10G-SR
  Xcvr 3      REV 01      740-031980      B10M00202         SFP+-10G-SR
...
ADC 0          REV 13      750-043596      ABBX5532          Adapter Card
Jedec Code:    0x7fb0             EEPROM Version: 0x02
P/N:           750-043596         S/N:           S/N ABBX5532
Assembly ID:   0x0b3d             Assembly Version: 01.13
Date:         09-12-2012         Assembly Flags:  0x00
Version:      REV 13             CLEI Code:     IPUCBA8CAA
ID: Adapter Card                  FRU Model Number: MX2000-LC-ADAPTER
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
  Address 0x60: 00 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff
  Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00 00
...

```

show chassis hardware models (MX2010 Router)

```
user@host > show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
FPM Board	REV 06	711-032349	ZX8744	711-032349
PSM 4	REV 0C	740-033727	VK00254	000000000000000000000000
PSM 5	REV 0B	740-033727	VG00015	000000000000000000000000
PSM 6	REV 0B	740-033727	VH00097	000000000000000000000000
PSM 7	REV 0C	740-033727	VJ00151	000000000000000000000000
PSM 8	REV 0C	740-033727	VJ00149	000000000000000000000000
PDM 0	REV 0B	740-038109	WA00008	
PDM 1	REV 0B	740-038109	WA00014	
Routing Engine 0	REV 02	740-041821	9009094134	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009094141	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3491	750-040257
CB 1	REV 08	750-040257	CAAB3489	750-040257
SFB 0	REV 06	711-032385	ZV1828	711-032385
SFB 1	REV 07	711-032385	ZZ2568	711-032385
SFB 2	REV 07	711-032385	ZZ2563	711-032385
SFB 3	REV 07	711-032385	ZZ2564	711-032385
SFB 4	REV 07	711-032385	ZZ2580	711-032385
SFB 5	REV 07	711-032385	ZZ2579	711-0323856
SFB 6	REV 07	711-032385	CAAB4882	711-044170
SFB 7	REV 07	711-032385	CAAB4898	711-044170
FPC 0	REV 33	750-028467	CAAB1919	MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205	ZG5027	MX-MPC3-3D
MIC 0	REV 03	750-033307	ZV6299	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	ZV6268	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	ZT9746	MX-MPC2-3D
MIC 0	REV 26	750-028392	ABBS1150	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	ABBR9582	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	ZL3591	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	CAAC2073	750-043596
ADC 1	REV 01	750-043596	ZV4117	750-043596
ADC 8	REV 01	750-043596	ZV4107	750-043596
ADC 9	REV 02	750-043596	ZW1555	750-043596
Fan Tray 0	REV 2A	760-046960	ACAY0015	
Fan Tray 1	REV 2A	760-046960	ACAY0019	
Fan Tray 2	REV 2A	760-046960	ACAY0020	
Fan Tray 3	REV 2A	760-046960	ACAY0021	

show chassis hardware clei-models (MX2010 Routers)

user@host > show chassis hardware clei-models

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
FPM Board	REV 06	711-032349	PROTOXCLEI	711-032349
PSM 4	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 5	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 6	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 7	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 8	REV 0C	740-033727	0000000000	000000000000000000000000
PDM 0	REV 0B	740-038109		
PDM 1	REV 0B	740-038109		
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 08	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 06	711-032385	PROTOXCLEI	711-032385
SFB 1	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 2	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 3	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 4	REV 07	711-032385	PROTOXCLEI	711-032385

SFB 5	REV 07	711-032385	PROTOXCLEI	711-0323856
SFB 6	REV 07	711-032385	PROTOXCLEI	711-044170
SFB 7	REV 07	711-032385	PROTOXCLEI	711-044170
FPC 0	REV 33	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205		MX-MPC3-3D
MIC 0	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	COUIBAYBAA	MX-MPC2-3D
MIC 0	REV 26	750-028392	COUIA15BAA	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	CMUIACGBAA	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	PROTOXCLEI	750-043596
ADC 1	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 8	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 9	REV 02	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 2A	760-046960		
Fan Tray 1	REV 2A	760-046960		
Fan Tray 2	REV 2A	760-046960		
Fan Tray 3	REV 2A	760-046960		

show chassis hardware (MX2020 Router)

user@host > show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ1579	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ1524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224006W	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 01	740-045050	1E02224005B	DC 52V Power Supply
Module				
PSM 10	REV 01	740-045050	1E02224005A	DC 52V Power Supply
Module				
PSM 11	REV 01	740-045050	1E022240052	DC 52V Power Supply
Module				
PSM 12	REV 01	740-045050	1E022240051	DC 52V Power Supply
Module				
PSM 13	REV 01	740-045050	1E022240058	DC 52V Power Supply
Module				
PSM 14	REV 01	740-045050	1E02224004L	DC 52V Power Supply
Module				
PSM 15	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 16	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				

PSM 17	REV 01	740-045050	1E02224005Z	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E012150028	DC Power Dist Module
PDM 3	REV 01	740-045234	1E012150045	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089704	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094138	RE-S-1800x4
CB 0	REV 14	750-040257	CAAF8430	Control Board
CB 1	REV 08	750-040257	CAAB3482	Control Board
SPMB 0	REV 01	711-041855	ZS2290	PMB Board
SPMB 1	REV 02	711-041855	CAAA6141	PMB Board
SFB 0	REV 03	711-044466	ABBV6789	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5666	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5678	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBX5687	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5609	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBX5675	Switch Fabric Board
SFB 6	REV 03	711-044466	ABBV6805	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5701	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR

Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 2	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 3	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 4	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR

Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 5	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 6	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 7	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 8	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 9	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 10	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
FPC 11	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 12	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 14	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 17	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02638	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 19	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5538	Adapter Card
ADC 11	REV 13	750-043596	ABBX5566	Adapter Card
ADC 12	REV 13	750-043596	ABBX5542	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0030	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0039	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0033	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0062	172mm FanTray - 6 Fans

show chassis hardware detail (MX2020 Router)

user@host> show chassis hardware detail

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ1821	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ1524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E02224006G	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 5	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 6	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 03	740-045050	1EDB2350095	DC 52V Power Supply
Module				
PSM 10	REV 03	740-045050	1EDB235009L	DC 52V Power Supply
Module				
PSM 11	REV 03	740-045050	1EDB2350092	DC 52V Power Supply
Module				
PSM 12	REV 03	740-045050	1EDB23500AT	DC 52V Power Supply
Module				
PSM 13	REV 03	740-045050	1EDB2350094	DC 52V Power Supply
Module				
PSM 15	REV 03	740-045050	1EDB235008X	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E262250072	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800x4
ad0	3998 MB	Virtium - TuffDisk	VCf3 20110825A021D0000064	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-499551-000347	Disk 1
usb0 (addr 1)		EHCI root hub 0	Intel	uhub0
usb0 (addr 2)		product 0x0020 32	vendor 0x8087	uhub1
DIMM 0		SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54	MFR ID-ce80	
DIMM 1		SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54	MFR ID-ce80	
DIMM 2		SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54	MFR ID-ce80	
DIMM 3		SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54	MFR ID-ce80	
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800x4
ad0	3831 MB	UGB30SFA4000T1	SFA4000T1 00000113	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-478612-001044	Disk 1
CB 0	REV 08	750-040257	CAAB3482	Control Board
CB 1	REV 04	750-040257	ZT2864	Control Board
SPMB 0	REV 02	711-041855	CAA6141	PMB Board
SPMB 1	REV 01	711-041855	ZS2275	PMB Board
SFB 0	REV 05	711-044466	ABBT2161	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBT2159	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX3718	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBT2152	Switch Fabric Board

SFB 4	REV 05	711-044466	ABBT2160	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBT2145	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBT2150	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBT2163	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 2	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 3	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 4	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 5	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP

Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 6	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 7	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03058	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02638	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 8	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 9	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
FPC 10	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR

Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 11	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 12	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR

Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 14	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZD06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 17	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 19	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5579	Adapter Card
ADC 11	REV 13	750-043596	ABBX5548	Adapter Card
ADC 12	REV 13	750-043596	ABBX5575	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 04	760-046960	ACAY0090	172mm FanTray - 6 Fans
Fan Tray 1	REV 04	760-046960	ACAY0088	172mm FanTray - 6 Fans
Fan Tray 2	REV 04	760-046960	ACAY0089	172mm FanTray - 6 Fans
Fan Tray 3	REV 04	760-046960	ACAY0108	172mm FanTray - 6 Fans

show chassis hardware models (MX2020 Router)

```

user@host > show chassis hardware models
Hardware inventory:

```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 27	750-040240	ABAB9384	750-040240
FPM Board	REV 06	760-040242	ABBT8837	760-040242
PSM 0	REV 01	740-045050	1E02224006G	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	1E022240053	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	1E02224004K	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	1E022240056	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	1E022240054	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	1E02224005H	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	1E02224006S	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	1E02224005M	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	1E022240062	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	1EDB2350095	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	1EDB235009L	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	1EDB2350092	MX2000-PSM-DC-S-A

PSM 12	REV 03	740-045050	1EDB23500AT	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	1EDB2350094	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	1EDB235008X	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234	1E012150033	
PDM 1	REV 01	740-045234	1E012150027	
PDM 2	REV 01	740-045234	1E262250072	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3482	750-040257
CB 1	REV 04	750-040257	ZT2864	750-040257
SFB 0	REV 05	711-044466	ABBT2161	MX2000-SFB-S
SFB 1	REV 05	711-044466	ABBT2159	MX2000-SFB-S
SFB 2	REV 05	711-044466	ABBX3718	MX2000-SFB-S
SFB 4	REV 05	711-044466	ABBT2160	MX2000-SFB-S
SFB 5	REV 05	711-044466	ABBT2145	MX2000-SFB-S
SFB 7	REV 05	711-044466	ABBT2163	MX2000-SFB-S
FPC 0	REV 30	750-028467	ABBN0284	MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467	ABBN0308	MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467	ABBN0316	MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467	ABBN6832	MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467	ABBN6811	MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467	ABBN6791	MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467	ABBM4592	MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467	ABBN6810	MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467	ABBM4739	MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467	ABBN6827	MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467	ABBN0302	MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467	ABBN6790	MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467	ZM5111	MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467	ABBN0208	MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467	YN2977	MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467	ABBN6798	MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467	ABBN0270	MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467	ABBN6796	MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467	ABBN0281	MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467	ABBN6813	MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	ABBX5561	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	ABBX5546	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	ABBX5535	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	ABBX5552	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	ABBX5581	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	ABBX5545	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	ABBX5554	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	ABBV7194	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	ABBV7251	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	ABBV7202	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	ABBX5579	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	ABBX5575	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	ABBX5539	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	ABBX5555	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	ABBX5557	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	ABBX5536	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	ABBX5559	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	ABBX5537	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	ABBW5685	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960	ACAY0090	
Fan Tray 1	REV 04	760-046960	ACAY0088	
Fan Tray 2	REV 04	760-046960	ACAY0089	
Fan Tray 3	REV 04	760-046960	ACAY0108	

show chassis hardware clei-models (MX2020 Router)

user@ host > show chassis hardware clei-models

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 27	750-040240	PROTOXCLEI	750-040240
FPM Board	REV 06	760-040242	PROTOXCLEI	760-040242
PSM 0	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234		
PDM 1	REV 01	740-045234		
PDM 2	REV 01	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 04	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467		MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467		MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY

ADC 6	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960		
Fan Tray 1	REV 04	760-046960		
Fan Tray 2	REV 04	760-046960		
Fan Tray 3	REV 04	760-046960		

show chassis hardware (MX Series routers with ATM MIC)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis                               JN115736EAFc  MX240
Midplane          REV 07    760-021404  ABAA5038      MX240 Backplane
FPM Board         REV 03    760-021392  ABBA2758      Front Panel Display
PEM 0             Rev 01    740-022697  QCS0937C07K   PS 1.2-1.7kW; 100-240V
AC in
PEM 1             Rev 01    740-022697  QCS0939C04X   PS 1.2-1.7kW; 100-240V
AC in
PEM 2             Rev 01    740-022697  QCS0937C06B   PS 1.2-1.7kW; 100-240V
AC in
PEM 3             Rev 01    740-022697  QCS0937C07U   PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0  REV 12    740-013063  9009042291    RE-S-2000
Routing Engine 1  REV 12    740-013063  9009042266    RE-S-2000
CB 0              REV 06    710-021523  ABBC1435      MX SCB
CB 1              REV 06    710-021523  ABBC1497      MX SCB
FPC 2             REV 14    750-031088  YH8446        MPC Type 2 3D Q
CPU               REV 06    711-030884  YH9612        MPC PMB 2G
MIC 0
MIC 1             REV 10    750-036132  ZP7062        2x0C12/8x0C3 CC-CE
PIC 2             BUILtIN   BUILtIN      2x0C12/8x0C3 CC-CE

Xcvr 0            NON-JNPR   23393-00492  UNKNOWN
Xcvr 1            NON-JNPR   23393-00500  UNKNOWN
Xcvr 2            NON-JNPR   23393-00912  UNKNOWN
Xcvr 3            REV 01    740-015638  22216-00575   Load SFP
Xcvr 4            REV 01    740-015638  24145-00110   Load SFP
Xcvr 5            REV 01    740-015638  24145-00016   Load SFP
Xcvr 6            REV 01    740-015638  24145-00175   Load SFP
Xcvr 7            NON-JNPR   23393-00627  UNKNOWN
QXM 0             REV 05    711-028408  YF4681        MPC QXM
QXM 1             REV 05    711-028408  YF4817        MPC QXM
Fan Tray 0        REV 01    710-021113  XL3645        MX240 Fan Tray

```

show chassis hardware (MX240, MX480, MX960 routers with Application Services Modular Line Card)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description

```


Chassis			JN11D969BAFA	MX960
Midplane	REV 03	710-013698	ACAA2362	MX960 Backplane
FPM Board	REV 03	710-014974	ZR0639	Front Panel Display
PDM	Rev 03	740-013110	QCS152250SX	Power Distribution Module
PEM 0	Rev 10	740-013683	QCS1512718W	DC Power Entry Module
PEM 1	Rev 10	740-013683	QCS1512702Y	DC Power Entry Module
Routing Engine 0	REV 15	740-013063	9012024667	RE-S-2000
Routing Engine 1	REV 15	740-013063	9012024649	RE-S-2000
CB 0	REV 14	750-031391	ZJ7749	Enhanced MX SCB
CB 1	REV 14	750-031391	ZJ7750	Enhanced MX SCB
CB 2	REV 14	750-031391	ZY9233	Enhanced MX SCB
FPC 0	REV 17	750-031089	YR7434	MPC Type 2 3D
CPU				
FPC 1	REV 11	750-037207	ZW9727	AS-MCC
CPU	REV 04	711-038173	ZW4817	AS-MCC-PMB
MIC 0	REV 01	750-037214	ZH3764	AS-MSC
PIC 0		BUILTIN	BUILTIN	AS-MSC
MIC 1	REV 01	711-028408	JZ9200	AS-MXC
PIC 2		BUILTIN	BUILTIN	AS-MXC
FPC 4	REV 30	750-028467	ABBN0232	MPC 3D 16x 10GE
CPU				
FPC 5	REV 04	750-037207	ZK9074	AS-MCC
CPU				
Fan Tray 0	REV 05	740-014971	VT5683	Fan Tray
Fan Tray 1	REV 05	740-014971	VT5684	Fan Tray

show chassis hardware extensive (MX240, MX480, MX960 routers with Application Services Modular Line Card)

user@host> show chassis hardware extensive

```
ID: AS-MCC                      FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU                               REV 04    711-038173    ZW4817    AS-MCC-PMB
Jedec Code: 0x7fb0               EEPROM Version: 0x02
P/N: 711-038173                  S/N: S/N ZW4817
Assembly ID: 0x0b38              Assembly Version: 01.04
Date: 12-30-2011                 Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC-PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0                             REV 01    750-037214    ZH3764    AS-MSC
```

```

Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 750-037214        S/N: S/N ZH3764
Assembly ID: 0x0a44     Assembly Version: 01.01
Date: 07-04-2011       Assembly Flags: 0x00
Version: REV 01
ID: AS-MSC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 0f
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 ff ff ff ff ff ff
Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN      AS-MSC
FPC 4          REV 30       750-028467    ABBN0232    MPC 3D 16x 10GE
Jedec Code: 0x7fb0      EEPROM Version: 0x01

```

show chassis hardware (T320 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			19093	T320
Midplane	REV 04	710-004339	BC1436	T320 Backplane
FPM GBUS	REV 03	710-004461	BC1407	T320 FPM Board
FPM Display	REV 04	710-002897	BE0763	FPM Display
CIP	REV 05	710-002895	BB2311	T Series CIP
PEM 0	Rev 01	740-004359	NB12546	Power Entry Module
SCG 0	REV 06	710-004455	AY4522	T320 Sonet
Clock Gen.				
Routing Engine 0				unknown
CB 0	REV 13	710-002728	BC1577	T Series
Control Board				
CB 1	REV 13	710-002728	BC1595	T Series
Control Board				
FPC 1	REV 09	710-007531	HS1572	FPC Type 2
CPU	REV 15	710-001726	HR8763	FPC CPU
PIC 0	REV 01	750-010618	CB5579	4x G/E SFP,
1000 BASE				
SFP 0	REV 01	740-007326	P5809Z1	SFP-SX
SFP 1	REV 01	740-007326	P4Q10XU	SFP-SX
SFP 2		NON-JNPR	RA45020031	SFP-SX
SFP 3		NON-JNPR	RA45020032	SFP-SX
PIC 1	REV 01	750-010618	CD9587	4x G/E SFP,
1000 BASE				
SFP 0		NON-JNPR	P5A08QZ	SFP-T
SFP 1	REV 01	740-007326	P4Q133K	SFP-SX
SFP 2	REV 01	740-007326	P5809YY	SFP-SX
SFP 3	REV 01	740-007327	4C81704	SFP-LX
MMB 1	REV 03	710-005555	HR9401	MMB-288mbit
PPB 0	REV 04	710-003758	HR2886	PPB Type 2
FPC 2	REV 07	710-005860	HP2392	FPC Type 1
CPU	REV 14	710-001726	HP7797	FPC CPU
PIC 0	REV 02	750-007643	HM0853	1x G/E QPP,
1000 BASE				
SFP 0	REV 01	740-007326	P11E9JJ	SFP-SX
MMB 1	REV 02	710-005555	HN2379	MMB-288mbit
PPB 0	REV 04	710-003758	HP8092	PPB Type 2

FPC 3	REV 07	710-005860	HP2393	FPC Type 1
CPU	REV 14	710-001726	HP0968	FPC CPU
PIC 0	REV 01	750-010240	CB5363	1x G/E SFP,
1000 BASE				
SFP 0	REV 01	740-007326	P4R0PNH	SFP-SX
PIC 1	REV 03	750-003034	HD2832	4x OC-3 SONET,
SMIR				
MMB 1	REV 02	710-005555	HN6307	MMB-288mbit
PPB 0	REV 04	710-003758	HP5051	PPB Type 2
FPC 4	REV 01	710-010845	JD3872	FPC Type 4
CPU	REV 02	710-011481	JB6042	FPC CPU
5	REV 01	710-005802	BC1566	FPC Type 2
CPU	REV 09	710-001726	AY4922	FPC CPU
PIC 0	REV 02	750-008155	BE2114	2x G/E QPP,
1000 BASE				
SFP 0	REV 01	740-007326	P4R0PMQ	SFP-SX
SFP 1	REV 01	740-007326	P4R0PN9	SFP-SX
PIC 1	REV 01	750-008155	BE2116	2x G/E QPP,
1000 BASE				
SFP 0	REV 01	740-007326	P4R0PNZ	SFP-SX
SFP 1		NON-JNPR	2908	SFP-T
MMB 1	REV 01	710-005555	AZ2246	MMB-288mbit
PPB 0	REV 03	710-003758	AY4839	PPB Type 2
FPC 7	REV 01	710-005803	AZ2123	FPC Type 3
...				

show chassis hardware (T640 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			19182	T640
Midplane	REV 04	710-002726	AX5608	T640 Backplane
FPM GBUS	REV 02	710-002901	HE3064	T640 FPM Board
FPM Display	REV 02	710-002897	HE7864	FPM Display
CIP	REV 05	710-002895	HA5024	T Series CIP
PEM 0	Rev 02	740-029522	VH26235	AC PEM 10kW US
PEM 1	Rev 02	740-029522	VH26230	AC PEM 10kW US
SCG 0	REV 03	710-003423	HA4508	T640 Sonet Clock Gen.
Routing Engine 0	REV 02	740-005022	210865700483	RE-3.0 (RE-600)
CB 0	REV 01	710-002728	HD3044	T Series Control Board
FPC 2	REV 04	710-001721	HD5572	FPC Type 3
CPU	REV 06	710-001726	HA4712	FPC CPU
PIC 1	REV 03	750-009567	HV2331	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202R103	XENPAK-SR
PIC 2	REV 03	750-009567	HV2332	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-011268	USC202R112	XENPAK-ZR
PIC 3	REV 03	750-009567	HX4416	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012056	434TC004	XENPAK-CX4
PIC 4	REV 03	750-009567	HX4420	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012058	434TC124	XENPAK-LX4
FPC 5	REV 01	710-013553	JE4839	E2-FPC Type 1
CPU	REV 01	710-013569	JW9163	FPC CPU
PIC 0	REV 01	750-009567	HX4419	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202RT05	XENPAK-LR
PIC 1	REV 03	750-009567	HN7426	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009550	03L90051	XENPAK-ER
PIC 2	REV 03	750-009467	HT7423	1x 10GE(LAN),XENPAK
SFP 0		NON-JNPR		UNKNOWN
PIC 3	REV 04	750-005100	AY4850	1x 10GE(LAN),DWDM
FPC 4	REV 01	710-010845	JD3872	FPC Type 4

CPU	REV 02	710-011481	JB6042	FPC CPU
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (T640 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 04   710-002726
FPM Display   REV 02   710-002897
CIP           REV 05   710-002895
PEM 0         Rev 01   740-002595
SCG 0         REV 04   710-003423
SCG 1         REV 04   710-003423
Routing Engine 0 REV 01   740-005022
Routing Engine 1 REV 07   740-005022
CB 0          REV 06   710-002726
CB 1          REV 06   710-002728
FPC 5         REV 05   710-007527
  PIC 0       REV 05   750-002510
  PIC 1       REV 05   750-001901
FPC 6         REV 03   710-001721
  PIC 1       REV 01   750-009553
SIB 4         REV 02   750-005486
Fan Tray 0
Fan Tray 1
Fan Tray 2
```

CHAS-BP-T640-S
CRAFT-T640-S
CIP-L-T640-S
PWR-T-DC-S
SCG-T-S
SCG-T-S
RE-600-2048-S
RE-600-2048-S
CHAS-BP-T640-S
CB-L-T-S
T640-FPC2
PB-2GE-SX
PB-40C12-SON-SMIR
T640-FPC3
PC-40C48-SON-SFP
SIB-I-T640-S
FANTRAY-T-S
FANTRAY-T-S
FAN-REAR-TX-T640-S

show chassis hardware extensive (T640 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          .....        S/N:          .....
Assembly ID:  0x0507          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
Version:      .....
ID: Gibson LCC Chassis
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 04   710-002726  AX5633
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          710-002726.      S/N:          S/N AX5633.
Assembly ID:  0x0127          Assembly Version: 01.04
Date:         06-27-2001      Assembly Flags:  0x00
Version:      REV 04.....
ID: Gibson Backplane
Board Information Record:
Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
```

```

Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS          REV 02   710-002901   HE3245
...
FPM Display      REV 02   710-002897   HA4873
...
CIP              REV 05   710-002895   HA4729
...
PEM 1            RevX02   740-002595   MD21815      Power Entry Module
...
SCG 0            REV 04   710-003423   HF6023
...
SCG 1            REV 04   710-003423   HF6061
...
Routing Engine 0 REV 01   740-005022   210865700292 RE-3.0
...
CB 0             REV 06   710-002728   HE3614
...
FPC 1            REV 01   710-002385   HE3009      FPC Type 1
...
                REV 06   710-001726   HC0010

```

show chassis hardware (T4000 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1172F25AHA  T4000
Midplane      REV 01   710-027486   RC8355        T-series Backplane
FPM GBUS      REV 13   710-002901   BBAE0927      T640 FPM Board
FPM Display   REV 01   710-021387   EF6764        T1600 FPM Display
CIP           REV 06   710-002895   BBAD9210      T-series CIP
PEM 0         REV 01   740-036442   VA00016       Power Entry Module 6x60
SCG 0         REV 18   710-003423   BBAD7248      T640 Sonet Clock Gen.
SCG 1         REV 18   710-003423   BBAE3874      T640 Sonet Clock Gen.
Routing Engine 0 REV 05   740-026941   P737F-002248  RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002653  RE-DUO-1800
CB 0          REV 09   710-022597   ED0295        LCC Control Board
CB 1          REV 09   710-022597   EA6050        LCC Control Board
FPC 0         REV 26   750-032819   EK1173        FPC Type 5-3D
CPU           REV 12   711-030686   EJ8584        SNG PMB
PIC 0         REV 07   750-034624   EF6837        12x10GE (LAN/WAN) SFPP
  Xcvr 0      REV 01   740-031980   123363A01145  SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   123363A01147  SFP+-10G-SR
  Xcvr 2      REV 01   740-031980   AJJ01P3       SFP+-10G-SR
  Xcvr 3      REV 01   740-031980   B10M03256     SFP+-10G-SR
  Xcvr 4      REV 01   740-031980   AJJ01M2       SFP+-10G-SR
  Xcvr 5      REV 01   740-031980   123363A01137  SFP+-10G-SR
  Xcvr 6      REV 01   740-031980   AJJ01PN       SFP+-10G-SR
  Xcvr 7      REV 01   740-031980   AJJ01NW       SFP+-10G-SR
  Xcvr 8      REV 01   740-031980   123363A01139  SFP+-10G-SR
  Xcvr 9      REV 01   740-031980   AJJ01KE       SFP+-10G-SR
  Xcvr 10     REV 01   740-031980   123363A01336  SFP+-10G-SR
  Xcvr 11     REV 01   740-031980   B10M01325     SFP+-10G-SR
PIC 1         REV 07   750-034624   EF6800        12x10GE (LAN/WAN) SFPP
  Xcvr 0      REV 01   740-031980   AJJ01SA       SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   AJJ01QZ       SFP+-10G-SR
  Xcvr 2      REV 01   740-031980   AJH0217       SFP+-10G-SR
  Xcvr 3      REV 01   740-031980   AJJ01TE       SFP+-10G-SR
  Xcvr 4      REV 01   740-031980   AJJ01KV       SFP+-10G-SR

```

Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJD0GV3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR

Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware (T4000 Router with 16 GB line card chassis (LCC) Routing Engine)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11BDF2CAHA	T1600
Midplane	REV 01	710-027486	ACAJ0774	T640 Backplane
FPM GBUS	REV 13	710-002901	BBAL6812	T640 FPM Board
FPM Display	REV 04	710-021387	BBAP2679	T1600 FPM Display
CIP	REV 06	710-002895	BBAP4758	T-series CIP
PEM 0	Rev 03	740-026384	XF86421	Power Entry Module 3x80
PEM 1	Rev 03	740-026384	XF86429	Power Entry Module 3x80
SCG 0	REV 18	710-003423	BBAP1896	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAN8659	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-042243	737F-002238	RE-DUO-1800-16G
Routing Engine 1	REV 01	740-042243	737F-002403	RE-DUO-1800-16G
CB 1	REV 11	710-022597	EK4526	LCC Control Board
CB 1	REV 11	710-022597	EK4527	LCC Control Board
FPC 0	REV 05	710-033871	EK5644	FPC Type 4-ES
CPU	REV 11	710-016744	EK3428	ST-PMB2
PIC 0	REV 20	750-017405	EJ3041	4x 10GE (LAN/WAN) XFPP
PIC 1	REV 17	750-026962	EH7536	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6039	ST-MMB2
MMB 1	REV 07	710-025563	EK6086	ST-MMB2
FPC 1	REV 05	710-033871	EK6583	FPC Type 4-ES
CPU	REV 11	710-016744	EK3401	ST-PMB2
PIC 0	REV 17	750-026962	EJ8948	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6202	ST-MMB2
MMB 1	REV 07	710-025563	EK6112	ST-MMB2
SPMB 1	REV 05	710-023321	EK4900	LCC Switch CPU
SIB 0	REV 11	710-013074	EK5958	SIB-I8-SF
SIB 1	REV 11	710-013074	EK4606	SIB-I8-SF
SIB 2	REV 11	710-013074	EK5971	SIB-I8-SF
SIB 3	REV 11	710-013074	EK4609	SIB-I8-SF
SIB 4	REV 11	710-013074	EK4602	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

show chassis hardware (T4000 Router with LSR FPC)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis              REV 01  710-027486  JN1173A24AHA  T4000
FPC 3                REV 01  750-048373  AN7797        FPC Type 5-LSR
CPU                  REV 10  711-030686  AN6649        SNG PMB
PIC 0                REV 07  750-034624  EF6830        12x10GE (LAN/WAN) SFPP

```

show chassis hardware clei-models (T4000 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 01  710-027486  IPMJ700DRD CHAS-BP-T1600-S
FPM Display   REV 01  710-021387  CRAFT-T1600-S
CIP           REV 06  710-002895  CIP-L-T640-S
PEM 0         REV 01  740-036442  IPUPAG6KAA  PWR-T-6-60-DC
SCG 0         REV 18  710-003423  SCG-T-S
SCG 1         REV 18  710-003423  SCG-T-S
Routing Engine 0 REV 05  740-026941  RE-DUO-C1800-8G-S
Routing Engine 1 REV 06  740-026941  RE-DUO-C1800-8G-S
CB 0          REV 09  710-022597  CB-LCC-S
CB 1          REV 09  710-022597  CB-LCC-S
FPC 3
  PIC 0        REV 08  750-035293  XXXXXXXXBB PF-1CGE-CFP
  PIC 1        REV 10  750-034624  XXXXXXXXCC PF-12XGE-SFPP
FPC 5          REV 03  710-033871  IPUCAMBCTD T1600-FPC4-ES
  PIC 1        REV 03  750-034781  IPUIBKLMAA PD-1CE-CFP-FPC4
FPC 6
  PIC 0        REV 10  750-034624  XXXXXXXXCC PF-12XGE-SFPP
Fan Tray 0     FANTRAY-T-S
Fan Tray 1     FANTRAY-T4000-S
Fan Tray 2     FANTRAY-TXP-R-S

```

show chassis hardware detail (T4000 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis              REV 01  710-027486  JN1172F25AHA  T4000
Midplane            REV 01  710-027486  RC8355        T-series Backplane
FPM GBUS            REV 13  710-002901  BBAE0927      T640 FPM Board
FPM Display         REV 01  710-021387  EF6764        T1600 FPM Display
CIP                 REV 06  710-002895  BBAD9210      T-series CIP
PEM 0              REV 01  740-036442  VA00016       Power Entry Module 6x60
SCG 0              REV 18  710-003423  BBAD7248      T640 Sonet Clock Gen.
SCG 1              REV 18  710-003423  BBAE3874      T640 Sonet Clock Gen.
Routing Engine 0    REV 05  740-026941  P737F-002248  RE-DUO-1800
  ad0  3823 MB  SMART CF      2009121602A661576157 Compact Flash
  ad1  59690 MB STEC MACH-8 SSD  STM000103FDB    Disk 1
Routing Engine 1    REV 06  740-026941  P737F-002653  RE-DUO-1800
  ad0  3823 MB  SMART CF      201011150153F52CF52C Compact Flash
  ad1  62720 MB SMART Lite SATA Drive 2010110900150A880A88 Disk 1
CB 0                REV 09  710-022597  ED0295        LCC Control Board
CB 1                REV 09  710-022597  EA6050        LCC Control Board
FPC 0              REV 26  750-032819  EK1173        FPC Type 5-3D
CPU                REV 12  711-030686  EJ8584        SNG PMB
PIC 0              REV 07  750-034624  EF6837        12x10GE (LAN/WAN) SFPP
Xcvr 0             REV 01  740-031980  123363A01145 SFP+-10G-SR

```


Xcvr 1	REV 01	740-031980	123363A01147	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01P3	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M03256	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01M2	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	123363A01137	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01PN	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01NW	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	123363A01139	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01KE	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01336	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10M01325	SFP+-10G-SR
PIC 1	REV 07	750-034624	EF6800	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJJ01SA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01QZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJH0217	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ01TE	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01KV	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJD0GV3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware models (T4000 Router)

user@host> show chassis hardware models

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 01	710-027486	RC8355	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	EF6764	CRAFT-T1600-S
CIP	REV 06	710-002895	BBAD9210	CIP-L-T640-S
PEM 0	REV 01	740-036442	VA00016	PWR-T-6-60-DC
SCG 0	REV 18	710-003423	BBAD7248	SCG-T-S
SCG 1	REV 18	710-003423	BBAE3874	SCG-T-S
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-C1800-8G-S
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-C1800-8G-S
CB 0	REV 09	710-022597	ED0295	CB-LCC-S
CB 1	REV 09	710-022597	EA6050	CB-LCC-S
FPC 3				
PIC 0	REV 08	750-035293	EF3657	PF-1CGE-CFP
PIC 1	REV 10	750-034624	BBAN4098	PF-12XGE-SFPP
FPC 5	REV 03	710-033871	BBAJ0768	T1600-FPC4-ES
PIC 1	REV 03	750-034781	EE6655	PD-1CE-CFP-FPC4
FPC 6				
PIC 0	REV 10	750-034624	BBAN4109	PF-12XGE-SFPP
Fan Tray 0				FANTRAY-T-S

Fan Tray 1
Fan Tray 2

FANTRAY-T4000-S
FAN-REAR-TXP-LCC

show chassis hardware lcc (TX Matrix Router)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			65751	T640
Midplane	REV 03	710-005608	RA1408	T640 Backplane
FPM GBUS	REV 09	710-002901	RA2784	T640 FPM Board
FPM Display	REV 05	710-002897	RA2825	FPM Display
CIP	REV 06	710-002895	HT0684	T Series CIP
PEM 0	Rev 11	740-002595	PM18483	Power Entry Module
PEM 1	Rev 11	740-002595	qb13984	Power Entry Module
SCG 0	REV 11	710-003423	HT0022	T640 Sonet Clock Gen.
Routing Engine 0	REV 13	740-005022	210865700363	RE-3.0 (RE-600)
CB 0	REV 03	710-007655	HW1195	Control Board (CB-T)
FPC 1	REV 05	710-007527	HM3245	FPC Type 2
CPU	REV 14	710-001726	HM1084	FPC CPU
PIC 0	REV 02	750-007218	AZ1112	2x OC-12 ATM2 IQ, SMIR
PIC 1	REV 02	750-007745	HG3462	4x OC-3 SONET, SMIR
PIC 2	REV 14	750-001901	BA5390	4x OC-12 SONET, SMIR
PIC 3	REV 09	750-008155	HS3012	2x G/E IQ, 1000 BASE
SFP 0		NON-JNPR	P1186TY	SFP-S
SFP 1	REV 01	740-007326	P11WLTf	SFP-SX
MMB 1	REV 02	710-005555	HL7514	MMB-288mbit
PPB 0	REV 04	710-003758	HM4405	PPB Type 2
PPB 1	REV 04	710-003758	AV1960	PPB Type 2
FPC 2	REV 08	710-010154	HZ3578	E-FPC Type 3
CPU	REV 05	710-010169	HZ3219	FPC CPU-Enhanced
PIC 0	REV 02	750-009567	HX2882	1x 10GE(LAN), XENPAK
SFP 0	REV 01	740-009898	USC202U709	XENPAK-LR
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 2	REV 01	750-004535	HC0235	1x OC-192 SM SR1
PIC 3	REV 07	750-007141	HX1699	10x 1GE(LAN), 1000 BASE
SFP 0	REV 01	740-007326	2441042	SFP-SX
SFP 1	REV 01	740-007326	2441027	SFP-SX
MMB 0	REV 03	710-010171	HV2365	MMB-5M3-288mbit
MMB 1	REV 03	710-010171	HZ3888	MMB-5M3-288mbit
SPMB 0	REV 09	710-003229	HW5245	T Series Switch CPU
SIB 3	REV 07	710-005781	HR5927	SIB-L8-F16
B Board	REV 06	710-005782	HR5971	SIB-L8-F16 (B)
SIB 4	REV 07	710-005781	HR5903	SIB-L8-F16
B Board	REV 06	710-005782	HZ5275	SIB-L8-F16 (B)

show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				TX Matrix
Midplane	REV 04	710-004396	RB0014	SCC Midplane
FPM GBUS	REV 04	710-004617	HW9141	SCC FPM Board
FPM Display	REV 04	710-004619	HS5950	SCC FPM
CIP 0	REV 01	710-010218	HV9151	SCC CIP

CIP 1	REV 01	710-010218	HV9152	SCC CIP
PEM 1	Rev 11	740-002595	QB13977	Power Entry Module
Routing Engine 0	REV 05	740-008883	P11123900153	RE-4.0 (RE-1600)
CB 0	REV 01	710-011709	HR5964	Control Board (CB-TX)
SPMB 0	REV 09	710-003229	HW5293	T Series Switch CPU
SIB 3				
SIB 4	REV 01	710-005839	HW1177	SIB-S8-F16
B Board	REV 01	710-005840	HW1202	SIB-S8-F16 (B)

show chassis hardware (T1600 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              B2703      T1600
Midplane             REV 03     710-005608   RC4137         T640 Backplane
FPM GBUS             REV 10     710-002901   DT7062         T640 FPM Board
FPM Display          REV 05     710-002897   DS3067         FPM Display
CIP                  REV 06     710-002895   DT3386         T-series CIP
PEM 0                Rev 07     740-017906   UA26344        Power Entry Module 3x80
PEM 1                Rev 18     740-002595   UF38441        Power Entry Module
SCG 0                REV 15     710-003423   DV0941         T640 Sonet Clock Gen.
Routing Engine 0     REV 08     740-014082   9009014502     RE-A-2000
Routing Engine 1     REV 07     740-014082   9009009591     RE-A-2000
CB 0                 REV 05     710-007655   JA9360         Control Board (CB-T)
CB 1                 REV 03     710-017707   DT3251         Control Board (CB-T)
FPC 0                REV 07     710-013558   DR4253         E2-FPC Type 2
CPU                  REV 05     710-013563   DS3902         FPC CPU-Enhanced
PIC 0                REV 01     750-010618   CB5446         4x G/E SFP, 1000 BASE
Xcvr 0               REV 01     740-011613   P9F11CW        SFP-SX
Xcvr 1               REV 01     740-011613   P9F15C2        SFP-SX
Xcvr 2               REV 01     740-011782   PB94K0L        SFP-SX
PIC 1                REV 06     750-001900   HB6399         1x OC-48 SONET, SMSR
PIC 2                REV 14     750-001901   AP1092         4x OC-12 SONET, SMIR
PIC 3                REV 07     750-001900   AR8275         1x OC-48 SONET, SMSR
MMB 1                REV 07     710-010171   DS1524         MMB-5M3-288mbit
FPC 1                REV 06     710-013553   DL9067         E2-FPC Type 1
CPU                  REV 04     710-013563   DM1685         FPC CPU-Enhanced
PIC 0                REV 08     750-001072   AB1688         1x G/E, 1000 BASE-SX
PIC 1                REV 10     750-012266   JX5519         4x 1GE(LAN), IQ2
Xcvr 0               REV 01     740-011613   AM0812S8UK6    SFP-SX
Xcvr 2               REV 01     740-011613   AM0812S8UK1    SFP-SX
Xcvr 3               REV 01     740-011782   P8N1YHG        SFP-SX
PIC 2                REV 22     750-005634   DP0083         1x CHOC12 IQ SONET, SMIR

MMB 1                REV 07     710-008923   DN1862         MMB 3M 288-bit
FPC 2                REV 01     710-005548   HJ9899         FPC Type 3
CPU                  REV 06     710-001726   HC0586         FPC CPU
PIC 0                REV 16     750-007141   NC9660         10x 1GE(LAN), 1000 BASE

Xcvr 0               REV 01     740-011613   AM0812S8XAR    SFP-SX
Xcvr 1               REV 01     740-011782   P920E7B        SFP-SX
Xcvr 2               REV 01     740-011613   AM0812S8XAU    SFP-SX
Xcvr 4               REV 01     740-011613   AM0812S8XAK    SFP-SX
Xcvr 5               REV 01     740-011613   AM0812S8XAA    SFP-SX
Xcvr 6               REV 01     740-011613   PAJ4NKY        SFP-SX
Xcvr 7               REV 01     740-011613   AM0812S8UJW    SFP-SX
Xcvr 8               REV 01     740-011782   PB81X89        SFP-SX
Xcvr 9               REV 01     740-011613   AM0812S8UJX    SFP-SX
PIC 1                REV 06     750-015217   DK3280         8x 1GE(TYPE3), IQ2
Xcvr 0               REV 01     740-011782   P8P0A3T        SFP-SX

```

Xcvr 1	REV 01	740-013111	5090002	SFP-T
Xcvr 2	REV 01	740-011613	AM0814S93BQ	SFP-SX
Xcvr 4		NON-JNPR	PDE0FAN	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q20XY	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8UJV	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UP7	SFP-SX
PIC 2	REV 05	750-004695	HT4383	1x Tunnel
PIC 3	REV 17	750-009553	RL0204	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T23	SFP-SR
Xcvr 1	REV 01	740-011785	P6Q0F3E	SFP-SR
MMB 0	REV 03	710-004047	HD5843	MMB-288mbit
MMB 1	REV 03	710-004047	HE3208	MMB-288mbit
PPB 0	REV 02	710-002845	HA4524	PPB Type 3
PPB 1	REV 02	710-002845	HA4766	PPB Type 3
FPC 3	REV 01	710-010154	HR0863	E-FPC Type 3
CPU	REV 01	710-010169	HN3422	FPC CPU-Enhanced
PIC 0	REV 07	750-012793	WF5096	1x 10GE(LAN/WAN) IQ2
Xcvr 0		NON-JNPR	M64294TP	XFP-10G-LR
PIC 1	REV 25	750-007141	DV2127	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	PFA6LTJ	SFP-SX
Xcvr 1	REV 01	740-011782	P9POXV4	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TNX	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0TTP	SFP-SX
Xcvr 5		NON-JNPR	PBS4LED	SFP-SX
PIC 2	REV 17	750-009553	RL0212	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T8G	SFP-SR
PIC 3	REV 32	750-003700	DL1279	1x OC-192 12xMM VSR
MMB 0	REV 01	710-010171	HR0821	MMB-288mbit
MMB 1	REV 01	710-010171	HR0818	MMB-288mbit
FPC 4	REV 16	710-013037	EB4919	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA4382	ST-PMB2
PIC 0	REV 03	711-029996	EB1569	100GE
PIC 1	REV 05	711-029999	EB9983	100GE CFP
Xcvr 0	REV 0	740-032210	J10G80746	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2235	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA7112	ST-MMB2
MMB 1	REV 04	710-025563	BBAA7149	ST-MMB2
FPC 5	REV 02	710-013037	DE3407	FPC Type 4-ES
CPU	REV 04	710-016744	DA2124	ST-PMB2
PIC 0	REV 16	750-012518	DF2554	4x OC-192 SONET XFP
Xcvr 0	REV 01	740-014279	AA0745N1FX8	XFP-OC192-SR
Xcvr 1	REV 01	740-014279	AA0748N1HN5	XFP-OC192-SR
Xcvr 2	REV 01	740-014279	AA0748N1HT6	XFP-OC192-SR
Xcvr 3	REV 01	740-014279	AA0744N1EC9	XFP-OC192-SR
PIC 1	REV 01	750-010850	JA0329	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DE9577	ST-MMB2
MMB 1	REV 04	710-016036	DK4060	ST-MMB2
FPC 6	REV 14	710-013037	DV1431	FPC Type 4-ES
CPU	REV 09	710-016744	DT9020	ST-PMB2
PIC 0	REV 11	750-017405	DM6261	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014289	C701XU05Q	XFP-10G-SR
Xcvr 1	REV 01	740-014279	AA0748N1HPT	XFP-10G-LR
Xcvr 2	REV 01	740-014289	T08E19189	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C715XU058	XFP-10G-SR
PIC 1	REV 13	750-017405	DP8772	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 02	740-011571	C850XJ037	XFP-10G-SR
Xcvr 1	REV 02	740-014289	C839XU0L9	XFP-10G-SR
Xcvr 2	REV 02	740-014289	C834XU05A	XFP-10G-SR
Xcvr 3	REV 02	740-014289	C810XU0CE	XFP-10G-SR
MMB 0	REV 01	710-025563	DT8454	ST-MMB2

MMB 1	REV 01	710-025563	DT8366	ST-MMB2
FPC 7	REV 09	710-007529	HZ7624	FPC Type 3
CPU	REV 15	710-001726	HZ1413	FPC CPU
PIC 0	REV 10	750-012793	DM5627	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 02	740-011571	C831XJ062	XFP-10G-SR
PIC 1	REV 01	750-015217	JT6762	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q25JU	SFP-SX
Xcvr 1	REV 01	740-011782	P9B0U0K	SFP-SX
PIC 2	REV 01	750-015217	JS4268	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8XBZ	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAP	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XBY	SFP-SX
Xcvr 3	REV 01	740-011613	AM0812S8XBX	SFP-SX
Xcvr 4	REV 01	740-011613	P9F1652	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q21YC	SFP-SX
Xcvr 6	REV 01	740-011782	P8Q27HQ	SFP-SX
Xcvr 7	REV 01	740-011613	P8E2SSU	SFP-SX
PIC 3	REV 15	750-009450	NB6790	1x OC-192 SM SR2
MMB 0	REV 03	710-005555	HZ3450	MMB-288mbit
MMB 1	REV 03	710-005555	HZ3415	MMB-288mbit
PPB 0	REV 04	710-002845	HP0887	PPB Type 3
PPB 1	REV 04	710-002845	HW5255	PPB Type 3
SPMB 0	REV 10	710-003229	HX3699	T-series Switch CPU
SPMB 1	REV 12	710-003229	DT3091	T-series Switch CPU
SIB 0	REV 07	710-013074	DS4747	SIB-I8-SF
SIB 1	REV 07	710-013074	DS4942	SIB-I8-SF
SIB 2	REV 07	710-013074	DS4965	SIB-I8-SF
SIB 3	REV 07	710-013074	DS4990	SIB-I8-SF
SIB 4	REV 07	710-013074	DS4944	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

show chassis hardware (TX Matrix Plus Router)

user@host> show chassis hardware

sfc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN113186EAHB	TXP
Midplane	REV 05	710-022574	TS3822	SFC Midplane
FPM Display	REV 03	710-024027	DW4701	TXP FPM Display
CIP 0	REV 05	710-023792	DW7998	TXP CIP
CIP 1	REV 05	710-023792	DW7999	TXP CIP
PEM 0	Rev 04	740-027463	UM26367	Power Entry Module
PEM 1	Rev 04	740-027463	UM26346	Power Entry Module
Routing Engine 0	REV 06	740-026942	737A-1081	RE-DUO-2600
Routing Engine 1	REV 06	740-026942	737A-1043	RE-DUO-2600
CB 0	REV 05	710-022606	DW4435	SFC Control Board
CB 1	REV 09	710-022606	DW6100	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	750-024564	DW5764	F13 SIB
B Board	REV 03	710-023431	DW9053	F13 SIB Mezz
SIB F13 3	REV 04	750-024564	DW5785	F13 SIB
B Board	REV 03	710-023431	DW9030	F13 SIB Mezz
SIB F13 6				
SIB F13 8	REV 04	750-024564	DW5752	F13 SIB
B Board	REV 03	710-023431	DW9051	F13 SIB Mezz
SIB F13 11	REV 04	750-024564	DW5782	F13 SIB

B Board	REV 03	710-023431	DW9058	F13 SIB Mezz
SIB F13 12	REV 03	750-024564	DT9466	F13 SIB
B Board	REV 02	710-023431	DT6556	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7898	F2S SIB
B Board	REV 05	710-023787	DW7625	F2S SIB Mezz
SIB F2S 0/2	REV 05	710-022603	DW7811	F2S SIB
B Board	REV 05	710-023787	DW7550	F2S SIB Mezz
SIB F2S 0/4	REV 04	710-022603	DW4873	F2S SIB
B Board	REV 05	710-023787	DW8509	F2S SIB Mezz
SIB F2S 0/6	REV 04	710-022603	DW4867	F2S SIB
B Board	REV 05	710-023787	DW8472	F2S SIB Mezz
SIB F2S 1/0	REV 04	710-022603	DW4871	F2S SIB
B Board	REV 05	710-023787	DW8497	F2S SIB Mezz
SIB F2S 1/2	REV 05	710-022603	DW7868	F2S SIB
B Board	REV 05	710-023787	DW7551	F2S SIB Mezz
SIB F2S 1/4	REV 04	710-022603	DW4854	F2S SIB
B Board	REV 05	710-023787	DW8496	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7889	F2S SIB
B Board	REV 05	710-023787	DW7496	F2S SIB Mezz
SIB F2S 2/0	REV 04	710-022603	DW4852	F2S SIB
B Board	REV 05	710-023787	DW8498	F2S SIB Mezz
SIB F2S 2/2	REV 04	710-022603	DW4845	F2S SIB
B Board	REV 05	710-023787	DW8457	F2S SIB Mezz
SIB F2S 2/4	REV 05	710-022603	DW7802	F2S SIB
B Board	REV 05	710-023787	DW7562	F2S SIB Mezz
SIB F2S 2/6	REV 04	710-022603	DW4822	F2S SIB
B Board	REV 05	710-023787	DW8467	F2S SIB Mezz
SIB F2S 3/0	REV 05	710-022603	DW7815	F2S SIB
B Board	REV 05	710-023787	DW7518	F2S SIB Mezz
SIB F2S 3/2	REV 03	710-022603	DV0068	F2S SIB
B Board	REV 03	710-023787	DT9974	F2S SIB Mezz
SIB F2S 3/4	REV 05	710-022603	DW7874	F2S SIB
B Board	REV 05	710-023787	DW7601	F2S SIB Mezz
SIB F2S 3/6	REV 03	710-022603	DV0033	F2S SIB
B Board	REV 03	710-023787	DT9969	F2S SIB Mezz
SIB F2S 4/0	REV 03	710-022603	DV0043	F2S SIB
B Board	REV 03	710-023787	DT9948	F2S SIB Mezz
SIB F2S 4/2	REV 05	710-022603	DW5446	F2S SIB
B Board	REV 05	710-023787	DW7611	F2S SIB Mezz
SIB F2S 4/4	REV 04	710-022603	DW4826	F2S SIB
B Board	REV 05	710-023787	DW8458	F2S SIB Mezz
SIB F2S 4/6	REV 03	710-022603	DV0026	F2S SIB
B Board	REV 03	710-023787	DT9963	F2S SIB Mezz
Fan Tray 0	REV 02	760-024497	DR8290	Front Fan Tray
Fan Tray 1	REV 02	760-024497	DR8293	Front Fan Tray
Fan Tray 2	REV 05	760-024502	DR8280	Rear Fan Tray
Fan Tray 3				
Fan Tray 4	REV 05	760-024502	DR8276	Rear Fan Tray
Fan Tray 5	REV 02	760-024502	DP5643	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11036F8AHA	T1600
Midplane	REV 03	710-017247	RC3799	T-series Backplane
FPM GBUS	REV 10	710-002901	DP7009	T640 FPM Board
FPM Display	REV 01	710-021387	DN7026	T1600 FPM Display
CIP	REV 06	710-002895	DP6024	T-series CIP
PEM 1	Rev 02	740-023211	WA50019	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DR6757	T640 Sonet Clock Gen.

SCG 1	REV 15	710-003423	DS2225	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1040	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1016	RE-DUO-1800
CB 0	REV 06	710-022597	DX4011	LCC Control Board
CB 1	REV 06	710-022597	DX4017	LCC Control Board
FPC 1	REV 07	710-013035	DN5847	FPC Type 3-ES
CPU	REV 08	710-016744	DP2570	ST-PMB2
PIC 0	REV 05	750-015217	DB0418	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q27ZG	SFP-SX
Xcvr 1		NON-JNPR	PDA1U0D	SFP-SX
Xcvr 2	REV 01	740-011613	P9F1ALW	SFP-SX
Xcvr 3	REV 01	740-011782	PBA403V	SFP-SX
Xcvr 4		NON-JNPR	PDE09DP	SFP-SX
Xcvr 5	REV 01	740-011782	PCH2P4K	SFP-SX
Xcvr 6	REV 01	740-011782	PB94K0F	SFP-SX
Xcvr 7	REV 01	740-011782	PBA2R2A	SFP-SX
PIC 1	REV 03	750-004424	HJ4020	1x 10GE(LAN), DWDM
PIC 2	REV 01	750-003336	HG6073	4x OC-48 SONET, SMR
MMB 0	REV 04	710-016036	DP3401	ST-MMB2
FPC 3	REV 12	710-013037	DR1169	FPC Type 4-ES
CPU	REV 08	710-016744	DP9429	ST-PMB2
PIC 0	REV 02	750-010850	JA0332	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DR0628	ST-MMB2
MMB 1	REV 04	710-016036	DR0592	ST-MMB2
FPC 4	REV 05	710-021534	DR7350	FPC Type 1-ES
CPU	REV 08	710-016744	DP8096	ST-PMB2
PIC 0	REV 04	750-014627	DP9171	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 02	740-011615	PDE2RVR	SFP-SR
PIC 1	REV 22	750-005634	DS5815	1x CHOC12 IQ SONET, SMIR
PIC 2	REV 09	750-002911	CF4539	4x F/E, 100 BASE-TX
PIC 3	REV 08	750-021652	DR2827	1x CHOC12 IQE SONET
Xcvr 0		NON-JNPR	8	UNKNOWN
MMB 0	REV 04	710-016036	DR0809	ST-MMB2
FPC 5	REV 07	710-007529	HS5608	FPC Type 3
CPU	REV 15	710-001726	HX4351	FPC CPU
PIC 0	REV 14	750-009567	WJ8961	1x 10GE(LAN), XENPAK
Xcvr 0	REV 01	740-013170	J05K05961	XENPAK-LR
PIC 1	REV 16	750-007141	JJ8146	10x 1GE(LAN), 1000 BASE
Xcvr 1	REV 01	740-011613	P9F117T	SFP-SX
Xcvr 2	REV 01	740-011782	PBA2VCL	SFP-SX
Xcvr 3	REV 01	740-011782	PB83DRB	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8UP8	SFP-SX
PIC 2	REV 12	750-009567	WF3566	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T07C94489	XENPAK-LR
MMB 0	REV 03	710-005555	HZ1907	MMB-288mbit
MMB 1	REV 03	710-005555	HW5283	MMB-288mbit
PPB 0	REV 04	710-002845	HZ7717	PPB Type 3
PPB 1	REV 04	710-002845	HS0110	PPB Type 3
FPC 6	REV 07	710-013035	DP7486	FPC Type 3-ES
CPU	REV 08	710-016744	DP2545	ST-PMB2
PIC 0	REV 09	750-009567	NE6323	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T09C71959	XENPAK-LR
PIC 1	REV 06	750-015217	DN4775	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P7E0T6M	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAY	SFP-SX
Xcvr 2	REV 01	740-011782	P7E0T6J	SFP-SX
Xcvr 3	REV 01	740-011782	PCH2P7D	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0QYT	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8WQJ	SFP-SX

Xcvr 6	REV 02	740-013111	9301220	SFP-T
Xcvr 7	REV 01	740-011782	P9B0TZ5	SFP-SX
PIC 2	REV 06	750-015217	DM6747	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	PAP0ZB2	SFP-SX
Xcvr 1	REV 01	740-013111	70191002	SFP-T
Xcvr 6	REV 01	740-011782	PBA29H8	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8WQG	SFP-SX
MMB 0	REV 04	710-016036	DP3238	ST-MMB2
FPC 7	REV 03	710-021540	DV3154	FPC Type 2-ES
CPU	REV 09	710-016744	DT9053	ST-PMB2
PIC 0	REV 13	750-001901	HB4225	4x OC-12 SONET, SMIR
PIC 1	REV 05	750-001900	AD3644	1x OC-48 SONET, SMSR
PIC 2	REV 10	750-008155	HV0335	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011782	PCH2UKF	SFP-SX
Xcvr 1	REV 01	740-011782	PCH2V19	SFP-SX
PIC 3	REV 03	750-014638	JS9493	1x OC-48-12-3 SFP
Xcvr 0	REV 01	740-011785	P6Q0ENK	SFP-SR
MMB 0	REV 05	710-016036	DP3323	ST-MMB2
SPMB 0	REV 04	710-023321	DX3004	LCC Switch CPU
SPMB 1	REV 04	710-023321	DX3009	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4195	LCC SIB
B Board	REV 07	710-023185	DW3930	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4179	LCC SIB
B Board	REV 07	710-023185	DW3919	LCC SIB Mezz
SIB 2				
SIB 3	REV 06	710-022594	DT8251	LCC SIB
B Board	REV 06	710-023185	DT5792	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8014	LCC SIB
B Board	REV 07	710-023185	DW3917	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1102270AHA	T1600
Midplane	REV 04	710-017247	RC5358	T-series Backplane
FPM GBUS	REV 10	710-002901	DS3443	T640 FPM Board
FPM Display	REV 01	710-021387	DS6411	T1600 FPM Display
CIP	REV 06	710-002895	DS4235	T-series CIP
PEM 0	Rev 02	740-023211	VM82438	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DS6649	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR6775	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1083	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1104	RE-DUO-1800
CB 0	REV 06	710-022597	DW8542	LCC Control Board
CB 1	REV 06	710-022597	DW8530	LCC Control Board
FPC 0	REV 02	710-010845	JE2392	FPC Type 4
CPU	REV 02	710-011481	JF6820	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP7259	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	AA0741N1C8T	XFP-10G-LR
Xcvr 1	REV 01	740-014279	AA0746N1GAM	XFP-10G-LR
Xcvr 2	REV 01	740-014279	AA0747N1H0B	XFP-10G-LR
Xcvr 3	REV 01	740-014279	AA0748N1HZ5	XFP-10G-LR
MMB 0	REV 03	710-010842	HY7601	ST-MMB
FPC 1	REV 16	710-013037	BBAA7398	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA2329	ST-PMB2
PIC 0	REV 03	711-029996	EB1575	100GE
PIC 1	REV 06	750-034781	EB9980	100GE CFP

MMB 0	REV 04	710-025563	BBAA5325	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5444	ST-MMB2
FPC 2	REV 16	710-013037	BBAA7185	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA3522	ST-PMB2
PIC 0	REV 03	711-029996	EB1557	100GE
PIC 1	REV 05	750-034781	EB4660	100GE CFP
Xcvr 0	REV 0	740-032210	J10F73666	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2237	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA5347	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5401	ST-MMB2
FPC 3	REV 10	710-021534	DZ0941	FPC Type 1-ES
CPU	REV 09	710-016744	DY6364	ST-PMB2
PIC 0	REV 13	750-012266	DK9192	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8WVD	SFP-SX
Xcvr 1		NON-JNPR	PDD63Q4	SFP-SX
Xcvr 2		NON-JNPR	PDE4G54	SFP-SX
Xcvr 3		NON-JNPR	PD40MAG	SFP-SX
PIC 1	REV 01	750-007641	HJ2003	1x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8WVG	SFP-SX
PIC 3	REV 17	750-007444	JB6873	1x CHSTM1 IQ SDH, SMIR
MMB 0	REV 04	710-025563	DZ0281	ST-MMB2
FPC 4	REV 06	710-013035	DK0614	FPC Type 3-ES
CPU	REV 07	710-016744	DK1616	ST-PMB2
PIC 0	REV 22	750-007141	DM1870	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	PCL3UKW	SFP-SX
Xcvr 1	REV 01	740-011782	P7E0T73	SFP-SX
Xcvr 2	REV 01	740-007326	P4TOWLR	SFP-SX
Xcvr 3	REV 01	740-011782	PAR1LRL	SFP-SX
Xcvr 4	REV 01	740-011782	P9M0U3Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9M0U0C	SFP-SX
Xcvr 6	REV 01	740-011782	P9M0TLG	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0U0F	SFP-SX
Xcvr 8	REV 01	740-011613	PFA6LAP	SFP-SX
Xcvr 9	REV 01	740-011782	PCH2POU	SFP-SX
PIC 1	REV 16	750-009450	CV2565	1x OC-192 SM SR2
PIC 2	REV 05	750-004424	HH3057	1x 10GE(LAN), 10GBASE-LR
PIC 3	REV 12	750-013423	DP0403	MultiServices 500
MMB 0	REV 04	710-016036	DK1988	ST-MMB2
FPC 5	REV 07	710-013560	DR0004	E2-FPC Type 3
CPU	REV 05	710-013563	DR0089	FPC CPU-Enhanced
PIC 0	REV 11	750-012793	DR6107	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 01	740-014289	C743XU074	XFP-10G-SR
PIC 1	REV 01	750-004695	HD5980	1x Tunnel
PIC 2	REV 32	750-003700	DL3770	1x OC-192 12xMM VSR
PIC 3	REV 12	750-009553	WB8901	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	P9D1GTQ	SFP-SR
Xcvr 1	REV 01	740-011785	PDSOMMB	SFP-SR
Xcvr 3	REV 01	740-011785	PDE1KXP	SFP-SR
MMB 0	REV 07	710-010171	DP7374	MMB-5M3-288mbit
MMB 1	REV 07	710-010171	DP7404	MMB-5M3-288mbit
FPC 6	REV 07	710-013035	DM0994	FPC Type 3-ES
CPU	REV 07	710-016744	DM3651	ST-PMB2
PIC 0	REV 07	750-015217	DN4743	8x 1GE(TYPE3), IQ2
Xcvr 3	REV 01	740-011613	AM0812S8XB0	SFP-SX
Xcvr 4	REV 01	740-011782	PB829RB	SFP-SX
Xcvr 5	REV 01	740-011782	P8J1SYX	SFP-SX
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 3	REV 02	750-012793	JM7665	1x 10GE(LAN/WAN) IQ2
MMB 0	REV 04	710-016036	DN6913	ST-MMB2

FPC 7	REV 08	710-010845	JM3958	FPC Type 4
CPU	REV 04	710-011481	JK3669	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP8837	4x 10GE (LAN/WAN) XFP
Xcvr 1	REV 01	740-014279	753019A00277	XFP-10G-LR
Xcvr 2	REV 02	740-011571	C850XJ00P	XFP-10G-SR
Xcvr 3	REV 01	740-014279	AA0813N1RTG	XFP-10G-LR
MMB 0	REV 04	710-010842	JN1971	ST-MMB
SPMB 0	REV 04	710-023321	DW3629	LCC Switch CPU
SPMB 1	REV 04	710-023321	DW3621	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4200	LCC SIB
B Board	REV 07	710-023185	DW3932	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4193	LCC SIB
B Board	REV 07	710-023185	DW3904	LCC SIB Mezz
SIB 2				
SIB 3	REV 07	710-022594	DW4210	LCC SIB
B Board	REV 06	710-023185	DT5780	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8019	LCC SIB
B Board	REV 06	710-023185	DT5795	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP
Midplane	REV 05	710-022574	TS4027	SFC Midplane
FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
CIP 0	REV 04	710-023792	DW4889	TXP CIP
CIP 1	REV 04	710-023792	DW4887	TXP CIP
PEM 0	Rev 07	740-027463	UM26368	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1064	SFC RE
Routing Engine 1	REV 01	740-026942	737A-1082	SFC RE
CB 0	REV 09	710-022606	DW6099	SFC Control Board
CB 1	REV 09	710-022606	DW6096	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	710-022600	DX0841	F13 SIB
B Board	REV 03	710-023431	DX0966	F13 SIB Mezz
SIB F13 1	REV 04	750-024564	DW5776	F13 SIB
B Board	REV 03	710-023431	DW9028	F13 SIB
SIB F13 3	REV 04	750-024564	DW5762	F13 SIB
B Board	REV 03	710-023431	DW9059	F13 SIB
SIB F13 4	REV 04	750-024564	DW5797	F13 SIB
B Board	REV 03	710-023431	DW9041	F13 SIB
SIB F13 6	REV 04	750-024564	DW5770	F13 SIB
B Board	REV 03	710-023431	DW9079	F13 SIB Mezz
SIB F13 7	REV 04	750-024564	DW5758	F13 SIB
B Board	REV 03	710-023431	DW9047	F13 SIB
SIB F13 8	REV 04	750-024564	DW5761	F13 SIB
B Board	REV 03	710-023431	DW9043	F13 SIB Mezz
SIB F13 9	REV 04	750-024564	DW5754	F13 SIB
B Board	REV 03	710-023431	DW9078	F13 SIB Mezz
SIB F13 11	REV 04	710-022600	DX0826	F13 SIB
B Board	REV 03	710-023431	DX0967	F13 SIB Mezz
SIB F13 12	REV 04	750-024564	DW5794	F13 SIB
B Board	REV 03	710-023431	DW9044	F13 SIB Mezz

SIB F2S 0/0	REV 05	710-022603	DW7897	F2S SIB
B Board	REV 05	710-023787	DW7657	NEO PMB
SIB F2S 0/2	REV 05	710-022603	DW7833	F2S SIB
B Board	REV 05	710-023787	DW7526	NEO PMB
SIB F2S 0/4	REV 05	710-022603	DW7875	F2S SIB
B Board	REV 05	710-023787	DW7588	NEO PMB
SIB F2S 0/6	REV 05	710-022603	DW7860	F2S SIB
B Board	REV 05	710-023787	DW7589	NEO PMB
SIB F2S 1/0	REV 04	710-022603	DW4820	F2S SIB
B Board	REV 05	710-023787	DW8510	NEO PMB
SIB F2S 1/2	REV 05	710-022603	DW7849	F2S SIB
B Board	REV 05	710-023787	DW7525	NEO PMB
SIB F2S 1/4	REV 05	710-022603	DW7927	F2S SIB
B Board	REV 05	710-023787	DW7556	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7866	F2S SIB
B Board	REV 05	710-023787	DW7651	NEO PMB
SIB F2S 2/0	REV 05	710-022603	DW7880	F2S SIB
B Board	REV 05	710-023787	DW7523	NEO PMB
SIB F2S 2/2	REV 05	710-022603	DW7895	F2S SIB
B Board	REV 05	710-023787	DW7591	NEO PMB
SIB F2S 2/4	REV 05	710-022603	DW7907	F2S SIB
B Board	REV 05	710-023787	DW7590	NEO PMB
SIB F2S 2/6	REV 05	710-022603	DW7785	F2S SIB
B Board	REV 05	710-023787	DW7524	NEO PMB
SIB F2S 3/0	REV 05	710-022603	DW7782	F2S SIB
B Board	REV 05	710-023787	DW7634	NEO PMB
SIB F2S 3/2	REV 05	710-022603	DW7793	F2S SIB
B Board	REV 05	710-023787	DW7548	NEO PMB
SIB F2S 3/4	REV 05	710-022603	DW7779	F2S SIB
B Board	REV 05	710-023787	DW7587	NEO PMB
SIB F2S 3/6	REV 05	710-022603	DW7930	F2S SIB
B Board	REV 05	710-023787	DW7505	NEO PMB
SIB F2S 4/0	REV 05	710-022603	DW7867	F2S SIB
B Board	REV 05	710-023787	DW7656	NEO PMB
SIB F2S 4/2	REV 05	710-022603	DW7917	F2S SIB
B Board	REV 05	710-023787	DW7640	NEO PMB
SIB F2S 4/4	REV 05	710-022603	DW7929	F2S SIB
B Board	REV 05	710-023787	DW7643	NEO PMB
SIB F2S 4/6	REV 05	710-022603	DW7870	F2S SIB
B Board	REV 05	710-023787	DW7635	NEO PMB
Fan Tray 0	REV 06	760-024497	DV7831	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9614	Front Fan Tray
Fan Tray 2	REV 06	760-024502	DV9618	Rear Fan Tray
Fan Tray 3	REV 06	760-024502	DV9616	Rear Fan Tray
Fan Tray 4	REV 06	760-024502	DV7807	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7828	Rear Fan Tray

show chassis hardware extensive (TX Matrix Plus Router)

```
user@host> show chassis hardware extensive
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP
Jedec Code:	0x7fb0	EEPROM Version:	0x02	
		S/N:	JN112F007AHB	
Assembly ID:	0x052c	Assembly Version:	00.00	
Date:	00-00-0000	Assembly Flags:	0x00	
ID:	TXP			
Board Information Record:				

```

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 05    710-022574    TS4027          SFC Midplane
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          710-022574      S/N:              S/N TS4027
Assembly ID:  0x0962          Assembly Version:  01.05
Date:         03-23-2009      Assembly Flags:    0x00
Version:      REV 05
ID: SFC Midplane
Board Information Record:
Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00
Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07
Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Display       REV 03    710-024027    DX0282          TXP FPM Display
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          710-024027      S/N:              S/N DX0282
Assembly ID:  0x096c          Assembly Version:  01.03
Date:         02-10-2009      Assembly Flags:    0x00
Version:      REV 03
ID: TXP FPM Display          FRU Model Number:  CRAFT-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00
Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 07
Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0             REV 04    710-023792    DW4889          TXP CIP
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          710-023792      S/N:              S/N DW4889
Assembly ID:  0x0969          Assembly Version:  01.04
Date:         01-26-2009      Assembly Flags:    0x00
Version:      REV 04
ID: TXP CIP          FRU Model Number:  CIP-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware clei-models (TX Matrix Plus Router)

```

user@host> show chassis hardware clei-models
sfc0-re0:

```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 03	710-024027		CRAFT-TXP-S
CIP 0	REV 05	710-023792		CIP-TXP-S
CIP 1	REV 05	710-023792		CIP-TXP-S
PEM 0	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
PEM 1	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
Routing Engine 0	REV 06	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 05	710-022606		CB-TXP-S
CB 1	REV 09	710-022606		CB-TXP-S
SIB F13 0	REV 04	750-024564		SIB-TXP-F13
SIB F13 3	REV 04	750-024564		SIB-TXP-F13
SIB F13 8	REV 04	750-024564		SIB-TXP-F13
SIB F13 11	REV 04	750-024564		SIB-TXP-F13
SIB F13 12	REV 03	750-024564		SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 0/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 1/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/6	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/2	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 3/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/2	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 3/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/6	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/0	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 4/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 4/6	REV 03	710-022603		SIB-TXP-F2S-S
Fan Tray 0	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 3				
Fan Tray 4	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 5	REV 02	760-024502		FANTRAY-TXP-V-S

lcc0-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 1	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 1	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 05	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-004424		PC-1XGE-LR

PIC 2	REV 01	750-003336	PC-40C48-SON-SMSR
FPC 3	REV 12	710-013037	T1600-FPC4-ES
PIC 0	REV 02	750-010850	PD-10C768-SON-SR
FPC 4	REV 05	710-021534	T640-FPC1-ES
PIC 0	REV 04	750-014627	PB-40C3-10C12-SON-SFP
PIC 1	REV 22	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 09	750-002911	PB-4FE-TX
PIC 3	REV 08	750-021652	PB-1CHOC12-STM4-IQE-SFP
FPC 5	REV 07	710-007529	T640-FPC3
PIC 0	REV 14	750-009567	PC-1XGE-XENPAK
PIC 1	REV 16	750-007141	PC-10GE-SFP
PIC 2	REV 12	750-009567	PC-1XGE-XENPAK
FPC 6	REV 07	710-013035	T640-FPC3-ES
PIC 0	REV 09	750-009567	PC-1XGE-XENPAK
PIC 1	REV 06	750-015217	PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 06	750-015217	PC-8GE-TYPE3-SFP-IQ2
FPC 7	REV 03	710-021540	T640-FPC2-ES
PIC 0	REV 13	750-001901	PB-40C12-SON-SMIR
PIC 1	REV 05	750-001900	PB-10C48-SON-SMSR
PIC 2	REV 10	750-008155	PB-2GE-SFP-QPP
PIC 3	REV 03	750-014638	PB-10C48-SON-B-SFP
SIB 0	REV 07	710-022594	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	SIB-TXP-T1600-S
SIB 3	REV 06	710-022594	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	SIB-TXP-T1600-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 0	REV 02	710-010845		T640-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
FPC 1	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP
FPC 2	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 05	750-034781		PD-1CE-CFP
FPC 3	REV 10	710-021534		T640-FPC1-ES
PIC 0	REV 13	750-012266		PB-4GE-TYPE1-SFP-IQ2
PIC 1	REV 01	750-007641		PE-1GE-SFP-QPP
PIC 3	REV 17	750-007444		PB-1CHSTM1-SMIR-QPP
FPC 4	REV 06	710-013035		T640-FPC3-ES
PIC 0	REV 22	750-007141		PC-10GE-SFP
PIC 1	REV 16	750-009450		PC-10C192-SON-SR2
PIC 2	REV 05	750-004424		PC-1XGE-LR
PIC 3	REV 12	750-013423		PC-MS-500-3
FPC 5	REV 07	710-013560		T640-FPC3-E2
PIC 0	REV 11	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-004695		PC-TUNNEL

PIC 2	REV 32	750-003700	PC-10C192-SON-VSR
PIC 3	REV 12	750-009553	PC-40C48-SON-SFP
FPC 6	REV 07	710-013035	T640-FPC3-ES
PIC 0	REV 07	750-015217	PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-003336	PC-40C48-SON-SMSR
PIC 3	REV 02	750-012793	PC-1XGE-TYPE3-XFP-IQ2
FPC 7	REV 08	710-010845	T640-FPC4-ES
PIC 0	REV 11	750-017405	PD-4XGE-XFP
SIB 0	REV 07	710-022594	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	SIB-TXP-T1600-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FANTRAY-TXP-R-S

show chassis hardware detail (TX Matrix Plus Router)

```
user@host> show chassis hardware detail
sfc0-re0:
```

----- Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN111B023AHB	TXP
Midplane	REV 01	710-022574	TR7990	SFC Midplane
FPM Display	REV 03	710-024027	DW4699	TXP FPM Display
CIP 0	REV 01	710-023792	DR1437	TXP CIP
CIP 1	REV 02	710-023792	DS4564	TXP CIP
PEM 0	Rev 07	740-027463	UM26360	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	200811050193CEB1CEB1	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A0762	Disk 1
Routing Engine 1	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	20081105004C19A019A0	Compact Flash
ad1	30533 MB	SAMSUNG MCBQE32G8MPP-0V	SY814A0794	Disk 1
CB 0	REV 03	710-022606	DR7134	SFC Control Board
CB 1	REV 01	710-022606	DP8890	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 03	750-024564	DT9478	F13 SIB
B Board	REV 02	710-023431	DT6554	F13 SIB
SIB F13 1	REV 03	750-024564	DT9454	F13 SIB
B Board	REV 02	710-023431	DT6551	F13 SIB
SIB F2S 0/0	REV 02	710-022603	DT2838	F2S SIB
B Board	REV 02	710-023787	DT1725	NEO PMB
SIB F2S 0/2	REV 02	710-022603	DT2824	F2S SIB
B Board	REV 02	710-023787	DT1706	NEO PMB
SIB F2S 0/4	REV 02	710-022603	DT2822	F2S SIB
B Board	REV 02	710-023787	DT1696	NEO PMB
SIB F2S 0/6	REV 02	710-022603	DT2823	F2S SIB
B Board	REV 02	710-023787	DT1717	NEO PMB
SIB F2S 1/0	REV 03	710-022603	DV0059	F2S SIB
B Board	REV 03	710-023787	DT9942	NEO PMB
SIB F2S 1/2	REV 02	710-022603	DT2826	F2S SIB
B Board	REV 02	710-023787	DT1713	NEO PMB
SIB F2S 1/4	REV 03	710-022603	DV0092	F2S SIB
B Board	REV 03	710-023787	DV0000	NEO PMB
SIB F2S 1/6	REV 03	710-022603	DV0079	F2S SIB
B Board	REV 03	710-023787	DT9972	NEO PMB
SIB F2S 2/0	REV 03	710-022603	DV0100	F2S SIB
B Board	REV 03	710-023787	DT9925	NEO PMB

SIB F2S 2/2	REV 03	710-022603	DV0050	F2S SIB
B Board	REV 03	710-023787	DV0005	NEO PMB
SIB F2S 2/4	REV 03	710-022603	DV0097	F2S SIB
B Board	REV 03	710-023787	DT9936	NEO PMB
Fan Tray 0	REV 02	760-024497	DR8286	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9624	Front Fan Tray
Fan Tray 2	REV 02	760-024502	DR8259	Rear Fan Tray
Fan Tray 3	REV 02	760-024502	DR8270	Rear Fan Tray
Fan Tray 4	REV 02	760-024502	DR8284	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7813	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1101F27AHA	T1600
Midplane	REV 04	710-017247	RC5317	T Series Backplane
FPM GBUS	REV 10	710-002901	DS8197	T640 FPM Board
FPM Display	REV 01	710-021387	DS6433	T1600 FPM Display
CIP	REV 06	710-002895	DS1493	T Series CIP
PEM 0	Rev 08	740-017906	UD26601	Power Entry Module 3x80
SCG 0	REV 15	710-003423	DP5847	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR0924	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026942	737F-1024	LCC RE
ad0 3887 MB SMART CF			2008110502B63E513E51	Compact Flash
ad1 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1208				Disk 1
Routing Engine 1	REV 01	740-026942	737F-1024	LCC RE
ad0 3887 MB SMART CF			2008110500F9A8A8A8A8	Compact Flash
ad1 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1076				Disk 1
CB 0	REV 05	710-022597	DV4264	LCC Control Board
CB 1	REV 03	710-022597	DP8558	LCC Control Board
FPC 0	REV 14	710-013037	DS9967	FPC Type 4-ES
CPU	REV 08	710-016744	DS3989	ST-PMB2
PIC 0	REV 12	750-013198	DL7506	1x Tunnel
PIC 1	REV 12	750-013198	DL7505	1x Tunnel
MMB 0	REV 01	710-025563	DS8524	ST-MMB2
MMB 1	REV 01	710-025563	DS8373	ST-MMB2
FPC 1	REV 14	710-013037	DT0027	FPC Type 4-ES
CPU	REV 09	710-016744	DS7684	ST-PMB2
PIC 0	REV 12	750-013198	DL7512	1x Tunnel
PIC 1	REV 12	750-013198	DL7498	1x Tunnel
MMB 0	REV 01	710-025563	DS8494	ST-MMB2
MMB 1	REV 01	710-025563	DS8436	ST-MMB2
SPMB 0	REV 04	710-023321	DV3867	LCC Switch CPU
SPMB 1	REV 02	710-023321	DP0238	LCC Switch CPU
SIB 0	REV 06	710-022594	DT8268	LCC SIB
B Board	REV 06	710-023185	DT5791	LCC SIB Mezz
SIB 1	REV 06	710-022594	DT8261	LCC SIB
B Board	REV 06	710-023185	DT5769	LCC SIB Mezz
SIB 2	REV 04	710-022594	DS2315	LCC SIB
B Board	REV 06	710-023185	DT5788	LCC SIB Mezz
SIB 3	REV 06	710-022594	DT8253	LCC SIB
B Board	REV 06	710-023185	DT5811	LCC SIB Mezz
SIB 4	REV 06	710-022594	DT8248	LCC SIB
B Board	REV 06	710-023185	DT5812	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
FPM Display	REV 03	710-024027	DX0282	CRAFT-TXP
CIP 0	REV 04	710-023792	DW4889	CIP-TXP
CIP 1	REV 04	710-023792	DW4887	CIP-TXP
PEM 0	Rev 07	740-027463	UM26368	yyyyyyyyyyyyyyyyyyyyyyyyyy
Routing Engine 0	REV 01	740-026942	737A-1064	RE-TXP-SFC-DUO-2600-16G
Routing Engine 1	REV 01	740-026942	737A-1082	RE-TXP-SFC-DUO-2600-16G
CB 0	REV 09	710-022606	DW6099	CB-TXP
CB 1	REV 09	710-022606	DW6096	CB-TXP
SIB F13 1	REV 04	750-024564	DW5776	SIB-TXP-F13
SIB F13 3	REV 04	750-024564	DW5762	SIB-TXP-F13
SIB F13 4	REV 04	750-024564	DW5797	SIB-TXP-F13
SIB F13 6	REV 04	750-024564	DW5770	SIB-TXP-F13
SIB F13 7	REV 04	750-024564	DW5758	SIB-TXP-F13
SIB F13 8	REV 04	750-024564	DW5761	SIB-TXP-F13
SIB F13 9	REV 04	750-024564	DW5754	SIB-TXP-F13
SIB F13 12	REV 04	750-024564	DW5794	SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603	DW7897	
SIB F2S 0/2	REV 05	710-022603	DW7833	
SIB F2S 0/4	REV 05	710-022603	DW7875	
SIB F2S 0/6	REV 05	710-022603	DW7860	
SIB F2S 1/0	REV 04	710-022603	DW4820	
SIB F2S 1/2	REV 05	710-022603	DW7849	
SIB F2S 1/4	REV 05	710-022603	DW7927	SIB-TXP-F2S
SIB F2S 1/6	REV 05	710-022603	DW7866	
SIB F2S 2/0	REV 05	710-022603	DW7880	
SIB F2S 2/2	REV 05	710-022603	DW7895	
SIB F2S 2/4	REV 05	710-022603	DW7907	
SIB F2S 2/6	REV 05	710-022603	DW7785	
SIB F2S 3/0	REV 05	710-022603	DW7782	
SIB F2S 3/2	REV 05	710-022603	DW7793	
SIB F2S 3/4	REV 05	710-022603	DW7779	
SIB F2S 3/6	REV 05	710-022603	DW7930	
SIB F2S 4/0	REV 05	710-022603	DW7867	
SIB F2S 4/2	REV 05	710-022603	DW7917	
SIB F2S 4/4	REV 05	710-022603	DW7929	
SIB F2S 4/6	REV 05	710-022603	DW7870	
Fan Tray 0	REV 06	760-024497	DV7831	FANTRAY-TXP-F
Fan Tray 1	REV 06	760-024497	DV9614	FANTRAY-TXP-F
Fan Tray 2	REV 06	760-024502	DV9618	FANTRAY-TXP-R
Fan Tray 3	REV 06	760-024502	DV9616	FANTRAY-TXP-R
Fan Tray 4	REV 06	760-024502	DV7807	FANTRAY-TXP-R
Fan Tray 5	REV 06	760-024502	DV7828	FANTRAY-TXP-R

```
lcc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3765	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN5441	CRAFT-T1600-S
CIP	REV 06	710-002895	DP6021	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UA26384	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UA26296	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DR0875	SCG-T-S
CB 0	REV 06	710-022597	DW8534	CB-LCC

CB 1	REV 06	710-022597	DW8527	CB-LCC
FPC 4	REV 12	710-013037	DJ8717	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8795	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8794	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS5335	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7634	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7637	PD-4XGE-XFP
FPC 7	REV 07	710-013035	DM0990	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8067	PC-10GE-SFP
PIC 1	REV 08	750-015749	WE9598	PC-10C192-SON-XFP
PIC 2	REV 10	750-009450	HX6466	PC-10C192-SON-SR2
SIB 0	REV 08	710-022594	DW8033	SIB-TXP-T1600-S
SIB 1	REV 08	710-022594	DW8044	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8020	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8063	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8064	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5361	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6430	CRAFT-T1600-S
CIP	REV 06	710-002895	DS4239	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26649	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5820	SCG-T-S
CB 0	REV 06	710-022597	DW8523	CB-LCC
CB 1	REV 06	710-022597	DW8528	CB-LCC
FPC 4	REV 12	710-013037	DP8509	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8808	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP7263	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS9961	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS5532	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7639	PD-4XGE-XFP
FPC 7	REV 03	710-013035	DF5564	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8063	PC-10GE-SFP
SIB 0	REV 08	710-022594	DW8035	SIB-TXP-T1600-S
SIB 1	REV 10	710-022594	DX7672	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8060	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8072	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8043	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3956	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN7030	CRAFT-T1600-S
CIP	REV 06	710-002895	DM3962	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26519	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26601	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP0277	SCG-T-S
CB 0	REV 06	710-022597	DW8524	CB-LCC
CB 1	REV 06	710-022597	DW8536	CB-LCC
FPC 4	REV 12	710-013037	DR1194	T1600-FPC4-ES

PIC 0	REV 11	750-017405	DP8811	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8823	PD-4XGE-XFP
FPC 5	REV 12	710-013037	DR1184	T1600-FPC4-ES
PIC 1	REV 11	750-017405	DP4744	PD-4XGE-XFP
FPC 6	REV 12	710-013037	DN8622	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9924	PD-40C192-SON-XFP
PIC 1	REV 11	750-017405	DP8776	PD-4XGE-XFP
FPC 7	REV 04	710-013560	JR3968	T640-FPC3-E2
PIC 0	REV 16	750-007141	NC9330	PC-10GE-SFP
SIB 0	REV 07	710-022594	DW4217	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4213	SIB-TXP-T1600-S
SIB 2	REV 07	710-022594	DW4189	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4173	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4201	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

```
lcc3-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5319	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6402	CRAFT-T1600-S
CIP	REV 06	710-002895	DR9973	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UC26496	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26599	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5831	SCG-T-S
CB 0	REV 06	710-022597	DW8533	CB-LCC
CB 1	REV 06	710-022597	DW8538	CB-LCC
FPC 0	REV 14	710-013037	DS5345	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7641	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS5479	PD-4XGE-XFP
FPC 1	REV 14	710-013037	DS7338	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7631	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7632	PD-4XGE-XFP
FPC 2	REV 14	710-013037	DS9962	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7581	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7627	PD-4XGE-XFP
FPC 4	REV 10	710-010845	JZ6573	T640-FPC4-ES
PIC 0	REV 14	750-012518	JT5124	PD-40C192-SON-XFP
FPC 5	REV 14	710-013037	DT0016	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9918	PD-40C192-SON-XFP
FPC 7	REV 07	710-013035	DM0967	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8059	PC-10GE-SFP
PIC 1	REV 13	750-004695	DM5712	PC-TUNNEL
SIB 0	REV 07	710-022594	DW4174	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4207	SIB-TXP-T1600-S
SIB 2	REV 06	710-022594	DT8231	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4175	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4209	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

show chassis hardware (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D

B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz

Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800
Routing Engine 1	REV 07	740-026941	P737F-002938	RE-DUO-1800
CB 0	REV 11	710-022597	EH4805	LCC Control Board
CB 1	REV 11	710-022597	EH4786	LCC Control Board
FPC 1	REV 01	710-033873	BBAH0320	FPC Type 3-ES
CPU	REV 11	710-016744	BBAF3281	ST-PMB2
MMB 0	REV 06	710-025563	BBAF5061	ST-MMB2
FPC 5	REV 04	710-033871	BBAM5070	FPC Type 4-ES
CPU	REV 11	710-016744	BBAM6653	ST-PMB2
PIC 1	REV 20	750-017405	BBAM1296	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10B42981	XFP-10G-SR
MMB 0	REV 07	710-025563	BBAN2631	ST-MMB2
MMB 1	REV 07	710-025563	BBAN2538	ST-MMB2
SPMB 0	REV 05	710-023321	EH3903	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3902	LCC Switch CPU
SIB 0	REV 01	750-041657	EH8019	LCC SIB 3D
B Board	REV 01	711-042424	EH7680	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB04F	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB04S	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04B	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB043	CXP Module
SIB 1	REV 01	750-041657	EH8012	LCC SIB 3D
B Board	REV 01	711-042424	EH7658	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05E	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01Z	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB018	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB054	CXP Module
SIB 2	REV 01	750-041657	EH7993	LCC SIB 3D
B Board	REV 01	711-042424	EH7678	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05C	CXP Module
Xcvr 2	REV 01	740-047547	XB47FB00N	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB05U	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05L	CXP Module
Fan Tray 0				Front Top Fan Tray

Fan Tray 1
Fan Tray 2

Front Bottom Fan Tray
Rear Fan Tray -- Rev 4

show chassis hardware clei-models (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware clei-models
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 09	710-024027		CRAFT-TXP-S
CIP 0	REV 12	710-023792		CIP-TXP-S
CIP 1	REV 12	710-023792		CIP-TXP-S
PEM 0	Rev 06	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC-S
Routing Engine 0	REV 07	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 07	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 13	710-022606		CB-TXP-S
CB 1	REV 14	710-022606		CB-TXP-S
SIB F13 0	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 1	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-048813		
Xcvr 10	REV 01	740-048813		
Xcvr 12	REV 01	740-048813		
Xcvr 14	REV 01	740-048813		
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D

Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 6	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 7	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 9	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 11	REV 10	750-035002	PROTOXCLEI	750-035002
Xcvr 0	REV 01	740-048813		

Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 12	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F2S 0/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/6	REV 08	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
Fan Tray 0	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 3	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 4	REV 10	760-024502		FANTRAY-TXP-V-S
Fan Tray 5	REV 10	760-024502		FANTRAY-TXP-V-S

lcc0-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-027486	IPMJ700DRD	CHAS-BP-T1600-S
FPM Display	REV 04	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
PEM 1	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
SCG 0	REV 18	710-003423		SCG-T-S
SCG 1	REV 18	710-003423		SCG-T-S

Routing Engine 0	REV 10	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 07	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 11	710-022597		CB-LCC-S
CB 1	REV 11	710-022597		CB-LCC-S
FPC 0	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 3	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 13	750-033423	XXXXXXXXDD	PF-12-24XGE-SFPP
FPC 4	REV 02	750-045173	IP9IAL4DAC	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 5	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 6	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 10	750-035293	IP9IAL3DAA	PF-1CGE-CFP
SIB 0	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 1	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 2	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 3	REV 07	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 4	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		

```

Xcvr 6          REV 01  740-048813
Xcvr 7          REV 01  740-048813
Fan Tray 0      FANTRAY-T-S
Fan Tray 1      FANTRAY-T-S
Fan Tray 2      FANTRAY-TXP3D-LCC-R-S
[Output Truncated]

```

show chassis hardware detail (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis hardware detail
sfc0-re0:

```

----- Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
ad0 3823 MB SMART CF			2011030400062C132C13	Compact Flash
ad1 62720 MB SMART Lite SATA Drive			201105100009A452A452	Disk 1
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
ad0 3823 MB SMART CF			20110508085EE471E471	Compact Flash
ad1 62720 MB SMART Lite SATA Drive			201110210089DF39DF39	Disk 1
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module

Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

1cc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
ad0	3823 MB	SMART CF	201103030490604E604E	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	20110729028B11D411D4	Disk 1
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
ad0	3823 MB	SMART CF	2011010504EB99649964	Compact Flash

```

ad1  62720 MB SMART Lite SATA Drive 201102140058934A934A Disk 1
CB 0          REV 11 710-022597 EH3611 LCC Control Board
CB 1          REV 11 710-022597 EH4798 LCC Control Board
FPC 5         REV 17 710-013037 BBAC5333 FPC Type 4-ES
CPU          REV 10 710-016744 BBAB7619 ST-PMB2
PIC 0         REV 18 750-017405 BBAE3420 4x 10GE (LAN/WAN) XFP
  Xcvr 0      REV 03 740-014289 T10C90659 XFP-10G-SR
  MMB 0       REV 05 710-025563 BBAB9538 ST-MMB2
  MMB 1       REV 05 710-025563 BBAB9502 ST-MMB2
FPC 7         REV 01 750-045173 BBAV0032 FPC Type 5-3D
CPU
SPMB 0       REV 05 710-023321 EG9434 LCC Switch CPU
SPMB 1       REV 05 710-023321 EH3878 LCC Switch CPU
SIB 0        REV 01 750-041657 EH7997 LCC SIB 3D
  B Board    REV 01 711-042424 EH7674 LCC SIB 3D Mezz
  Xcvr 0     REV 01 740-047547 XB48FB014 CXP Module
  Xcvr 2     REV 01 740-047547 XB48FB05A CXP Module
  Xcvr 4     REV 01 740-047547 XB48FB052 CXP Module
  Xcvr 6     REV 01 740-047547 XB48FB01B CXP Module
SIB 1        REV 01 750-041657 EH8023 LCC SIB 3D
  B Board    REV 01 711-042424 EH7659 LCC SIB 3D Mezz
  Xcvr 0     REV 01 740-047547 XB48FB05J CXP Module
  Xcvr 2     REV 01 740-047547 XB48FB01E CXP Module
  Xcvr 4     REV 01 740-047547 XB48FB01J CXP Module
  Xcvr 6     REV 01 740-047547 XB48FB02S CXP Module
SIB 2        REV 03 750-041657 EJ6554 LCC SIB 3D
  B Board    REV 02 711-042424 EJ5756 LCC SIB 3D Mezz
  Xcvr 0     REV 01 740-047547 XB34FB01Z CXP Module
  Xcvr 2     REV 01 740-047547 XB34FB013 CXP Module
  Xcvr 4     REV 01 740-047547 XB48FB04Z CXP Module
  Xcvr 6     REV 01 740-047547 XB48FB05N CXP Module
Fan Tray 0   Front Top Fan Tray
Fan Tray 1   Front Bottom Fan Tray
Fan Tray 2   Rear Fan Tray -- Rev 4

```

```
lcc2-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800
ad0 3823 MB SMART CF			201103030356329E329E	Compact Flash
ad1 62720 MB SMART Lite SATA Drive			2011051000488D8B8D8B	Disk 1
Routing Engine 1	REV 07	740-026941	P737F-002938	RE-DUO-1800
ad0 3823 MB SMART CF			20110304000F02680268	Compact Flash
ad1 62720 MB SMART Lite SATA Drive			201105300A70F325F325	Disk 1
CB 0	REV 11	710-022597	EH4805	LCC Control Board
CB 1	REV 11	710-022597	EH4786	LCC Control Board
FPC 1	REV 01	710-033873	BBAH0320	FPC Type 3-ES
CPU	REV 11	710-016744	BBAF3281	ST-PMB2
MMB 0	REV 06	710-025563	BBAF5061	ST-MMB2
FPC 5	REV 04	710-033871	BBAM5070	FPC Type 4-ES
CPU	REV 11	710-016744	BBAM6653	ST-PMB2

PIC 1	REV 20	750-017405	BBAM1296	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10B42981	XFP-10G-SR
MMB 0	REV 07	710-025563	BBAN2631	ST-MMB2
MMB 1	REV 07	710-025563	BBAN2538	ST-MMB2
SPMB 0	REV 05	710-023321	EH3903	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3902	LCC Switch CPU
SIB 0	REV 01	750-041657	EH8019	LCC SIB 3D
B Board	REV 01	711-042424	EH7680	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB04F	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB04S	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04B	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB043	CXP Module
SIB 1	REV 01	750-041657	EH8012	LCC SIB 3D
B Board	REV 01	711-042424	EH7658	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05E	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01Z	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB018	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB054	CXP Module
SIB 2	REV 01	750-041657	EH7993	LCC SIB 3D
B Board	REV 01	711-042424	EH7678	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05C	CXP Module
Xcvr 2	REV 01	740-047547	XB47FB00N	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB05U	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05L	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

show chassis hardware lcc (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

```
-----
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module

Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

show chassis hardware sfc (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module

SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN112D865AFA	MX960
Midplane	REV 03	710-013698	TS3339	MX960 Backplane
FPM Board	REV 03	710-014974	WW6267	Front Panel Display
PDM	Rev 03	740-013110	QCS12485026	Power Distribution

Module					
in	PEM 0	Rev 04	740-013682	QCS12434086	PS 1.7kW; 200-240VAC
in	PEM 1	Rev 04	740-013682	QCS1243408Z	PS 1.7kW; 200-240VAC
in	PEM 2	Rev 04	740-013682	QCS1243407X	PS 1.7kW; 200-240VAC
	Routing Engine 0	REV 07	740-015113	9009009677	RE-S-1300
	Routing Engine 1	REV 07	740-015113	9009011510	RE-S-1300
	CB 0	REV 03	710-021523	XF0394	MX SCB
	CB 1	REV 03	710-021523	XF0550	MX SCB
	CB 2	REV 03	710-021523	XD7455	MX SCB
	FPC 4	REV 02	750-028467	JR6127	MPC M 16x 10GE
	CPU	REV 02	711-029089	JX0129	AS PMB
	PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
	PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
	PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
	PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
	Fan Tray 0	REV 05	740-014971	TP9990	Fan Tray
	Fan Tray 1	REV 05	740-014971	VS1709	Fan Tray

show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1101AFEAFB	MX480
Midplane	REV 05	710-017414	TR4444	MX480 Midplane
FPM Board	REV 02	710-017254	KG6056	Front Panel Display
PEM 0	Rev 03	740-017330	QCS082090FC	PS 1.2-1.7kW; 100-240V
PEM 1	Rev 03	740-017330	QCS082090FD	PS 1.2-1.7kW; 100-240V
Routing Engine 0	REV 07	740-013063	9009004124	RE-S-2000
Routing Engine 1	REV 07	740-013063	9009005569	RE-S-2000
CB 0	REV 07	710-021523	XZ3587	MX SCB
CB 1	REV 03	710-021523	KH8306	MX SCB
FPC 1	REV 04.1.07	750-033205	P1240	MPC Type 3
CPU	REV 01	711-035209	YL0504	HMPC PMB 2G
MIC 1	REV 10	750-033199	YX4495	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	C22CQNE	CFP-100G-LR4
FPC 2	REV 26	750-016670	KH0045	DPCE 40x 1GE R EQ
CPU	REV 07	710-013713	KF5448	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PF21JHU	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 9	REV 01	740-011613	AM0813S8ZL6	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 02	740-011613	PGL2KYF	SFP-SX
Xcvr 2	REV 01	740-011613	AM0806S8N4P	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 5	REV 01	740-011613	AM0815S967N	SFP-SX
Xcvr 7	REV 01	740-011613	AM0806S8N1X	SFP-SX
Xcvr 8	REV 01	740-011613	AM0815S967J	SFP-SX
Xcvr 9	REV 01	740-011613	AM0815S967M	SFP-SX
FPC 3	REV 12.2.09	750-033205	YR9443	MPC Type 3
CPU	REV 03	711-035209	YL6931	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3269	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP

Xcvr 0	REV 01	740-032210	ULH0KG3	CFP-100G-LR4
MIC 1	REV 02	750-033199	YG3245	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULH0KGF	CFP-100G-LR4
FPC 4	REV 12.3.09	750-033205	YR9437	MPC Type 3
CPU	REV 03	711-035209	YT5857	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3295	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12000187	CFP-100G-SR10
MIC 1	REV 10	750-033199	YX4518	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00008	CFP-100G-SR10
FPC 5	REV 06	750-024884	JW9769	MPC Type 2 3D EQ
CPU	REV 02	711-028401	JR6158	MPC PMB 2G Proto
MIC 0	REV 05	750-028387	JR6197	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71112	XFP-10G-SR
Xcvr 1	REV 02	740-014289	T08L85610	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
MIC 1	REV 22	750-028392	YM0053	3D 20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0703S005B	SFP-SX
Xcvr 1	REV 01	740-011613	E07L01352	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 5	REV 01	740-013111	6500217	SFP-T
Xcvr 9	REV 02	740-013111	8499527	SFP-T
Fan Tray				Left Fan Tray

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

show chassis hardware (QFX3500 Switches)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               QFX3500
Routing Engine 0                               QFX Routing Engine
FPC 0         REV 04   750-044071   BBAR3902      QFX3500-48S4Q-AFI
CPU                               FPC CPU
PIC 0                               48x 10G-SFP+
PIC 1                               15x 10G-SFP+
MGMT BRD      REV 02   750-044063   BBAR0398      QFX3500-MGMT-SFP-AF0
Xcvr 0        REV 01   740-011614   AC0946S0BD1   SFP-LX10
Xcvr 1        REV 02   740-013111   A281922       SFP-T
Power Supply 0 Rev 04   740-032091   UI00677       JPSU-650W-AC-AFI
Power Supply 1 REV 00   740-041741   VJ00162       JPSU-650W-AC-AF0
Fan Tray 0                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 1                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 2                               QFX Fan Tray, Back to
Front Airlfow

```

show chassis hardware detail (QFX3500 Switches)

```

user@switch> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN000TEST5    QFX3500

```

Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
FPC 0	REV 05	750-036931	EE0823	QFX3500-48S4Q-AFI
CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	48x 10G-SFP+
Xcvr 0	REV 01	740-030589	S99E270079	SFP+-10G-LPBK
Xcvr 1	REV 01	740-030589	S9AK450099	SFP+-10G-LPBK
Xcvr 2	REV 01	740-030589	S99E270078	SFP+-10G-LPBK
Xcvr 3	REV 01	740-030589	S9AK450098	SFP+-10G-LPBK
Xcvr 4	REV 01	740-030589	S99E270075	SFP+-10G-LPBK
Xcvr 5	REV 01	740-030589	S9AK450093	SFP+-10G-LPBK
Xcvr 6	REV 01	740-030589	S9AK450097	SFP+-10G-LPBK
Xcvr 7	REV 01	740-030589	S9AK450095	SFP+-10G-LPBK
Xcvr 8	REV 01	740-030589	S99E270072	SFP+-10G-LPBK
Xcvr 9	REV 01	740-030589	S99E270073	SFP+-10G-LPBK
Xcvr 10	REV 01	740-030589	S99E270080	SFP+-10G-LPBK
Xcvr 11	REV 01	740-030589	S9AK450169	SFP+-10G-LPBK
Xcvr 12	REV 01	740-030589	S99E270076	SFP+-10G-LPBK
Xcvr 13	REV 01	740-030589	S9AK450167	SFP+-10G-LPBK
Xcvr 14	REV 01	740-030589	S9AK450170	SFP+-10G-LPBK
Xcvr 15	REV 01	740-030589	S9AK450166	SFP+-10G-LPBK
Xcvr 16	REV 01	740-030589	S9AK450092	SFP+-10G-LPBK
Xcvr 17	REV 01	740-030589	S9AK450163	SFP+-10G-LPBK
Xcvr 18	REV 01	740-030589	S9AK450094	SFP+-10G-LPBK
Xcvr 19	REV 01	740-030589	S9AK450100	SFP+-10G-LPBK
Xcvr 20	REV 01	740-030589	S9AK450168	SFP+-10G-LPBK
Xcvr 21	REV 01	740-030589	S9AK450165	SFP+-10G-LPBK
Xcvr 22	REV 01	740-030589	S9AK450073	SFP+-10G-LPBK
Xcvr 23	REV 01	740-030589	S9AK450164	SFP+-10G-LPBK
Xcvr 24	REV 01	740-030589	S9AK450074	SFP+-10G-LPBK
Xcvr 25	REV 01	740-030589	SA62270195	SFP+-10G-LPBK
Xcvr 26	REV 01	740-030589	S9AK450078	SFP+-10G-LPBK
Xcvr 27	REV 01	740-030589	S9AK450024	SFP+-10G-LPBK
Xcvr 28	REV 01	740-030589	S9AK450027	SFP+-10G-LPBK
Xcvr 29	REV 01	740-030589	S9AK450080	SFP+-10G-LPBK
Xcvr 30	REV 01	740-030589	S9AK450030	SFP+-10G-LPBK
Xcvr 31	REV 01	740-030589	S9AK450025	SFP+-10G-LPBK
Xcvr 32	REV 01	740-030589	S9AK450023	SFP+-10G-LPBK
Xcvr 33	REV 01	740-030589	S9AK450075	SFP+-10G-LPBK
Xcvr 34	REV 01	740-030589	S9AK450161	SFP+-10G-LPBK
Xcvr 35	REV 01	740-030589	S9AK450071	SFP+-10G-LPBK
Xcvr 36	REV 01	740-030589	S9AK450072	SFP+-10G-LPBK
Xcvr 37	REV 01	740-030589	S9AK450022	SFP+-10G-LPBK
Xcvr 38	REV 01	740-030589	S9AK450021	SFP+-10G-LPBK
Xcvr 39	REV 01	740-030589	S9AK450175	SFP+-10G-LPBK
Xcvr 40	REV 01	740-030589	S9AK450162	SFP+-10G-LPBK
Xcvr 41	REV 01	740-030589	S99E270074	SFP+-10G-LPBK
Xcvr 42	REV 01	740-030589	S9AK450174	SFP+-10G-LPBK
Xcvr 43	REV 01	740-030589	S9AK450077	SFP+-10G-LPBK
Xcvr 44	REV 01	740-030589	S9AK450076	SFP+-10G-LPBK
Xcvr 45	REV 01	740-030589	S9AK450026	SFP+-10G-LPBK
Xcvr 46	REV 01	740-030589	S9AK450079	SFP+-10G-LPBK
Xcvr 47	REV 01	740-030589	S9AK450029	SFP+-10G-LPBK
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
Xcvr 1	REV 01	740-032986	QA170087	QSFP+-40G-SR4
Xcvr 4	REV 01	740-032986	QA360442	QSFP+-40G-SR4
Xcvr 8	REV 01	740-032986	QA170091	QSFP+-40G-SR4
Xcvr 12	REV 01	740-032986	QA170042	QSFP+-40G-SR4
MGMT BRD	REV 08	750-036946	EE0731	QFX3500-MB
Power Supply 0	Rev 04	740-032091	UI00690	QFX PS 650W AC

```

Power Supply 1  Rev 04  740-032091  UI00679      QFX PS 650W AC
Fan Tray 0                               QFX Fan Tray
Fan Tray 1                               QFX Fan Tray

```

show chassis hardware models (QFX3500 Switches)

```

user@switch> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Routing Engine 0      BUILTIN  BUILTIN
FPC 0          REV 02   711-032234  EC4074
Power Supply 0  PSMI 2C  11-d65800  --

```

show chassis hardware clei-models (QFX3500 Switches)

```

user@switch> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Routing Engine 0      BUILTIN
FPC 0          REV 02   711-032234
Power Supply 0  PSMI 2C  11-d65800

```

show chassis hardware interconnect-device (QFabric Systems)

```

user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 07
Midplane      REV 07   750-021261  BH0208188289  QFX Midplane
CB 0          REV 07   750-021261  BH0208188289  QFXIC08-CB4S

```

show chassis hardware node-device (QFabric Systems)

```

user@switch> show chassis hardware node-device node1
Routing Engine 0  BUILTIN  BUILTIN  QFX Routing Engine
node1            REV 05   711-032234  ED3694      QFX3500-48S4Q-AFI

CPU
PIC 0            BUILTIN  BUILTIN
Xcvr 8           REV 01   740-030658  AD0946A028B  48x 10G-SFP+
SFP+-10G-USR
...

```

show chassis hardware (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 03   711-031896  JN11D1FD7AJA  PTX5000
Midplane      REV 08   760-030647  EG1679        Midplane-8S
FPM           REV 05   740-032019  ZE00006       Front Panel Display
PDU 0         Rev 05   740-032022  ZJ00018       DC Power Dist Unit
PSM 0         Rev 05   740-032022  ZJ00018       DC 12V Power Supply
PSM 1         Rev 04   740-032022  ZC00052       DC 12V Power Supply
PSM 2         Rev 04   740-032022  ZD00051       DC 12V Power Supply
PSM 3         Rev 05   740-032022  ZJ00060       DC 12V Power Supply
CCG 0         REV 04   750-030653  EG3703        Clock Generator
CCG 1         REV 04   750-030653  EG3698        Clock Generator
Routing Engine 0 REV 05   740-026942  P737A-002231  RE-DUO-2600
Routing Engine 1 REV 06   740-026942  P737A-002438  RE-DUO-2600
CB 0          REV 08   750-030625  EG5519        Control Board
CB 1          REV 08   750-030625  EG5516        Control Board

```

FPC 0	REV 18	750-036844	EJ3080	FPC
CPU	REV 12	711-030686	EJ3260	SNG PMB
FPC 2	REV 13	750-036844	EG5065	FPC
CPU	REV 09	711-030686	EG4082	SNG PMB
PIC 0	REV 14	750-031913	EG5127	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	143363A00240	SFP+-10G-SR
Xcvr 1	REV 01	740-031981	UK90PZ1	SFP+-10G-LR
Xcvr 2	REV 01	740-031980	AD1141A04XH	SFP+-10G-SR
Xcvr 3	REV 01	740-031981	UK90Q46	SFP+-10G-LR
Xcvr 4	REV 01	740-031980	AD1141A04X4	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11H02560	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11C01589	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AD1141A04XF	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01094	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LKF	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	183363A01528	SFP+-10G-SR
Xcvr 14	REV 01	740-031980	193363A01079	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	AK80MC8	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	AJCOBHC	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08D26856	SFP+-10G-LR
Xcvr 21	REV 01	740-031980	AK80KCT	SFP+-10G-SR
Xcvr 22	REV 01	740-031981	UK90PZL	SFP+-10G-LR
Xcvr 23	REV 01	740-031980	AK80N1V	SFP+-10G-SR
FPC 3	REV 13	750-036844	EG5074	FPC
CPU	REV 09	711-030686	EG4064	SNG PMB
PIC 1	REV 10	750-031903	EG0325	SNG Load
FPC 5	REV 06	750-036844	EH3198	FPC
CPU				
PIC 0	REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S

SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware clei-models (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware clei-models
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
FPM	REV 08	760-030647	PROTOXCLEI	CRAFT-PTX5000-S
PDU 0	Rev 05	740-032019	IPUPAHLKAA	PWR-SAN-PDU-DC
PSM 0	Rev 05	740-032022	IPUPAHNKAA	PSM-PTX-DC-120-S
PSM 1	Rev 04	740-032022	032022XXXX	PWR-SAN-12-DC
PSM 2	Rev 04	740-032022	032022XXXX	PWR-SAN-12-DC
PSM 3	Rev 05	740-032022	IPUPAHNKAA	PSM-PTX-DC-120-S
CCG 0	REV 04	750-030653	PROTOXCLEI	CCG-PTX-S
CCG 1	REV 04	750-030653	PROTOXCLEI	CCG-PTX-S
Routing Engine 0	REV 05	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 08	750-030625	PROTOXCLEI	CB-PTX-S
CB 1	REV 08	750-030625	PROTOXCLEI	CB-PTX-S
FPC 0	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
FPC 2	REV 13	750-036844	PROTOXCLEI	FPC-PTX-P1-A
PIC 0	REV 14	750-031913	PROTOXCLEI	P1-PTX-24-10GE-SFPP
FPC 3	REV 13	750-036844	PROTOXCLEI	FPC-PTX-P1-A
FPC 5				
PIC 0	REV 14	750-031913	PROTOXCLEI	P1-PTX-24-10GE-SFPP
FPC 6	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
FPC 7	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
SIB 0	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 1	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 2	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 3	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 4	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 5	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 6	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 7	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 8	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
Fan Tray 1	REV 04	760-030642	PROTOXCLEI	FAN-PTX-H-S

show chassis hardware detail (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware detail
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN11D1FD7AJA	PTX5000
Midplane	REV 03	711-031896	ABAC5589	Midplane-8S
FPM	REV 08	760-030647	EG1679	Front Panel Display
PDU 0	Rev 05	740-032019	ZE00006	DC Power Dist Unit
PSM 0	Rev 05	740-032022	ZJ00018	DC 12V Power Supply
PSM 1	Rev 04	740-032022	ZC00052	DC 12V Power Supply
PSM 2	Rev 04	740-032022	ZD00051	DC 12V Power Supply
PSM 3	Rev 05	740-032022	ZJ00060	DC 12V Power Supply
CCG 0	REV 04	750-030653	EG3703	Clock Generator
CCG 1	REV 04	750-030653	EG3698	Clock Generator
Routing Engine 0	REV 05	740-026942	P737A-002231	RE-DUO-2600
ad0	3823 MB	SMART CF	201006190039C02DC02D	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	2011042300CF4C6B4C6B	Disk 1
Routing Engine 1	REV 06	740-026942	P737A-002438	RE-DUO-2600
ad0	3823 MB	SMART CF	20100619053455F055F0	Compact Flash

ad1	62720 MB	SMART	Lite SATA Drive	20110423000AE8E7E8E7	Disk 1
CB 0		REV 08	750-030625	EG5519	Control Board
CB 1		REV 08	750-030625	EG5516	Control Board
FPC 0		REV 18	750-036844	EJ3080	FPC
CPU		REV 12	711-030686	EJ3260	SNG PMB
FPC 2		REV 13	750-036844	EG5065	FPC
CPU		REV 09	711-030686	EG4082	SNG PMB
PIC 0		REV 14	750-031913	EG5127	24x 10GE(LAN) SFP+
Xcvr 0		REV 01	740-031980	143363A00240	SFP+-10G-SR
Xcvr 1		REV 01	740-031981	UK90PZ1	SFP+-10G-LR
Xcvr 2		REV 01	740-031980	AD1141A04XH	SFP+-10G-SR
Xcvr 3		REV 01	740-031981	UK90Q46	SFP+-10G-LR
Xcvr 4		REV 01	740-031980	AD1141A04X4	SFP+-10G-SR
Xcvr 6		REV 01	740-031980	B11H02560	SFP+-10G-SR
Xcvr 7		REV 01	740-031980	B11C01589	SFP+-10G-SR
Xcvr 8		REV 01	740-031980	AD1141A04XF	SFP+-10G-SR
Xcvr 10		REV 01	740-031980	123363A01094	SFP+-10G-SR
Xcvr 11		REV 01	740-031980	AK80LKF	SFP+-10G-SR
Xcvr 12		REV 01	740-031980	183363A01528	SFP+-10G-SR
Xcvr 14		REV 01	740-031980	193363A01079	SFP+-10G-SR
Xcvr 15		REV 01	740-031980	AK80MC8	SFP+-10G-SR
Xcvr 16		REV 01	740-031980	AJC0BHC	SFP+-10G-SR
Xcvr 19		REV 01	740-021309	J08D26856	SFP+-10G-LR
Xcvr 21		REV 01	740-031980	AK80KCT	SFP+-10G-SR
Xcvr 22		REV 01	740-031981	UK90PZL	SFP+-10G-LR
Xcvr 23		REV 01	740-031980	AK80N1V	SFP+-10G-SR
FPC 3		REV 13	750-036844	EG5074	FPC
CPU		REV 09	711-030686	EG4064	SNG PMB
PIC 1		REV 10	750-031903	EG0325	SNG Load
FPC 5		REV 06	750-036844	EH3198	FPC
CPU					
PIC 0		REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0		REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1		REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2		REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5		REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6		REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7		REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10		REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11		REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12		REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15		REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16		REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18		REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19		REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22		REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23		REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1		REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0		REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1		REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6		REV 18	750-036844	EJ4391	FPC
CPU		REV 12	711-030686	EJ3257	SNG PMB
FPC 7		REV 18	750-036844	EJ4382	FPC
CPU		REV 12	711-030686	EJ3238	SNG PMB
SPMB 0		REV 10	711-030686	EG5418	SNG PMB
SPMB 1		REV 09	711-030686	EG5373	SNG PMB
SIB 0		REV 07	750-030631	EG4858	SIB-I-8S
SIB 1		REV 07	750-030631	EG4872	SIB-I-8S
SIB 2		REV 07	750-030631	EG4866	SIB-I-8S
SIB 3		REV 07	750-030631	EG6011	SIB-I-8S
SIB 4		REV 07	750-030631	EG4907	SIB-I-8S

SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware models (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
FPM                 REV 08    760-030647   EG1679         CRAFT-PTX5000-S
PDU 0              Rev 05    740-032019   ZE00006        PWR-SAN-PDU-DC
  PSM 0            Rev 05    740-032022   ZJ00018        PSM-PTX-DC-120-S
  PSM 1            Rev 04    740-032022   ZC00052        PWR-SAN-12-DC
  PSM 2            Rev 04    740-032022   ZD00051        PWR-SAN-12-DC
  PSM 3            Rev 05    740-032022   ZJ00060        PSM-PTX-DC-120-S
CCG 0              REV 04    750-030653   EG3703         CCG-PTX-S
CCG 1              REV 04    750-030653   EG3698         CCG-PTX-S
Routing Engine 0   REV 05    740-026942   P737A-002231   RE-DUO-C2600-16G-S
Routing Engine 1   REV 06    740-026942   P737A-002438   RE-DUO-C2600-16G-S
CB 0               REV 08    750-030625   EG5519         CB-PTX-S
CB 1               REV 08    750-030625   EG5516         CB-PTX-S
FPC 0              REV 18    750-036844   EJ3080         FPC-PTX-P1-A
FPC 2              REV 13    750-036844   EG5065         FPC-PTX-P1-A
  PIC 0            REV 14    750-031913   EG5127         P1-PTX-24-10GE-SFPP
FPC 3              REV 13    750-036844   EG5074         FPC-PTX-P1-A
FPC 5
  PIC 0            REV 14    750-031913   EG5134         P1-PTX-24-10GE-SFPP
FPC 6              REV 18    750-036844   EJ4391         FPC-PTX-P1-A
FPC 7              REV 18    750-036844   EJ4382         FPC-PTX-P1-A
SIB 0              REV 07    750-030631   EG4858         SIB-I-PTX5008
SIB 1              REV 07    750-030631   EG4872         SIB-I-PTX5008
SIB 2              REV 07    750-030631   EG4866         SIB-I-PTX5008
SIB 3              REV 07    750-030631   EG6011         SIB-I-PTX5008
SIB 4              REV 07    750-030631   EG4907         SIB-I-PTX5008
SIB 5              REV 07    750-030631   EG4879         SIB-I-PTX5008
SIB 6              REV 07    750-030631   EG4864         SIB-I-PTX5008
SIB 7              REV 07    750-030631   EG4899         SIB-I-PTX5008
SIB 8              REV 07    750-030631   EG4880         SIB-I-PTX5008
Fan Tray 1         REV 04    760-030642   EG1335         FAN-PTX-H-S

```

show chassis hardware extensive (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware extensive
Hardware inventory:
Item                Version  Part number  Serial number  Description
.....
PDU 0              Rev 04    740-032019   UE0003         DC Power Dist Unit
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-032019        S/N: S/N UE0003
Assembly ID: 0x043d     Assembly Version: 04.00
Date: 11-29-2010       Assembly Flags: 0x00
Version: Rev 04        CLEI Code: 032022XXXX
ID: DC Power Dist Unit FRU Model Number: PWR-SAN-PDU-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00

```

```

Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0          Rev 04    740-032022  YG00065          DC 12V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-032022      S/N:              S/N YG00065
Assembly ID:   0x0440          Assembly Version:  04.00
Date:          07-30-2010      Assembly Flags:    0x00
Version:       Rev 04          CLEI Code:         032022XXXX
ID: DC 12V Power Supply Module FRU Model Number: PWR-SAN-12-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20
Address 0x60: 20 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1100FB1AFB  MX480
Midplane      REV 05   710-017414   TR3310         MX480 Midplane
FPM Board     REV 02   710-017254   KG1872         Front Panel Display
PEM 2         Rev 02   740-017343   QCS0812A00N    DC Power Entry Module
PEM 3         Rev 02   740-017343   QCS0812A00U    DC Power Entry Module
Routing Engine 0 REV 07   740-015113   1000740938     RE-S-1300
CB 0          REV 03   710-021523   KF4630         MX SCB
FPC 1         REV 11   750-037207   ZW9726         AS-MCC
CPU           REV 04   711-038173   ZW4819         AS-MCC PMB
MIC 0         REV 06   750-037214   ZW3574         AS-MSC
PIC 0         BUILTIN BUILTIN      AS-MSC
MIC 1         REV 00   750-037211   BUILTIN        AS-MXC
PIC 2         BUILTIN BUILTIN      AS-MXC

```

show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware extensive
FPC 1          REV 11   750-037207   ZW9726          AS-MCC
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037207      S/N:              S/N ZW9726
Assembly ID:   0x0b37          Assembly Version:  01.11
Date:          02-17-2012      Assembly Flags:    0x00
Version:       REV 11          CLEI Code:         PROTOXCLEI
ID: AS-MCC      FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00

```

```

Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 04    711-038173    ZW4819          AS-MCC-PMB
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         711-038173      S/N:          S/N ZW4819
Assembly ID: 0x0b38          Assembly Version: 01.04
Date:        12-30-2011      Assembly Flags: 0x00
Version:     REV 04
ID: AS-MCC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 06    750-037214    ZW3574          AS-MSC
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         750-037214      S/N:          S/N ZW3574
Assembly ID: 0x0a44          Assembly Version: 01.06
Date:        02-19-2012      Assembly Flags: 0x00
Version:     REV 06          CLEI Code:      PROTOXCLEI
ID: AS-MSC          FRU Model Number: 750-037214
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00
PIC 0          BUILTIN    BUILTIN          AS-MSC
MIC 1          REV 00    750-037211          AS-MXC
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         750-037211
Assembly ID: 0x0a43          Assembly Version: 01.00
Date:        255-255-65535    Assembly Flags: 0x00
Version:     REV 00
ID: AS-MXC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 ff ff ff
Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
PIC 2 BUILTIN BUILTIN AS-MXC

show chassis lcd

List of Syntax	show chassis lcd (EX Series) on page 785 show chassis lcd (QFX Series and QFabric Systems) on page 785
show chassis lcd (EX Series)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <menu <(all-members local member <i>member-id</i>)>></pre>
show chassis lcd (QFX Series and QFabric Systems)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <interconnect-device <i>device-id</i>> <node-device <i>device-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>menu option introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.1 for QFabric systems.</p>
Description	<p>Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.</p>
Options	<p>none—Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.</p> <p>fpc-slot <<i>fpc-slot-number</i>>—(Optional) Display the information as follows:</p> <ul style="list-style-type: none"> (EX3200, EX3300, EX4200, and EX4500 switches, or the QFX Series) Display the information that appears on the LCD panel for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by fpc-slot 0. fpc-slot refers to the switch itself and 0 is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the none option. <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no <i>fpc-slot-number</i> value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the none option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the information for the specified member. <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> (EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the <i>fpc-slot-number</i> value.

Also display the status of the currently selected port parameter of the Status LED for each network port.

interconnect-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

menu—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member *member-id*—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

node-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

Required Privilege Level

view

Related Documentation

- *LCD Panel in EX3200 Switches*
- *LCD Panel in EX4200 Switches*
- *LCD Panel in EX4500 Switches*
- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- [set chassis display message on page 1469](#)

List of Sample Output

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 787](#)
[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 789](#)
[show chassis lcd \(EX8200 Switch\) on page 789](#)
[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 791](#)
[show chassis lcd menu \(EX4200 Switch\) on page 791](#)
[show chassis lcd menu \(EX8200 Switch\) on page 791](#)
[show chassis lcd \(QFX3500 Switches\) on page 792](#)
[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 792](#)
[show chassis lcd interconnect-device \(QFabric Systems\) on page 795](#)

[show chassis lcd node-device \(QFabric Systems\) on page 797](#)

Output Fields [Table 64 on page 787](#) lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

Table 64: show chassis lcd Output Fields

Field Name	Field Description
membernumber (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
Front panel contents for slot	FPC slot number of the switch whose content is being displayed. The number is always 0 , except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
Front panel contents (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the member number field for each member device in the EX8200 Virtual Chassis.</p>
LCD screen	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays RE and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet (EX3200 and EX4200 switches only)
LEDs status	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: The XRE200 External Routing Engine always displays the NA parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • NA—Not applicable. • POE—Power over Ethernet
fpcx	On standalone EX Series and QFX Series switches, always 0 . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

Sample Output

show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
```

Front panel contents for slot: 0

```

-----
LCD screen:
  00:BK switch1
  LED:SPD ALARM 00
LEDs status:
  Alarms LED: Off
  System LED: Green
  Master LED: Off
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off

```

Front panel contents for slot: 1

```

-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off
ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off

```



```

ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off

```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```

user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off
ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off
ge-1/0/15     Off
ge-1/0/16     Off
ge-1/0/17     Off
ge-1/0/18     Off
ge-1/0/19     Off
ge-1/0/20     Off
ge-1/0/21     Off
ge-1/0/22     Off
ge-1/0/23     Off

```

show chassis lcd (EX8200 Switch)

```

user@switch> show chassis lcd
Front panel contents:
-----
LCD screen:
  RE st-8200-r
  LED:ADM ALARM 01

```

LEDs status:

Alarms LED: Yellow

System LED: Yellow

Master LED: Green

Interface	LED(ADM/SPD/DPX)
-----------	------------------

ge-0/0/0	Off
ge-0/0/1	Off
ge-0/0/2	Off
ge-0/0/3	Off
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	Off
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	Off
ge-0/0/41	Off
ge-0/0/42	Off
ge-0/0/43	Off
ge-0/0/44	Off
ge-0/0/45	Off
ge-0/0/46	Off
ge-0/0/47	Off
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off

```

xe-2/0/7      Off
xe-3/0/0      Off
xe-3/0/1      Off
xe-3/0/2      Off
xe-3/0/3      Off
xe-3/0/4      Off
xe-3/0/5      Off
xe-3/0/6      Off
xe-3/0/7      Off
xe-5/0/0      Off
xe-5/0/1      Off
xe-5/0/2      Off
xe-5/0/3      Off
xe-5/0/4      Off
xe-5/0/5      Off
xe-5/0/6      On
xe-5/0/7      On
xe-7/0/5      Off

```

show chassis lcd fpc-slot 2 (EX8200 Switch)

```
show chassis lcd fpc-slot 2
```

Interface	LED (ADM/SPD/DPX)
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
```

```

-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default

```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu** **all-members** command is the same as the output for the **show chassis lcd menu** command.

show chassis lcd menu (EX8200 Switch)

```

user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu

```

```
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off
```

show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis lcd
member0:
-----
Front panel contents:
-----
LCD screen:
  RE ex8200-member0
  LED:ADM ALARM 04
LEDs status:
  Alarms LED: Red
  System LED: Yellow
  Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
  BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
  09:RE xre200-member9
  LED: NA ALARM 01
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0       On
ge-0/0/1       On
ge-0/0/2       On
ge-0/0/3       On
```

ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On
ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off
ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off

ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On
ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On
ge-22/0/17	Off
ge-22/0/18	On
ge-22/0/19	Off
ge-22/0/20	On
ge-22/0/21	Off
ge-22/0/22	On
ge-22/0/23	Off

```

ge-22/0/24      On
ge-22/0/25      Off
ge-22/0/26      Off
ge-22/0/27      Off
ge-22/0/28      Off
ge-22/0/29      Off
ge-22/0/30      Off
ge-22/0/31      Off
ge-22/0/32      On
ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off

```

show chassis lcd interconnect-device (QFabric Systems)

```

show chassis lcd interconnect-device IC-F1012
      Front Panel Module Information
      -----
      LCD screen:
      IC-F1012      3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off
  PEM 5 LED : off

      LED info for: CB - 0
      -----

LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface      STATUS LED      LINK/ACTIVITY LED

```

```

-----
IC-F1012:pme0 :      Green      N/A
IC-F1012:pme1 :      Green      N/A
IC-F1012:pme2 :      off        N/A
IC-F1012:pme3 :      off        N/A

```

LED info for: CB - 1

```

-----
LEDs status:
  Status LED: Green
  Mastership LED: Amber

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

```

-----
LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

LED info for: FC 1 FPC - 1

```

-----
LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

LED info for: RC 0 FPC - 8

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 1 FPC - 9

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 2 FPC - 10

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 3 FPC - 11


```

-----
LEDs status:
  Status LED: Green

  LED info for: RC 4 FPC - 12
-----
LEDs status:
  Status LED: Green

  LED info for: RC 5 FPC - 13
-----
LEDs status:
  Status LED: Green

  LED info for: RC 6 FPC - 14
-----
LEDs status:
  Status LED: Green

  LED info for: RC 7 FPC - 15
-----
LEDs status:
  Status LED: Green

```

show chassis lcd node-device (QFabric Systems)

```

show chassis lcd node-device P3774-C
  Front panel contents for: P3774-C
  -----
  LCD screen:
  P3774-C
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

show chassis led

List of Syntax	show chassis led (EX Series) on page 798 show chassis led (QFX Series) on page 798
show chassis led (EX Series)	<code>show chassis led</code> <code><fpc-slot <fpc-slot-number>></code>
show chassis led (QFX Series)	<code>show chassis led</code> <code><fpc-slot <fpc-slot-number>></code> <code>interconnect-device <i>name</i></code> <code>node-device <i>name</i></code>
Release Information	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the status and colors of the chassis LEDs on the front panel of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.
Options	<p>none—Display the status of the chassis status LEDs (for EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members).</p> <p>fpc-slot <fpc-slot-number>—(Optional) (Not on EX2200 switches) Display the information as follows:</p> <ul style="list-style-type: none">(EX3200, standalone EX4200, standalone QFX3500, and EX4500 switches) Display the status of the chassis status LEDs for either an FPC slot with no fpc-slot-number value specified or for the FPC slot specified by fpc-slot 0. fpc-slot refers to the switch itself and 0 is the only valid value for fpc-slot-number. Output for these options is the same as for the none option.(EX4200 switches in a Virtual Chassis with two or more members) If no fpc-slot-number value is specified, display the status of the chassis status LEDs for all members of the Virtual Chassis. Output for this option is the same as for the none option. If the fpc-slot-number value is specified (it equals the member-id value), display the status of the chassis status LEDs for the specified member.(EX8200 switches)—Display the status of the chassis status LEDs for the line card in the line-card slot specified by the fpc-slot-number value. <p>interconnect-device <i>name</i>—</p> <p>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Interconnect device.</p> <p>node-device <i>name</i>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Node device.</p>

Required Privilege Level view

- Related Documentation**
- *Chassis Status LEDs in EX2200 Switches*
 - *Chassis Status LEDs in EX3200 Switches*
 - *Chassis Status LEDs in EX4200 Switches*
 - *Chassis Status LEDs in EX4500 Switches*
 - *Chassis Status LEDs in an EX8200 Switch*
 - *Chassis Status LEDs on a QFX3500 Device*
 - *Chassis Status LEDs in the QFX3600 and QFX3600-I Device*
 - *Management Port LEDs on a QFX3500 Device*
 - *Management Port LEDs in the QFX3600 and QFX3600-I Device*
 - *Chassis Status LEDs on a QFX3008-I Interconnect Device*
 - *Control Board LEDs on a QFX3008-I Interconnect Device*

List of Sample Output

[show chassis led \(EX2200 Switch\) on page 802](#)
[show chassis led on page 803](#)
[show chassis led fpc-slot 0 on page 804](#)
[show chassis led \(EX Series\) on page 804](#)
[show chassis led node-device \(QFabric System Node Device\) on page 805](#)
[show chassis led interconnect-device \(QFabric System - QFX3600-I Interconnect Device\) on page 805](#)
[show chassis led interconnect-device \(QFabric System - QFX3008-I Interconnect Device\) on page 806](#)

Output Fields [Table 51 on page 464](#) lists the output fields for the **show chassis led** command. Output fields are listed in the approximate order in which they appear.

Table 65: show chassis led Output Fields

Field Name	Field Description
Front panel contents for slot	FPC slot number of the device whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
Front panel contents (EX8200 Switches)	
Front Panel Module Information (QFabric system QFX3008-I Interconnect device)	On EX8200 switches, no slot number is displayed.
Front panel contents for (QFabric system Node devices and QFX3600-I Interconnect devices)	On QFabric system Node devices, the name of the Node device whose content is being displayed.

Table 65: show chassis led Output Fields (*continued*)

Field Name	Field Description
Alarms LED	<p>(EX Series switches only) Displays status of the ALM LED:</p> <ul style="list-style-type: none"> • Off—No alarm has been configured. • Green—No alarm has been triggered. • Red—Major alarm. • Yellow—Minor alarm
System LED	<p>(EX Series switches only) Displays status of the SYS LED:</p> <ul style="list-style-type: none"> • Off—Switch is powered off. • Green—Switch is operating normally. • Yellow—Switch is booting.
Master LED:	<p>Displays status of the MST LED (on EX3200, EX4200, and EX8200 switches):</p> <ul style="list-style-type: none"> • Green—On an EX4200 Virtual Chassis switch, indicates the switch is the master in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational. • Off <ul style="list-style-type: none"> • On an EX4200 Virtual Chassis switch, indicates that this switch is not the master in the Virtual Chassis configuration. • On EX3200, standalone EX4200, and EX8200 switches, indicates that the Routing Engine is not operational.
Mode LED:	<p>(EX Series switches only) On an EX2200 switch only, displays the currently selected port parameter of the Status LED:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet
Status/Beacon LED	<p>(QFX Series only) Displays the system status as indicated by the Status LED on the chassis. For more information, see:</p> <ul style="list-style-type: none"> • <i>Chassis Status LEDs on a QFX3500 Device</i> • <i>Chassis Status LEDs in the QFX3600 and QFX3600-I Device</i>
LINK/SPEED LED	<p>(QFX Series only) Displays the link status and speed of a management port. For more information, see:</p> <ul style="list-style-type: none"> • <i>Management Port LEDs on a QFX3500 Device</i> • <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i>
ACTIVITY LED	<p>(QFX Series only) Displays the activity status of a management port. For more information, see:</p> <ul style="list-style-type: none"> • <i>Management Port LEDs on a QFX3500 Device</i> • <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i>

Table 65: show chassis led Output Fields (*continued*)

Field Name	Field Description
STATUS LED	<p>(QFX Series only) Displays the link status of an interface as indicated by the ST LED. For more information, see:</p> <ul style="list-style-type: none"> Control Board LEDs on a QFX3008-I Interconnect Device Access Port and Uplink Port LEDs on a QFX3500 Device Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
LINK/ACTIVITY LED	<p>(QFX Series only) Displays link activity or faults on an interface as indicated by the LA LED. For more information, see:</p> <ul style="list-style-type: none"> Access Port and Uplink Port LEDs on a QFX3500 Device Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
Status LED	<p>(QFX3008-I Interconnect device only)</p> <ul style="list-style-type: none"> Displays the system status as indicated by the STATUS LED on the front panel of the chassis. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>. Displays the status of a Control Board as indicated by the STATUS LED on the Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i>.
Power LED	<p>(QFX3008-I Interconnect device only) Displays the status of system power on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Major Alarm LED	<p>(QFX3008-I Interconnect device only) Displays whether a critical error condition that requires immediate action exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Minor Alarm LED	<p>(QFX3008-I Interconnect device only) Displays whether a noncritical condition that requires monitoring or maintenance exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Fan 0 LED	<p>(QFX3008-I Interconnect device only) Displays the status of fan trays on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
Fan 1 LED	
Fan 2 LED	
Fan 3 LED	
Fan 4 LED	
Fan 5 LED	
Fan 6 LED	
Fan 7 LED	
Fan 8 LED	

Table 65: show chassis led Output Fields (*continued*)

Field Name	Field Description
PEM 0 LED	(QFX3008-I Interconnect device only) Displays the status of power supplies on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i> .
PEM 1 LED	
PEM 2 LED	
PEM 3 LED	
PEM 4 LED	
LED info for	(QFX3008-I Interconnect device only) Displays the LED information for a Control Board.
Mastership LED	(QFX3008-I Interconnect device only) Displays status of the MASTER LED on a Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i> .
Interface	Names of the interfaces on the device.
LED (ADM/SPD/DPX/POE)	<p>(EX Series switches only) State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: EX4500 and EX8200 switches do not have the POE port parameter.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet

Sample Output

show chassis led (EX2200 Switch)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  System LED: Green
  Mode LED : Duplex
Interface    LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex

```

```

ge-0/0/12      Full Duplex
ge-0/0/13      Full Duplex
ge-0/0/14      Full Duplex
ge-0/0/15      Full Duplex
ge-0/0/16      Full Duplex
ge-0/0/17      Full Duplex
ge-0/0/18      Full Duplex
ge-0/0/19      Full Duplex
ge-0/0/20      Full Duplex
ge-0/0/21      Full Duplex
ge-0/0/22      Off
ge-0/0/23      Off
ge-0/0/24      Full Duplex
ge-0/0/25      Full Duplex
ge-0/0/26      Off
ge-0/0/27      Off
ge-0/0/28      Full Duplex
ge-0/0/29      Full Duplex

```

show chassis led

```
user@switch> show chassis led
```

```
Front panel contents for slot: 0
```

```
-----
LEDs status:
```

```
  Alarms LED: Off
```

```
  System LED: Green
```

```
  Master LED: Green
```

```
Interface      LED (ADM/SPD/DPX/POE)
```

```
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex
ge-0/0/12     Full Duplex
ge-0/0/13     Full Duplex
ge-0/0/14     Full Duplex
ge-0/0/15     Full Duplex
ge-0/0/16     Full Duplex
ge-0/0/17     Full Duplex
ge-0/0/18     Full Duplex
ge-0/0/19     Full Duplex
ge-0/0/20     Full Duplex
ge-0/0/21     Full Duplex
ge-0/0/22     Off
ge-0/0/23     Off
ge-0/0/24     Full Duplex
ge-0/0/25     Full Duplex
ge-0/0/26     Off
ge-0/0/27     Off
ge-0/0/28     Full Duplex
ge-0/0/29     Full Duplex

```

show chassis led fpc-slot 0

```
user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Red
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off
```

show chassis led (EX Series)

```
user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  Status LED: Green
  Mode LED : Duplex
Interface LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0 Off
ge-0/0/1 Full Duplex
ge-0/0/2 Full Duplex
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Full Duplex
ge-0/0/6 Full Duplex
ge-0/0/7 Full Duplex
ge-0/0/8 Full Duplex
ge-0/0/9 Full Duplex
ge-0/0/10 Full Duplex
ge-0/0/11 Full Duplex
ge-0/0/12 Full Duplex
ge-0/0/13 Full Duplex
```



```

ge-0/0/14 Full Duplex
ge-0/0/15 Full Duplex
ge-0/0/16 Full Duplex
ge-0/0/17 Full Duplex
ge-0/0/18 Full Duplex
ge-0/0/19 Full Duplex
ge-0/0/20 Full Duplex
ge-0/0/21 Full Duplex
ge-0/0/22 Off
ge-0/0/23 Off
ge-0/0/24 Full Duplex
ge-0/0/25 Full Duplex
ge-0/0/26 Off
ge-0/0/27 Off
ge-0/0/28 Full Duplex
ge-0/0/29 Full Duplex

```

show chassis led node-device (QFabric System Node Device)

```

user@switch> show chassis led node-device node1
Front panel contents for: node1
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
node1:me5	Green	N/A
node1:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
node1:xe-0/0/8	Green	Green
node1:ge-0/0/10	Green	Green
node1:ge-0/0/12	Green	Green
node1:ge-0/0/24	Green	Green
node1:ge-0/0/25	Green	Green
node1:ge-0/0/26	Green	Green
node1:ge-0/0/27	Green	Green
node1:ge-0/0/28	Green	Green
node1:ge-0/0/29	Green	Green
node1:ge-0/0/30	Green	Green
node1:ge-0/0/31	Green	Green
node1:ge-0/0/32	Green	Green
node1:ge-0/0/33	Green	Green
node1:ge-0/0/34	Green	Green
node1:ge-0/0/35	Green	Green
node1:ge-0/0/36	Green	Green
node1:ge-0/0/37	Green	Green
node1:ge-0/0/38	Green	Green
node1:ge-0/0/39	Green	Green
node1:fte-0/1/0	Green	Green Blinking
node1:fte-0/1/2	Green	Green Blinking

show chassis led interconnect-device (QFabric System - QFX3600-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
Front panel contents for: FPC 0
-----
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
IC-EG0712:me5	Green	N/A
IC-EG0712:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
IC-EG0712:fte-0/1/0	Green	Green
IC-EG0712:fte-0/1/1	Green	Green Blinking
IC-EG0712:fte-0/1/2	Green	Green
IC-EG0712:fte-0/1/3	Green	Green Blinking
IC-EG0712:fte-0/1/4	Green	Green
IC-EG0712:fte-0/1/5	Green	Green Blinking
IC-EG0712:fte-0/1/6	Green	Green
IC-EG0712:fte-0/1/7	Green	Green
IC-EG0712:fte-0/1/8	Green	Green Blinking
IC-EG0712:fte-0/1/9	Green	Green Blinking
IC-EG0712:fte-0/1/10	Green	Green Blinking

show chassis led interconnect-device (QFabric System - QFX3008-I Interconnect Device)

```
user@switch> show chassis led interconnect-device IC-EG0712
Front Panel Module Information
```

LEDs status:

```
Status LED: Green
Power LED : Yellow Blinking
Major Alarm LED: Red
Minor Alarm LED: Yellow
Fan 0 LED : Green
Fan 1 LED : Green
Fan 2 LED : Green
Fan 3 LED : Green
Fan 4 LED : Green
Fan 5 LED : Green
Fan 6 LED : Green
Fan 7 LED : Green
Fan 8 LED : Green
Fan 9 LED : Green
PEM 0 LED : Green
PEM 1 LED : Green
PEM 2 LED : Green
PEM 3 LED : off
PEM 4 LED : Yellow Blinking
PEM 5 LED : off
```

```
LED info for: CB - 0
```

LEDs status:

```
Status LED: Green
Mastership LED: Green
```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

```
LED info for: CB - 1
```

LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-0/0/0	Green	N/A
IC-F4899:fte-0/0/1	Green	N/A
IC-F4899:fte-0/0/2	Green	N/A
IC-F4899:fte-0/0/3	Green	N/A
IC-F4899:fte-0/0/4	Green	N/A
IC-F4899:fte-0/0/5	Green	N/A
IC-F4899:fte-0/0/6	Green	N/A
IC-F4899:fte-0/0/7	Green	N/A
IC-F4899:fte-0/0/8	Green	N/A
IC-F4899:fte-0/0/9	Green	N/A
IC-F4899:fte-0/0/10	Green	N/A
IC-F4899:fte-0/0/11	Green	N/A
IC-F4899:fte-0/0/12	Green	N/A
IC-F4899:fte-0/0/13	Green	N/A
IC-F4899:fte-0/0/14	Green	N/A
IC-F4899:fte-0/0/15	Green	N/A

LED info for: FC 1 FPC - 1

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-1/0/0	Green	N/A
IC-F4899:fte-1/0/1	Green	N/A

LED info for: RC 2 FPC - 10

LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

LEDs status:

Status LED: Green

show chassis location

List of Syntax	Syntax on page 808 Syntax (TX Matrix Router) on page 808 Syntax (TX Matrix Plus Router) on page 808 Syntax (MX Series Router) on page 808 Syntax (QFX Series) on page 808
Syntax	show chassis location
Syntax (TX Matrix Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc number lcc number) lcc number scc>
Syntax (TX Matrix Plus Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc number lcc number) lcc number sfc number>
Syntax (MX Series Router)	show chassis location <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show chassis location <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the physical location of the chassis. This command can only be used on the master Routing Engine.
Options	none —Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached routers. all-members —(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration. fpc —(TX Matrix router and TX Matrix Plus router only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs). interconnect-device <i>name</i> —(QFabric systems only) (Optional) Display the physical location of the Interconnect device. interface by-name <i>name</i> —(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (line-card chassis) number associated

with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and router (line-card chassis) number associated with the specified interface.

interface by-slot fpc *number* lcc *number*—(TX Matrix and TX Matrix Plus router only)

(Optional) On a TX Matrix router, display the global FPC number of an interface by specifying its local FPC number and T640 router (line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and router (line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. On TX Matrix Plus router with 3D SIBs, the value is **0** through **63**. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace *number* with a value from **0** through **7**.
- For **lcc**, replace *number* with a value from **0** through **7**.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the physical location of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display the physical location of the Node device.

scc—(TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

List of Sample Output [show chassis location on page 810](#)
[show chassis location fpc \(TX Matrix Router\) on page 810](#)
[show chassis location interface by-slot \(TX Matrix Router\) on page 811](#)
[show chassis location fpc \(TX Matrix Plus Router\) on page 811](#)
[show chassis location interface by-slot \(TX Matrix Plus Router\) on page 811](#)
[show chassis location \(QFX3500 Switches\) on page 811](#)
[show chassis location \(QFabric Systems\) on page 811](#)

Output Fields [Table 66 on page 810](#) lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

Table 66: show chassis location Output Fields

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is 0 through 31. On TX Matrix Plus router with 3D SIBs the value is 0 through 63.
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular router connected to the TX Matrix Plus router.
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular router.

Sample Output

show chassis location

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location fpc (TX Matrix Router)

```
user@host> show chassis location fpc
Global FPC   LCC   Local FPC
    17         2       1
    21         2       5
```

show chassis location interface by-slot (TX Matrix Router)

```
user@host> show chassis location interface by-slot fpc 1 lcc 1
Global FPC: 9
```

show chassis location fpc (TX Matrix Plus Router)

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    0          0        0
    1          0        1
```

show chassis location interface by-slot (TX Matrix Plus Router)

```
user@host> show chassis location interface by-slot fpc 2 lcc 1
Global FPC: 10
```

show chassis location (QFX3500 Switches)

```
user@switch> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location (QFabric Systems)

```
user@switch> show chassis location interconnect-device interconnect1
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis mac-addresses

List of Syntax	Syntax on page 812 Syntax (TX Matrix Router) on page 812 Syntax (TX Matrix Plus Router) on page 812 Syntax (MX Series Router) on page 812 Syntax (MX2010 3D Universal Edge Routers) on page 812 Syntax (MX2020 3D Universal Edge Routers) on page 812 Syntax (QFX Series) on page 812 Syntax (ACX Series Universal Access Routers) on page 812
Syntax	show chassis mac-addresses
Syntax (TX Matrix Router)	show chassis mac-addresses <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis mac-addresses <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show chassis mac-addresses <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis mac-addresses
Syntax (MX2020 3D Universal Edge Routers)	show chassis mac-addresses
Syntax (QFX Series)	show chassis mac-addresses <interconnect-device <i>name</i> > <node-group <i>name</i> >
Syntax (ACX Series Universal Access Routers)	show chassis mac-addresses
Release Information	Command introduced before JUNOS Release 7.4. Command introduced in JUNOS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display the media access control (MAC) addresses for the router, switch chassis, or switch.

Options **none**—(TX Matrix, TX Matrix Plus routers, and the QFX Series) Display the MAC addresses for the router chassis or switch. On a TX Matrix router, display MAC addresses on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display MAC addresses on the TX Matrix Plus router and its attached routers.

all-members—(MX Series routers only) (Optional) Display the MAC addresses for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display the MAC addresses for the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display MAC addresses for a specified T640 router (line-card chassis) that is connected to the TX Matrix Plus router. On a TX Matrix Plus router, display MAC addresses for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the MAC addresses for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the MAC addresses for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display the MAC addresses for the specified Node group.

scc—(TX Matrix routers only) (Optional) Display MAC addresses for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display MAC addresses for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

Related Documentation [• ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping](#)

List of Sample Output [show chassis mac-addresses on page 814](#)

[show chassis mac-addresses \(MX2010 Router\) on page 814](#)
[show chassis mac-addresses \(MX2020 Router\) on page 814](#)
[show chassis mac-addresses \(TX Matrix Router\) on page 815](#)
[show chassis mac-addresses \(TX Matrix Plus Router\) on page 815](#)
[show chassis mac-addresses \(QFX3500 Switches\) on page 816](#)
[show chassis mac-addresses interconnect-device \(QFabric Systems\) on page 816](#)
[show chassis mac-addresses node-group \(QFabric Systems\) on page 816](#)
[show chassis mac-addresses \(ACX2000 Universal Access Router\) on page 816](#)

Output Fields Table 67 on page 814 lists the output fields for the **show chassis mac-addresses** command. Output fields are listed in the approximate order in which they appear.

Table 67: show chassis mac-addresses Output Fields

Field Name	Field Description
MAC address information	
Public base address	Base address of the MAC addresses allocated to this router or switch.
Public count	Number of allocated public addresses.
Private base address	Base address of the private MAC addresses allocated to this router or switch.
Private count	Number of allocated private addresses.

Sample Output

show chassis mac-addresses

```

user@host> show chassis mac-addresses
MAC address information
  Public base address  0:90:69:0:4:0
  Public count         1008
  Private base address 0:90:69:0:7:f0
  Private count        16

```

show chassis mac-addresses (MX2010 Router)

```

user@host> show chassis mac-addresses
MAC address information:
  Public base address  64:87:88:04:50:00
  Public count         1984
  Private base address 64:87:88:04:57:c0
  Private count        64

```

show chassis mac-addresses (MX2020 Router)

```

user@host> show chassis mac-addresses
MAC address information:
  Public base address  2c:21:72:70:20:00
  Public count         4032
  Private base address 2c:21:72:70:2f:c0
  Private count        64

```

show chassis mac-addresses (TX Matrix Router)

```

user@host> show chassis mac-addresses
scc-re0:
-----
MAC address information:
  Public base address  00:05:85:9e:cc:00
  Public count         8064
  Private base address 00:05:85:9e:eb:80
  Private count        128
lcc0-re0:
-----
MAC address information:
  Public base address  00:05:85:68:98:00
  Public count         2032
  Private base address 00:05:85:68:9f:f0
  Private count        16
lcc2-re0:
-----
MAC address information:
  Public base address  00:05:85:68:78:00
  Public count         2032
  Private base address 00:05:85:68:7f:f0
  Private count        16

```

show chassis mac-addresses (TX Matrix Plus Router)

```

user@host> show chassis mac-addresses
sfc0-re0:
-----
MAC address information:
  Public base address  00:1d:b5:14:00:00
  Public count         65023
  Private base address 00:1d:b5:14:fd:ff
  Private count        512
lcc0-re0:
-----
MAC address information:
  Public base address  00:1f:12:7a:84:00
  Public count         2032
  Private base address 00:1f:12:7a:8b:f0
  Private count        16
lcc1-re0:
-----
MAC address information:
  Public base address  00:22:83:42:48:00
  Public count         2032
  Private base address 00:22:83:42:4f:f0
  Private count        16
lcc2-re0:
-----
MAC address information:
  Public base address  00:1f:12:c3:58:00
  Public count         2032
  Private base address 00:1f:12:c3:5f:f0
  Private count        16
lcc3-re0:

```

```
-----  
MAC address information:  
Public base address    00:21:59:ef:b8:00  
Public count           2032  
Private base address   00:21:59:ef:bf:f0  
Private count           16
```

show chassis mac-addresses (QFX3500 Switches)

```
user@switch> show chassis mac-addresses  
MAC address information:  
Public base address 02:00:08:00:00:00  
Public count 512  
Private base address 02:00:00:00:00:00  
Private count 64
```

show chassis mac-addresses interconnect-device (QFabric Systems)

```
user@switch> show chassis mac-addresses interconnect-device interconnect1  
MAC address information:  
Public base address    00:1f:12:30:9c:c0  
Public count           58  
Private base address   00:1f:12:30:9c:fa  
Private count           6
```

show chassis mac-addresses node-group (QFabric Systems)

```
user@switch> show chassis mac-addresses node-group NW-NG-0  
MAC address information:  
-----  
RE:  
FC MAC base    00:11:00:00:00:00  
FC MAC count   2  
VLAN MAC       00:11:00:00:00:09  
EC6007  
Base address   00:00:01:76:00:00  
Count          64  
EC6008  
Base address   00:22:83:22:52:ae  
Count          260
```

show chassis mac-addresses (ACX2000 Universal Access Router)

```
user@switch> show chassis mac-addresses  
MAC address information:  
Public base address    84:18:88:c0:2b:00  
Public count           112  
Private base address   84:18:88:c0:2b:70  
Private count           16
```

show chassis nonstop-upgrade node-group

Syntax	show chassis nonstop-upgrade node-group <i>node-group-name</i>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display the status of the Node group after the most recent nonstop software upgrade (NSSU).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 105 • request system software nonstop-upgrade on page 410
List of Sample Output	show chassis nonstop-upgrade node-group on page 817
Output Fields	Table 68 on page 817 lists the output fields for the show chassis nonstop-upgrade node-group command. Output fields are listed in the approximate order in which they appear.

Table 68: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```

user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
Item           Status           Reason
P1550-C       Online

```

show chassis pic

List of Syntax	Syntax on page 818 Syntax (TX Matrix and TX Matrix Plus Routers) on page 818 Syntax (MX Series Routers) on page 818 Syntax (MX2010 3D Universal Edge Routers) on page 818 Syntax (MX2020 3D Universal Edge Routers) on page 818 Syntax (QFX Series) on page 818 Syntax (ACX Series Universal Access Routers) on page 818
Syntax	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (TX Matrix and TX Matrix Plus Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <fcc <i>number</i>></code>
Syntax (MX Series Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <all-members></code> <code><local></code> <code><member <i>member-id</i>></code>
Syntax (MX2010 3D Universal Edge Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (MX2020 3D Universal Edge Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Syntax (QFX Series)	<code>show chassis pic</code> <code><interconnect-device <i>name</i> (fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>)></code> <code><node-device <i>name</i> pic-slot <i>slot-number</i>></code>
Syntax (ACX Series Universal Access Routers)	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.
Options	fpc-slot <i>slot-number</i> —Display information about the PIC in this particular FPC slot: <ul style="list-style-type: none">On a TX Matrix router, if you specify the number of the T640 router by using the fcc <i>number</i> option (the recommended method), replace <i>slot-number</i> with a value from 0 through 7. Otherwise, replace <i>slot-number</i> with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
user@host> show chassis pic fpc-slot 9 pic-slot 1
```

- M120 routers only—Replace **slot-number** with a value from 0 through 5.
- MX80 routers only—Replace **slot-number** with a value from 0 through 1.
- MX240 routers only—Replace **slot-number** with a value from 0 through 2.
- MX480 routers only—Replace **slot-number** with a value from 0 through 5.
- MX960 routers only—Replace **slot-number** with a value from 0 through 11.
- MX2020 routers only—Replace **slot-number** with a value from 0 through 19.
- Other routers—Replace **slot-number** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **slot-number** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9 (switch's member ID).
 - EX8208 switches—Replace **slot-number** with a value from 0 through 7 (line card).
 - EX8216 switches—Replace **slot-number** with a value from 0 through 15 (line card).
- QFX Series:
 - QFX3500 switches—Replace **slot-number** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
 - QFabric systems—Replace **slot-number** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

all-members—(MX Series routers only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

interconnect-device name—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display PIC information for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

pic-slot *slot-number*—Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

Required Privilege Level

view

Related Documentation

- *request chassis pic*
- [show chassis hardware on page 667](#)
- *Configuring the PIC Type*
- *100-Gigabit Ethernet PIC Overview*

List of Sample Output

[show chassis pic fpc-slot pic-slot on page 823](#)
[show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 823](#)
[show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 823](#)
[show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 823](#)
[show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 823](#)
[show chassis pic fpc-slot 3 pic-slot 0 \(M120 Router\) on page 823](#)
[show chassis pic fpc-slot pic-slot \(MX960 Router Bidirectional Optics\) on page 824](#)
[show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 824](#)
[show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 824](#)

[show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 824](#)
[show chassis pic fpc-slot pic-slot \(MX2020 Router\) on page 825](#)
[show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 825](#)
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 825](#)
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 825](#)
[show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 826](#)
[show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 826](#)
[show chassis pic fpc-slot 0 pic-slot 1 \(4x CHOC3 SONET CE SFP\) on page 826](#)
[show chassis pic fpc-slot 0 pic-slot 0 \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 827](#)
[show chassis pic fpc-slot 3 pic-slot 0 \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 827](#)
[show chassis pic fpc-slot 5 pic-slot 0 \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 827](#)
[show chassis pic fpc-slot 1 pic-slot 0 \(1-port OC192/STM64 MIC with XFP\) on page 827](#)
[show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 828](#)
[show chassis pic fpc-slot pic-slot \(OTN\) on page 828](#)
[show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 828](#)
[show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 828](#)
[show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 828](#)
[show chassis pic fpc-slot 0 pic-slot 1 \(ACX2000 Universal Access Router\) on page 829](#)
[show chassis pic FPC-slot 1 PIC-slot 0 \(MX Routers with Media Services Blade \[MSB\]\) on page 829](#)
[show chassis pic FPC slot 1, PIC slot 2 \(MX Routers with Media Services Blade \[MSB\]\) on page 829](#)

Output Fields [Table 69 on page 821](#) lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 69: show chassis pic Output Fields

Field Name	Field Description
Type	<p>PIC type.</p> <p>NOTE: On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as MIC-3D-1STM64-XFP and with the SONET framing mode, the type is displayed as MIC-3D-1OC192-XFP. By default, the 1-port OC192/STM64 MICs displays the type as MIC-3D-1OC192-XFP.</p>
ASIC type	Type of ASIC on the PIC.
State	<p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> • Online— PIC is online and running. • Offline—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.

Table 69: show chassis pic Output Fields (*continued*)

Field Name	Field Description
Package	(Multiservices PICs only) Services package supported: Layer-2 or Layer-3 .
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: LH , LX , or SX .
PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)	Port-level information for the PIC. <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction
PIC Port Information (MX960 Router Bidirectional Optics)	Port-level information for the PIC. <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> • BX10-10-km bidirectional optics. • BX40-40-km bidirectional optics. • SFP-LX-40-km SFP optics. • Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.
PIC Port Information (Next-Generation SONET/SDH SFP)	Port-level information for the next-generation SONET/SDH SFP PIC. <ul style="list-style-type: none"> • Port—Port number. • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. • Fiber type—Type of fiber: SM (single-mode) or MM (multimode). • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.
Multirate Mode	Rate-selectability status for the MIC: Enabled or Disabled .

Table 69: show chassis pic Output Fields (*continued*)

Field Name	Field Description
Channelization	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

Sample Output

show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
  Type                10x 1GE(LAN), 1000 BASE
  ASIC type           H chip
  State               Online
  PIC version         1.1
  Uptime              1 day, 50 minutes, 58 seconds
PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type        Vendor Name  Part Number
  0         GIGE 1000EX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000EX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (PIC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
  State              Offline

```

show chassis pic fpc-slot pic-slot (FPC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

show chassis pic fpc-slot pic-slot (PIC Not Present)

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
  Type                2x G/E IQ, 1000 BASE
  ASIC type           IQ GE 2 VLAN-TAG FPGA
  State               Online
  PIC version         1.16
  Uptime              3 hours, 3 minutes
PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type        Vendor Name  Part Number
  0         GIGE 1000SX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000SX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (MX960 Router Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                10x 1GE(LAN)
  State                Online
  PIC version          0.0
  Uptime               18 days, 5 hours, 41 minutes, 54 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
1	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
2	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
3	SFP-1000BASE-BX10-D	SM	OCF	TRXBG1LXDBVM2-JW	1490 nm
4	SFP-1000BASE-BX10-D	SM	OCF	TRXBG1LXDBVM2-JW	1490 nm
5	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
6	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
7	SFP-1000BASE-BX10-U	SM	OCF	TRXBG1LXDBBMH-J1	1310 nm
8	SFP-1000BASE-BX10-U	SM	OCF	TRXBG1LXDBBMH-J1	1310 nm
9	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm

show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                1X100GE CFP
  State                Online
  PIC version          2.10
  Uptime               4 minutes, 48 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	100GBASE LR4	SM	FINISAR CORP.	FTLC1181RDN5-J3	1310 nm

Xcvr vendor
 firmware version
 1.8

show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                AS-MXC
  State                Online
  PIC version          1.0
  Uptime               11 hours, 18 minutes, 3 seconds

```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```

user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                1X100GE CFP
  State                Online
  PIC version          0.0
  Uptime               14 hours, 51 seconds

```

show chassis pic fpc-slot pic-slot (MX2020 Router)

```

user@host>show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
  Type                4x 10GE(LAN) SFP+
  State                Online
  PIC version          0.0
  Uptime               1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wave-length	Xcvr
0	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
1	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
2	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
3	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0

show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                100GE SLOT1
  ASIC type            Brooklyn 100GE FPGA
  State                Online
  PIC version          1.3
  Uptime               10 minutes, 44 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	100GBASE LR4	SM	Opnext Inc.	TRC5E20ENFSF000F	1310 nm

show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:
-----
PIC fpc slot 1 pic slot 1 information:
  Type                4x OC-3 SONET, SMIR
  ASIC type            D chip
  State                Online
  PIC version          1.2
  Uptime               5 days, 2 hours, 12 minutes, 8 seconds

```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```

user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:
-----
FPC slot 8, PIC slot 0 information:
  Type                1x 10GE(LAN/WAN)
  State                Online
  Uptime               2 hours, 46 minutes, 23 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	10GBASE ZR	SM	Opnext Inc.	TRF7061BN-LF150	1550 nm
0	10GBASE ZR	SM	FINISAR CORP.	FTRX-1811-3-J2	1550 nm

show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

user@host> show chassis pic fpc-slot 4 pic-slot 0

FPC slot 4, PIC slot 0 information:

```

Type                4x OC-3 1x OC-12 SFP
ASIC type            D FPGA
State                Online
PIC version          1.3
Uptime               1 day, 50 minutes, 4 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	OC48 short reach	SM	FINISAR CORP.	FTRJ1321P1BTL-J2	1310 nm
1	OC3 short reach	MM	OCF	TRPA03MM3BAS-JE	1310 nm
2	OC3 short reach	MM	OCF	TRXA03MM3BAS-JW	1310 nm
3	OC12 inter reach	SM	FINISAR CORP.	FTLF1322P1BTR	1310 nm

show chassis pic fpc-slot pic-slot (12-Port T1/E1)

user@host> show chassis pic fpc-slot 0 pic-slot 3

FPC slot 0, PIC slot 3 information:

```

Type                12x T1/E1 CE
State                Online
PIC version          1.1
CPU load average     1 percent
Interrupt load average 0 percent
Total DRAM size      128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime               1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

user@host> show chassis pic fpc-slot 0 pic-slot 1

FPC slot 0, PIC slot 1 information:

```

Type                4x CHOC3 SONET CE SFP
State                Online
PIC version          1.3
CPU load average     1 percent
Interrupt load average 0 percent
Total DRAM size      128 MB
Memory buffer utilization 99 percent
Memory heap utilization 4 percent
Uptime               1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
1	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
3	OC3 long reach	SM	OPNEXT INC	TRF5456AVLB314	1310 nm

show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                MIC-3D-80C30C12-40C48
  State                Online
  PIC version          1.8
  Uptime               3 days, 22 hours, 3 minutes, 50 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

```

Multirate Mode      Enabled

```

show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                MIC-3D-8CHOC3-4CHOC12
  State                Online
  PIC version          1.9
  Uptime               1 hour, 21 minutes, 24 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
1	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J2	1310 nm
4	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
5	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
6	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
FPC slot 5, PIC slot 0 information:
  Type                MIC-3D-4CHOC3-2CHOC12
  State                Online
  PIC version          1.9
  Uptime               1 hour, 21 minutes

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
3	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                MIC-3D-10C192-XFP
  State                Online
  PIC version          1.2
  Uptime               1 day, 11 hours, 4 minutes, 6 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC192 short reach	n/a	FINISAR CORP.	FTLX1412M3BCL-J3	1310 nm

show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type          MIC-3D-8DS3-E3
  State         Online
  PIC version   1.10
  Uptime       4 days, 1 hour, 29 minutes, 19 seconds
  Channelization Mode Disabled

```

show chassis pic fpc-slot pic-slot (OTN)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
  Type          1x10GE(LAN),OTN
  ASIC type     H chip
  State         Online
  PIC version   1.0
  Uptime       5 minutes, 50 seconds

```

show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type 48x 10G-SFP+ Builtin
  State Online
  Uptime 3 days, 3 hours, 5 minutes, 20 seconds

```

show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```

user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type          16x 40G-GE Builtin
  State         Online
  Uptime       2 hours, 47 minutes, 40 seconds

```

show chassis pic node-device fpc-slot pic-slot (QFabric System)

```

user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
  Type          48x 10G-SFP+Builtin
  State         Online
  Uptime       2 hours, 52 minutes, 37 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
1	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
2	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
3	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
4	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
5	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
6	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
7	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
8	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
9	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

10	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
11	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
12	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
13	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
14	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
15	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
16	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
17	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
18	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
19	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
20	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
21	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
22	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
23	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
24	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
25	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
26	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
27	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
28	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
29	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
30	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
31	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
32	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
33	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
34	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
35	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
36	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
37	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
38	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
39	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
40	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
41	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
42	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
43	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
44	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
45	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
46	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
47	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type                               8x 1GE(LAN) RJ45 Built-in
State                               Online
Uptime                             6 days, 2 hours, 51 minutes, 11 seconds

```

show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
Type                               AS-MSC
State                               Online
PIC version                         1.6
Uptime                             11 hours, 17 minutes, 56 seconds

```

show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                               AS-MXC
State                               Online

```

PIC version	1.0
Uptime	11 hours, 18 minutes, 3 seconds

show chassis routing-engine

List of Syntax	Syntax on page 831 Syntax (EX Series Switches) on page 831 Syntax (T Series routers) on page 831 Syntax (TX Matrix Routers) on page 831 Syntax (TX Matrix Plus Routers) on page 831 Syntax (QFX Series) on page 831 Syntax (MX Series Routers) on page 831 Syntax (MX2010 3D Universal Edge Routers) on page 831 Syntax (MX2020 3D Universal Edge Routers) on page 831 Syntax (ACX Series Universal Access Routers) on page 831
Syntax	show chassis routing-engine <bios <i>slot</i> >
Syntax (EX Series Switches)	show chassis routing-engine < <i>slot</i> >
Syntax (T Series routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (TX Matrix Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> sfc <i>number</i> >
Syntax (QFX Series)	show chassis routing-engine <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (MX Series Routers)	show chassis routing-engine <bios <i>slot</i> > <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (MX2020 3D Universal Edge Routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (ACX Series Universal Access Routers)	show chassis routing-engine

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release in 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display the status of the Routing Engine.
Options	<p>none—Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached routers.</p> <p>all-members—(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.</p> <p>bios—(Optional) Display the (BIOS) firmware version.</p> <p>interconnect-device <i>number</i>—(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none">• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. <p>local—(MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.</p> <p>member <i>member-id</i>—(MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace <i>member-id</i> with a value of 0 or 1.</p> <p>node-device <i>number</i>—(QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.</p>

scc—(TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with 0.

slot—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace **slot** with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.

Required Privilege Level view

Related Documentation

- [request chassis routing-engine master on page 367](#)
- [Configuring Routing Engine Redundancy](#)
- [Switching the Global Master and Backup Roles in a Virtual Chassis Configuration](#)

List of Sample Output

[show chassis routing-engine \(M5 Router\) on page 835](#)
[show chassis routing-engine \(M10 Router\) on page 836](#)
[show chassis routing-engine \(M20 Router\) on page 836](#)
[show chassis routing-engine \(M40 Router\) on page 837](#)
[show chassis routing-engine \(M120 Router\) on page 837](#)
[show chassis routing-engine \(M160 Router\) on page 838](#)
[show chassis routing-engine \(MX240 Router\) on page 838](#)
[show chassis routing-engine \(MX480 Router\) on page 839](#)
[show chassis routing-engine \(MX960 Router\) on page 839](#)
[show chassis routing-engine \(MX2010 Router\) on page 840](#)
[show chassis routing-engine \(MX2020 Router\) on page 840](#)
[show chassis routing-engine \(T320 router\) on page 841](#)
[show chassis routing-engine \(T640 router\) on page 842](#)
[show chassis routing-engine \(T1600 router\) on page 843](#)
[show chassis routing-engine \(T4000 router\) on page 843](#)
[show chassis routing-engine \(TX Matrix Router\) on page 844](#)
[show chassis routing-engine lcc \(TX Matrix Router\) on page 845](#)
[show chassis routing-engine bios \(TX Matrix Router\) on page 846](#)
[show chassis routing-engine \(TX Matrix Plus Router\) on page 846](#)
[show chassis routing-engine lcc \(TX Matrix Plus Router\) on page 847](#)
[show chassis routing-engine bios \(TX Matrix Plus Router\) on page 848](#)
[show chassis routing-engine \(QFX Series\) on page 848](#)
[show chassis routing engine interconnect-device \(QFabric systems\) on page 849](#)
[show chassis routing-engine \(PTX Series Packet Transport Switch\) on page 849](#)
[show chassis routing-engine \(ACX2000 Universal Access Router\) on page 850](#)
[show chassis routing-engine \(ACX1000 Universal Access Router\) on page 850](#)

Output Fields [Table 70 on page 834](#) lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 70: show chassis routing-engine Output Fields

Field Name	Field Description
Slot	(Systems with single and multiple Routing Engines) Slot number.
Current state	(Systems with multiple Routing Engines) Current state of the Routing Engine: Master , Backup , or Disabled .
Election priority	(Systems with multiple Routing Engines) Election priority for the Routing Engine: Master or Backup .
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	Temperature of the CPU.
DRAM	Total DRAM available to the Routing Engine's processor. Starting with Junos OS Release 12.3R1, the DRAM field displays both available memory and installed memory.
Memory utilization	Percentage of Routing Engine memory being used.
CPU utilization	Information about the Routing Engine's CPU utilization: <ul style="list-style-type: none"> • User—Percentage of CPU time being used by user processes. • Background—Percentage of CPU time being used by background processes. • Kernel—Percentage of CPU time being used by kernel processes. • Interrupt—Percentage of CPU time being used by interrupts. • Idle—Percentage of CPU time that is idle.
Model	Routing Engine model number.
Serial ID	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
Start time	Time at which the Routing Engine started running.
Uptime	How long the Routing Engine has been running.
Routing Engine BIOS Version	BIOS version being run by the Routing Engine.

Table 70: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
Last reboot reason	<p>Reason for last reboot, including:</p> <ul style="list-style-type: none"> • power cycle/failure—Halt of the Routing Engine using the halt command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the request system halt command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard. • watchdog—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered. • reset-button reset—(Not available on the J Series router or EX Series switch) Reboot due to pressing of the reset button on the Routing Engine. • power-button hard power off—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the request system power-off command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software. • misc hardware reason—Reboot due to miscellaneous hardware reasons. • thermal shutdown—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations. • hard disk failure—Reboot due to a hard disk or solid-state drive (SSD) failure. • reset from debugger—Reboot due to reset from the debugger. • chassis control reset—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the restart chassis-control command. • bios auto recovery reset—Reboot due to a BIOS auto-recovery reset. • could not be determined—Reboot due to an undetermined reason. • Router rebooted after a normal shutdown—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the request system reboot command. You can enter this command to reboot the chassis or specific Routing Engines.
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

Sample Output

show chassis routing-engine (M5 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                       768 MB
  Memory utilization         21 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent

```

```

Idle 100 percent
Model RE-2.0
Serial ID 31000007349bf701
Start time 2003-12-04 09:42:17 PST
Uptime 26 days, 1 hour, 12 minutes, 27 seconds
Last reboot reason Router rebooted after a normal shutdown
Load averages: 1 minute 5 minute 15 minute
                0.00 0.01 0.00

```

show chassis routing-engine (M10 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature 25 degrees C / 77 degrees F
  DRAM 768 MB
  Memory utilization 21 percent
  CPU utilization:
    User 0 percent
    Background 0 percent
    Kernel 0 percent
    Interrupt 0 percent
    Idle 100 percent
  Model RE-2.0
  Serial ID 31000007349bf701
  Start time 2003-12-04 09:42:17 PST
  Uptime 26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages: 1 minute 5 minute 15 minute
                  0.00 0.01 0.00

```

show chassis routing-engine (M20 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state Master
    Election priority Master (default)
    Temperature 29 degrees C / 84 degrees F
    DRAM 768 MB
    Memory utilization 20 percent
    CPU utilization:
      User 1 percent
      Background 0 percent
      Kernel 2 percent
      Interrupt 0 percent
      Idle 97 percent
    Model RE-2.0
    Serial ID 58000007348d9a01
    Start time 2003-12-30 07:05:47 PST
    Uptime 3 hours, 41 minutes, 14 seconds
    Last reboot reason Router rebooted after a normal shutdown
    Load averages: 1 minute 5 minute 15 minute
                    0.00 0.02 0.00

  Routing Engine status:
    Slot 1:
      Current state Backup
      Election priority Backup (default)
      Temperature 29 degrees C / 84 degrees F
      DRAM 768 MB
      Memory utilization 0 percent
      CPU utilization:

```



```

User                0 percent
Background          0 percent
Kernel              1 percent
Interrupt           0 percent
Idle                99 percent
Model               RE-2.0
Serial ID           d800000734745701
Start time          2003-06-17 16:37:33 PDT
Uptime              195 days, 18 hours, 47 minutes, 9 seconds
Last reboot reason   Router rebooted after a normal shutdown

```

show chassis routing-engine (M40 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature        25 degrees C / 77 degrees F
  DRAM               768 MB
  Memory utilization  21 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle              100 percent
  Model              RE-2.0
  Serial ID           31000007349bf701
  Start time          2003-12-04 09:42:17 PST
  Uptime              26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason   Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute  15 minute
                      0.00        0.01    0.00

```

show chassis routing-engine (M120 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        46 degrees C / 114 degrees F
  CPU temperature     44 degrees C / 111 degrees F
  DRAM               2048 MB
  Memory utilization  18 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             5 percent
    Interrupt          0 percent
    Idle              95 percent
  Model              RE-A-1000
  Serial ID           1000621154
  Start time          2006-10-31 17:10:05 PST
  Uptime              14 minutes, 31 seconds
  Last reboot reason   Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute  15 minute
                      0.02        0.07    0.07

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        45 degrees C / 113 degrees F

```

```
CPU temperature      42 degrees C / 107 degrees F
DRAM                2048 MB
Memory utilization   15 percent
CPU utilization:
  User               0 percent
  Background         0 percent
  Kernel             0 percent
  Interrupt          0 percent
  Idle               100 percent
Model               RE-A-1000
Serial ID            1000621151
Start time           2006-10-31 17:10:04 PST
Uptime               14 minutes, 30 seconds
Last reboot reason   Router rebooted after a normal shutdown
```

show chassis routing-engine (M160 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        43 degrees C / 109 degrees F
  DRAM               2048 MB
  Memory utilization  11 percent
  CPU utilization:
    User             1 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             97 percent
  Model              RE-3.0
  Serial ID          210865700403
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 24 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:    1 minute   5 minute   15 minute
                   0.24       0.13       0.04

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        40 degrees C / 104 degrees F
  DRAM               2048 MB
  Memory utilization  9 percent
  CPU utilization:
    User             0 percent
    Background       0 percent
    Kernel           0 percent
    Interrupt        0 percent
    Idle             100 percent
  Model              RE-3.0
  Serial ID          210865700332
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 21 seconds
  Last reboot reason Router rebooted after a normal shutdown
```

show chassis routing-engine (MX240 Router)

```
user@host> show chassis routing-engine
```

```

Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
  Temperature             40 degrees C / 104 degrees F
  CPU temperature         47 degrees C / 116 degrees F
  DRAM                    3584 MB
  Memory utilization      7 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  100 percent
  Model                   RE-S-2000
  Serial ID               1000703522
  Start time              2007-12-19 10:35:40 PST
  Uptime                  16 days, 3 hours, 15 minutes, 23 seconds
  Last reboot reason      Router rebooted after a normal shutdown

```

show chassis routing-engine (MX480 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             41 degrees C / 105 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                    2048 MB
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                   RE-S-1300
  Serial ID               1000697044
  Start time              2008-01-04 06:46:08 PST
  Uptime                  8 hours, 17 minutes, 16 seconds
  Last reboot reason      Router rebooted after a normal shutdown

```

show chassis routing-engine (MX960 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             37 degrees C / 98 degrees F
  CPU temperature         37 degrees C / 98 degrees F
  DRAM                    2048 MB
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             0 percent
    Idle                  96 percent
  Model                   RE-S-1300

```

Serial ID	1000617944
Start time	2006-10-26 12:37:13 PDT
Uptime	6 days, 4 hours, 59 minutes, 40 seconds
Last reboot reason	Router rebooted after a normal shutdown
Load averages:	1 minute 5 minute 15 minute
	0.16 0.08 0.02

show chassis routing-engine (MX2010 Router)

```
user@host> show chassis routing-engine
```

Routing Engine status:

Slot 0:

Current state	Master
Election priority	Master (default)
Temperature	3 degrees C / 37 degrees F
CPU temperature	3 degrees C / 37 degrees F
DRAM	17152 MB
Memory utilization	13 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	4 percent
Interrupt	2 percent
Idle	95 percent
Model	RE-S-1800x4
Serial ID	9009099704
Start time	2012-10-02 14:33:32 PDT
Uptime	14 hours, 39 minutes, 39 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.06 0.05 0.01

Routing Engine status:

Slot 1:

Current state	Backup
Election priority	Backup (default)
Temperature	1 degrees C / 33 degrees F
CPU temperature	2 degrees C / 35 degrees F
DRAM	17152 MB
Memory utilization	11 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	0 percent
Interrupt	0 percent
Idle	100 percent
Model	RE-S-1800x4
Serial ID	9009099706
Start time	2012-10-02 10:36:06 PDT
Uptime	18 hours, 36 minutes, 57 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.01 0.00 0.00

show chassis routing-engine (MX2020 Router)

```
user@host> show chassis routing-engine
```

Routing Engine status:

Slot 0:

Current state	Master
Election priority	Master (default)

```

Temperature          6 degrees C / 42 degrees F
CPU temperature       6 degrees C / 42 degrees F
DRAM                 17152 MB
Memory utilization    14 percent
CPU utilization:
  User                1 percent
  Background          0 percent
  Kernel              7 percent
  Interrupt            2 percent
  Idle                91 percent
Model                RE-S-1800x4
Serial ID             9009089704
Start time            2012-10-02 11:05:24 PDT
Uptime                2 days, 15 hours, 49 minutes, 13 seconds
Last reboot reason    Router rebooted after a normal shutdown.
Load averages:        1 minute   5 minute   15 minute
                      0.10       0.05       0.01

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority     Backup (default)
  Temperature          7 degrees C / 44 degrees F
  CPU temperature       5 degrees C / 41 degrees F
  DRAM                 17152 MB
  Memory utilization    12 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              0 percent
    Interrupt            0 percent
    Idle                99 percent
  Model                RE-S-1800x4
  Serial ID             9009094138
  Start time            2012-10-02 11:09:57 PDT
  Uptime                2 days, 15 hours, 44 minutes, 27 seconds
  Last reboot reason    Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute
                        0.00       0.00       0.00

```

show chassis routing-engine (T320 router)

```

user@host> show chassis routing-engine
Slot 0:
  Current state        Master
  Election priority     Master (default)
  Temperature          51 degrees C / 123 degrees F
  CPU temperature       55 degrees C / 131 degrees F
  DRAM                 3584 MB
  Memory utilization    11 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              2 percent
    Interrupt            0 percent
    Idle                97 percent
  Model                RE-A-2000
  Serial ID             9009010618
  Start time            2012-10-10 01:24:05 PDT
  Uptime                5 days, 10 hours, 49 minutes, 23 seconds
  Last reboot reason    0x1:power cycle/failure
  Load averages:        1 minute   5 minute   15 minute

```

```

                                0.00      0.05      0.04
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  45 degrees C / 113 degrees F
  CPU temperature              48 degrees C / 118 degrees F
  DRAM                        3584 MB
  Memory utilization           9 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 0 percent
    Idle                      100 percent
  Model                       RE-A-2000
  Serial ID                   9009003642
  Start time                  2012-10-10 01:24:04 PDT
  Uptime                      5 days, 10 hours, 49 minutes, 28 seconds
  Last reboot reason          0x1:power cycle/failure

```

show chassis routing-engine (T640 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                  50 degrees C / 122 degrees F
  CPU temperature              58 degrees C / 136 degrees F
  DRAM                        3584 MB
  Memory utilization           14 percent
  CPU utilization:
    User                      1 percent
    Background                0 percent
    Kernel                    4 percent
    Interrupt                 1 percent
    Idle                      95 percent
  Model                       RE-A-2000
  Serial ID                   1000686556
  Start time                  2012-10-10 01:24:02 PDT
  Uptime                      5 days, 10 hours, 50 minutes, 27 seconds
  Last reboot reason          0x1:power cycle/failure
  Load averages:              1 minute   5 minute   15 minute
                                1.24      0.33      0.12

Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  44 degrees C / 111 degrees F
  CPU temperature              49 degrees C / 120 degrees F
  DRAM                        3584 MB
  Memory utilization           12 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 1 percent
    Idle                      99 percent
  Model                       RE-A-2000
  Serial ID                   1000702739

```

```

Start time          2012-10-10 01:24:02 PDT
Uptime              5 days, 10 hours, 50 minutes, 26 seconds
Last reboot reason  0x1:power cycle/failure

```

show chassis routing-engine (T1600 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            48 degrees C / 118 degrees F
  CPU temperature        58 degrees C / 136 degrees F
  DRAM                   3584 MB
  Memory utilization     13 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               3 percent
    Interrupt            1 percent
    Idle                 96 percent
  Model                  RE-A-2000
  Serial ID              1000704521
  Start time             2012-10-10 01:23:41 PDT
  Uptime                 5 days, 10 hours, 46 minutes, 56 seconds
  Last reboot reason     0x1:power cycle/failure
  Load averages:        1 minute   5 minute   15 minute
                        0.05        0.03        0.01

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            44 degrees C / 111 degrees F
  CPU temperature        48 degrees C / 118 degrees F
  DRAM                   3584 MB
  Memory utilization     12 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-A-2000
  Serial ID              9009006579
  Start time             2012-10-10 01:23:42 PDT
  Uptime                 5 days, 10 hours, 46 minutes, 54 seconds
  Last reboot reason     0x1:power cycle/failure

```

show chassis routing-engine (T4000 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            33 degrees C / 91 degrees F
  CPU temperature        50 degrees C / 122 degrees F
  DRAM                   8960 MB
  Memory utilization     18 percent
  CPU utilization:
    User                 0 percent

```

```

        Background          0 percent
        Kernel              4 percent
        Interrupt           1 percent
        Idle                95 percent
    Model                  RE-DUO-1800
    Serial ID              P737F-002248
    Start time              2012-02-09 22:49:53 PST
    Uptime                  2 hours, 21 minutes, 35 seconds
    Last reboot reason      Router rebooted after a normal shutdown.
    Load averages:         1 minute   5 minute   15 minute
                           0.00        0.04        0.00

Routing Engine status:
Slot 1:
    Current state          Backup
    Election priority      Backup (default)
    Temperature            32 degrees C / 89 degrees F
    CPU temperature        46 degrees C / 114 degrees F
    DRAM                   8960 MB
    Memory utilization     24 percent
    CPU utilization:
        User               0 percent
        Background         0 percent
        Kernel             0 percent
        Interrupt          0 percent
        Idle               99 percent
    Model                  RE-DUO-1800
    Serial ID              P737F-002653
    Start time              2012-02-08 20:12:51 PST
    Uptime                  1 day, 4 hours, 58 minutes, 28 seconds
    Last reboot reason      Router rebooted after a normal shutdown.

```

show chassis routing-engine (TX Matrix Router)

```

user@host> show chassis routing-engine
scc-re0:
-----
Routing Engine status:
Slot 0:
    Current state          Master
    Election priority      Master (default)
    Temperature            34 degrees C / 93 degrees F
    CPU temperature        33 degrees C / 91 degrees F
    DRAM                   2048 MB
    Memory utilization     12 percent
    CPU utilization:
        User               0 percent
        Background         0 percent
        Kernel             2 percent
        Interrupt          0 percent
        Idle               98 percent
    Model                  RE-4.0
    Serial ID              P11123900153
    Start time              2004-08-05 18:42:05 PDT
    Uptime                  9 days, 22 hours, 49 minutes, 50 seconds
    Last reboot reason      Router rebooted after a normal shutdown
    Load averages:         1 minute   5 minute   15 minute
                           0.00        0.08        0.07

lcc0-re0:
-----
Routing Engine status:

```



```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             33 degrees C / 91 degrees F
  CPU temperature         30 degrees C / 86 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-3.0
  Serial ID               210865700363
  Start time             2004-08-05 18:42:05 PDT
  Uptime                 9 days, 22 hours, 48 minutes, 20 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute  5 minute 15 minute
                        0.00      0.02   0.00

```

```
lcc2-re0:
```

```
-----
Routing Engine status:
```

```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             34 degrees C / 93 degrees F
  CPU temperature         35 degrees C / 95 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-4.0
  Serial ID               P11123900126
  Start time             2004-08-05 18:42:05 PDT
  Uptime                 9 days, 22 hours, 49 minutes, 4 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute  5 minute 15 minute
                        0.01      0.01   0.0

```

show chassis routing-engine lcc (TX Matrix Router)

```
user@host> show chassis routing-engine 0 lcc 0
```

```
lcc0-re0:
```

```
-----
Routing Engine status:
```

```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             33 degrees C / 91 degrees F
  CPU temperature         30 degrees C / 86 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent

```

```

Kernel                1 percent
Interrupt              0 percent
Idle                  98 percent
Model                 RE-3.0
Serial ID              210865700363
Start time             2004-08-05 18:42:05 PDT
Uptime                7 days, 22 hours, 49 minutes, 6 seconds
Last reboot reason     Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                       0.00       0.00       0.00

```

show chassis routing-engine bios (TX Matrix Router)

```

user@host> show chassis routing-engine bios
scc-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.0
lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.17
lcc2-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.0

```

show chassis routing-engine (TX Matrix Plus Router)

```

user@host> show chassis routing-engine
sfc0-re0:

```

```

-----
Routing Engine status:

```

Slot 0:

```

Current state          Master
Election priority       Master (default)
Temperature             27 degrees C / 80 degrees F
CPU temperature         42 degrees C / 107 degrees F
DRAM                   3327 MB
Memory utilization      12 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                2 percent
  Interrupt              0 percent
  Idle                  98 percent
Model                  RE-TXP-SFC
Serial ID               737A-1024
Start time              2009-05-11 17:39:49 PDT
Uptime                  3 hours, 45 minutes, 25 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:         1 minute   5 minute   15 minute
                       0.00       0.00       0.00

```

```

Routing Engine status:

```

Slot 1:

```

Current state          Backup
Election priority       Backup (default)
Temperature             29 degrees C / 84 degrees F
CPU temperature         43 degrees C / 109 degrees F
DRAM                   3327 MB
Memory utilization      11 percent
CPU utilization:
  User                  0 percent
  Background            0 percent

```

```

Kernel                0 percent
Interrupt              0 percent
Idle                  100 percent
Model                 RE-TXP-SFC
Serial ID              737A-1024
Start time             2009-05-11 17:08:54 PDT
Uptime                 4 hours, 16 minutes, 52 seconds
Last reboot reason     0x1:power cycle/failure

lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state        Master
  Election priority    Master (default)
  Temperature           30 degrees C / 86 degrees F
  CPU temperature       43 degrees C / 109 degrees F
  DRAM                  3327 MB
  Memory utilization    9 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              2 percent
    Interrupt           0 percent
    Idle                98 percent
  Model                 RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-11 17:40:32 PDT
  Uptime                 3 hours, 44 minutes, 51 seconds
  Last reboot reason     Router rebooted after a normal shutdown.
  Load averages:        1 minute  5 minute 15 minute
                        0.00      0.00    0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)
  Temperature           30 degrees C / 86 degrees F
  CPU temperature       43 degrees C / 109 degrees F
  DRAM                  3327 MB
  Memory utilization    9 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              0 percent
    Interrupt           0 percent
    Idle                100 percent
  Model                 RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-06 17:31:32 PDT
  Uptime                 5 days, 3 hours, 54 minutes, 19 seconds
  Last reboot reason     Router rebooted after a normal shutdown.

```

show chassis routing-engine lcc (TX Matrix Plus Router)

```

user@host> show chassis routing-engine 0 lcc 0
lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state        Master
  Election priority    Master (default)

```

```

Temperature          30 degrees C / 86 degrees F
CPU temperature       43 degrees C / 109 degrees F
DRAM                 3327 MB
Memory utilization    9 percent
CPU utilization:
  User                0 percent
  Background          0 percent
  Kernel              2 percent
  Interrupt            0 percent
  Idle                98 percent
Model                RE-TXP-LCC
Serial ID             737F-1024
Start time            2009-05-11 17:40:32 PDT
Uptime                3 hours, 45 minutes, 26 seconds
Last reboot reason    Router rebooted after a normal shutdown.
Load averages:        1 minute   5 minute   15 minute
                      0.00       0.00       0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)
  Temperature          30 degrees C / 86 degrees F
  CPU temperature       43 degrees C / 109 degrees F
  DRAM                 3327 MB
  Memory utilization    9 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              0 percent
    Interrupt            0 percent
    Idle                100 percent
  Model                RE-TXP-LCC
  Serial ID             737F-1024
  Start time            2009-05-06 17:31:32 PDT
  Uptime                5 days, 3 hours, 54 minutes, 59 seconds
  Last reboot reason    Router rebooted after a normal shutdown.

```

show chassis routing-engine bios (TX Matrix Plus Router)

```

user@host> show chassis routing-engine bios
sfc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.Z

```

```

lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.N

```

show chassis routing-engine (QFX Series)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state Master
  Election priority Master (default)
  DRAM 2820 MB
  Memory utilization 49 percent
  CPU utilization:
    User 1 percent
    Background 0 percent
    Kernel 1 percent

```

```

Interrupt 0 percent
Idle 97 percent
Model QFX3500-48S4Q
Serial ID S/N ED3709
Uptime 3 days, 4 hours, 29 minutes, 42 seconds
Last reboot reason 0x200:chassis control reset
Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19

```

show chassis routing engine interconnect-device (QFabric systems)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             48 degrees C / 118 degrees F
  DRAM                   3312 MB
  Memory utilization      63 percent
  CPU utilization:
    User                  14 percent
    Background            0 percent
    Kernel                5 percent
    Interrupt             0 percent
    Idle                  81 percent
  Model                  RE-QFXC08-CB4S
  Serial ID              BUILTIN
  Start time              2011-07-06 13:26:15 UTC
  Uptime                  11 hours, 24 minutes, 57 seconds
  Last reboot reason      0x4:reset-button reset
  Load averages:         1 minute   5 minute   15 minute
                        2.62       2.31       2.28

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             39 degrees C / 102 degrees F
  DRAM                   3312 MB
  Memory utilization      59 percent
  CPU utilization:
    User                  9 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             0 percent
    Idle                  91 percent
  Model                  RE-QFXC08-CB4S
  Serial ID              BUILTIN
  Start time              2011-07-06 13:24:58 UTC
  Uptime                  11 hours, 26 minutes, 18 seconds
  Last reboot reason      0x4:reset-button reset

```

show chassis routing-engine (PTX Series Packet Transport Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             60 degrees C / 140 degrees F
  CPU temperature         76 degrees C / 168 degrees F

```

```

DRAM                                17152 MB
Memory utilization                    11 percent
CPU utilization:
  User                               0 percent
  Background                         0 percent
  Kernel                             4 percent
  Interrupt                           0 percent
  Idle                               95 percent
Model                                RE-DUO-2600
Serial ID                            P737A-002231
Start time                           2011-12-21 16:54:37 PST
Uptime                               25 minutes, 44 seconds
Last reboot reason                    Router rebooted after a normal shutdown.
Load averages:                       1 minute   5 minute  15 minute
                                      0.01       0.02    0.06

Routing Engine status:
Slot 1:
  Current state                       Backup
  Election priority                    Backup (default)
  Temperature                          50 degrees C / 122 degrees F
  CPU temperature                      64 degrees C / 147 degrees F
  DRAM                                17152 MB
  Memory utilization                    10 percent
  CPU utilization:
    User                               0 percent
    Background                         0 percent
    Kernel                             0 percent
    Interrupt                           0 percent
    Idle                               99 percent
  Model                                RE-DUO-2600
  Serial ID                            P737A-002438
  Start time                           2011-12-21 16:52:26 PST
  Uptime                               27 minutes, 49 seconds
  Last reboot reason                    Router rebooted after a normal shutdown.

```

show chassis routing-engine (ACX2000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                          53 degrees C / 127 degrees F
  DRAM                                1536 MB
  Memory utilization                    25 percent
  CPU utilization:
    User                               0 percent
    Background                         0 percent
    Kernel                             0 percent
    Interrupt                           1 percent
    Idle                               99 percent
  Model                                RE-ACX-2000
  Start time                           2012-05-09 00:57:07 PDT
  Uptime                               5 days, 3 hours, 16 minutes, 15 seconds
  Last reboot reason                    Router rebooted after a normal shutdown.
  Load averages:                       1 minute   5 minute  15 minute
                                      0.00       0.03    0.05

```

show chassis routing-engine (ACX1000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                          36 degrees C / 96 degrees F
  DRAM                                768 MB

```

```
Memory utilization      50 percent
CPU utilization:
  User                  3 percent
  Background            0 percent
  Kernel                6 percent
  Interrupt             0 percent
  Idle                  91 percent
Model                  RE-ACX-1000
Start time              2012-05-10 07:12:23 PDT
Uptime                  4 days, 10 hours, 46 minutes, 53 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:          1 minute   5 minute  15 minute
                        0.00       0.00    0.00
```

show chassis temperature-thresholds

List of Syntax	Syntax on page 852 Syntax (TX Matrix Routers) on page 852 Syntax (TX Matrix Plus Routers) on page 852 Syntax (MX Series Routers) on page 852 Syntax (MX2010 3D Universal Edge Routers) on page 852 Syntax (MX2020 3D Universal Edge Routers) on page 852 Syntax (QFX Series) on page 852
Syntax	show chassis temperature-thresholds
Syntax (TX Matrix Routers)	show chassis temperature-thresholds <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis temperature-thresholds <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis temperature-thresholds <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis temperature-thresholds
Syntax (MX2020 3D Universal Edge Routers)	show chassis temperature-thresholds
Syntax (QFX Series)	show chassis temperature-thresholds <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc command introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1 for T4000 Core Routers. Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display chassis temperature threshold settings, in degrees Celsius.
Options	none —Display the temperature threshold details. all-members —(MX Series routers only) (Optional) Display the chassis temperature threshold settings of all member routers in the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the temperature threshold details of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the temperature threshold details of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display the chassis temperature threshold settings of the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display the chassis temperature threshold settings of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Node device.

scc—(TX Matrix routers only) (Optional) Display the temperature threshold details of the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) On TX Matrix Plus routers, display the temperature threshold details of the TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* with 0.

Required Privilege Level view

Related Documentation • [Defining Alarm Thresholds for System Temperature Sensors](#)

List of Sample Output [show chassis temperature-thresholds on page 855](#)
[show chassis temperature-thresholds \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 855](#)
[show chassis temperature-thresholds \(MX2010 Router\) on page 856](#)
[show chassis temperature-thresholds \(MX2020 Router\) on page 858](#)
[show chassis temperature-thresholds \(T4000 Core Routers\) on page 861](#)
[show chassis temperature-thresholds \(TX Matrix Plus Router\) on page 861](#)

[show chassis temperature-thresholds lcc \(TX Matrix Plus Router\) on page 862](#)
[show chassis temperature-thresholds sfc \(TX Matrix Plus Router\) on page 863](#)
[show chassis temperature-thresholds \(TX Matrix Plus routers with 3D SIBs\) on page 863](#)
[show chassis temperature-thresholds \(QFX3500 Switch and QFX3600\) on page 865](#)
[show chassis temperature-thresholds interconnect-device \(QFabric System\) on page 865](#)
[show chassis temperature-thresholds \(PTX5000 Packet Transport Switch\) on page 865](#)
[show chassis temperature-thresholds \(MX Routers with Media Services Blade \[MSB\]\) on page 866](#)

Output Fields Table 71 on page 854 lists the output fields for the **show chassis temperature-thresholds** command. Output fields are listed in the approximate order in which they appear.

Table 71: show chassis temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.
Fan speed	<p>NOTE: On the QFX3500 switch and QFX3600 switch, there are four fan speeds: low, medium-low, medium-high, and high. The fan speed changes at the threshold when going from a low speed to a higher speed. When the fan speed changes from a higher speed to a lower speed, the temperature changes two degrees below the threshold.</p> <p>Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds.</p> <ul style="list-style-type: none"> • Normal—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. <p>NOTE: On a TX Matrix Plus router with 3D SIBs, the threshold temperature at the XF junction is set to 70°C for Normal fan speed, which is less than or equal to 4800 RPM.</p> <ul style="list-style-type: none"> • High—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. <p>NOTE: On a TX Matrix Plus router with 3D SIBs, the threshold temperature at the XF junction is set to 75°C for High fan speed, which is greater than or equal to 5000 RPM.</p> <p>NOTE: For MX480 Routers, there are three fan speeds: Low, Medium, and High.</p> <p>An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.</p>
Yellow alarm	<p>Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm.</p> <ul style="list-style-type: none"> • Normal—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.

Table 71: show chassis temperature-thresholds Output Fields (*continued*)

Field name	Field Description
Red alarm	<p>Temperature threshold settings, in degrees Celsius, that trigger a red alarm.</p> <ul style="list-style-type: none"> • Normal—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.
Fire Shutdown	(T4000 routers, TX Matrix Plus router with 3D SIBs, and PTX Series Packet Transport Switches only)—Temperature threshold settings, in degrees Celsius, for the network device to shut down.

Sample Output

show chassis temperature-thresholds

```
user@host> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	70	80	95	95	110	110
Routing Engine 1	70	80	95	95	110	110
FPC 0	55	60	75	65	90	80
FPC 1	55	60	75	65	90	80
FPC 2	55	60	75	65	90	80
FPC 3	55	60	75	65	90	80
FPC 4	55	60	75	65	90	80
FPC 5	55	60	75	65	90	80
FPC 6	55	60	75	65	90	80
FPC 7	55	60	75	65	90	80
FPC 8	55	60	75	65	90	80
FPC 9	55	60	75	65	90	80
FPC 10	55	60	75	65	90	80
FPC 11	55	60	75	65	90	80

show chassis temperature-thresholds (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host>show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65		
Routing Engine 0	70	80	95	95	110	110		
Routing Engine 1	70	80	95	95	110	110		
FPC 0	55	60	75	65	90	80		
FPC 1	55	60	75	65	90	80		
FPC 2	55	60	75	65	90	80		

```

95
FPC 4          55    60    75    65    90    80
95
FPC 5          55    60    75    65    90    80
95

```

show chassis temperature-thresholds (MX2010 Router)

```

user@host> show chassis temperature-thresholds

```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal	Bad fan
Routing Engine 0	70	80	95	95	110	110	112	
Routing Engine 1	70	80	95	95	110	110	112	
CB 0 IntakeA-Zone0	60	65	78	75	85	80	95	
CB 0 IntakeB-Zone1	60	65	78	75	85	80	95	
CB 0 IntakeC-Zone0	60	65	78	75	85	80	95	
CB 0 ExhaustA-Zone0	60	65	78	75	85	80	95	
CB 0 ExhaustB-Zone1	60	65	78	75	85	80	95	
CB 0 TCBC-Zone0	60	65	78	75	85	80	95	
CB 1 IntakeA-Zone0	60	65	78	75	85	80	95	
CB 1 IntakeB-Zone1	60	65	78	75	85	80	95	
CB 1 IntakeC-Zone0	60	65	78	75	85	80	95	
CB 1 ExhaustA-Zone0	60	65	78	75	85	80	95	
CB 1 ExhaustB-Zone1	60	65	78	75	85	80	95	
CB 1 TCBC-Zone0	60	65	78	75	85	80	95	
SPMB 0 Intake	56	62	75	63	83	76	95	
SPMB 1 Intake	56	62	75	63	83	76	95	
SFB 0 Intake-Zone0	56	62	75	63	82	70	87	
SFB 0 Exhaust-Zone1	56	62	75	63	82	70	87	
SFB 0 IntakeA-Zone0	56	62	75	63	82	70	87	
SFB 0 IntakeB-Zone1	56	62	75	63	82	70	87	
SFB 0 Exhaust-Zone0	56	62	75	63	82	70	87	
SFB 0 SFB-XF2-Zone1	70	80	90	90	107	107	115	
SFB 0 SFB-XF1-Zone0	70	80	90	90	107	107	115	
SFB 0 SFB-XF0-Zone0	70	80	90	90	107	107	115	
SFB 1 Intake-Zone0	56	62	75	63	82	70	87	
SFB 1 Exhaust-Zone1	56	62	75	63	82	70	87	
SFB 1 IntakeA-Zone0	56	62	75	63	82	70	87	
SFB 1 IntakeB-Zone1	56	62	75	63	82	70	87	
SFB 1 Exhaust-Zone0	56	62	75	63	82	70	87	
SFB 1 SFB-XF2-Zone1	70	80	90	90	107	107	115	
SFB 1 SFB-XF1-Zone0	70	80	90	90	107	107	115	
SFB 1 SFB-XF0-Zone0	70	80	90	90	107	107	115	
SFB 2 Intake-Zone0	56	62	75	63	82	70	87	
SFB 2 Exhaust-Zone1	56	62	75	63	82	70	87	
SFB 2 IntakeA-Zone0	56	62	75	63	82	70	87	
SFB 2 IntakeB-Zone1	56	62	75	63	82	70	87	
SFB 2 Exhaust-Zone0	56	62	75	63	82	70	87	
SFB 2 SFB-XF2-Zone1	70	80	90	90	107	107	115	
SFB 2 SFB-XF1-Zone0	70	80	90	90	107	107	115	
SFB 2 SFB-XF0-Zone0	70	80	90	90	107	107	115	
SFB 3 Intake-Zone0	56	62	75	63	82	70	87	
SFB 3 Exhaust-Zone1	56	62	75	63	82	70	87	
SFB 3 IntakeA-Zone0	56	62	75	63	82	70	87	
SFB 3 IntakeB-Zone1	56	62	75	63	82	70	87	
SFB 3 Exhaust-Zone0	56	62	75	63	82	70	87	
SFB 3 SFB-XF2-Zone1	70	80	90	90	107	107	115	
SFB 3 SFB-XF1-Zone0	70	80	90	90	107	107	115	
SFB 3 SFB-XF0-Zone0	70	80	90	90	107	107	115	
SFB 4 Intake-Zone0	56	62	75	63	82	70	87	

SFB 4 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 4 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 4 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 4 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 4 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 4 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 5 Intake-Zone0	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 5 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 5 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 5 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 5 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 5 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 6 Intake-Zone0	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 6 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 6 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 6 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 6 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 6 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 7 Intake-Zone0	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 7 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 7 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 7 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 7 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 7 SFB-XF0-Zone0	70	80	90	90	107	107	115
FPC 0	55	60	75	65	95	80	100
FPC 1	55	60	75	65	90	80	95
FPC 2	55	60	75	65	95	80	100
FPC 3	55	60	75	65	90	80	95
FPC 4	55	60	75	65	90	80	95
FPC 5	55	60	75	65	95	80	100
FPC 6	55	60	75	65	90	80	95
FPC 7	55	60	75	65	95	80	100
FPC 8	55	60	75	65	90	80	95
FPC 9	55	60	75	65	95	80	100
ADC 0 Intake	56	62	75	63	83	76	95
ADC 0 Exhaust	56	62	75	63	83	76	95
ADC 0 ADC-XF1	70	80	90	90	107	107	115
ADC 0 ADC-XF0	70	80	90	90	107	107	115
ADC 1 Intake	56	62	75	63	83	76	95
ADC 1 Exhaust	56	62	75	63	83	76	95
ADC 1 ADC-XF1	70	80	90	90	107	107	115
ADC 1 ADC-XF0	70	80	90	90	107	107	115
ADC 2 Intake	56	62	75	63	83	76	95
ADC 2 Exhaust	56	62	75	63	83	76	95
ADC 2 ADC-XF1	70	80	90	90	107	107	115
ADC 2 ADC-XF0	70	80	90	90	107	107	115
ADC 3 Intake	56	62	75	63	83	76	95
ADC 3 Exhaust	56	62	75	63	83	76	95
ADC 3 ADC-XF1	70	80	90	90	107	107	115
ADC 3 ADC-XF0	70	80	90	90	107	107	115
ADC 4 Intake	56	62	75	63	83	76	95
ADC 4 Exhaust	56	62	75	63	83	76	95
ADC 4 ADC-XF1	70	80	90	90	107	107	115
ADC 4 ADC-XF0	70	80	90	90	107	107	115

ADC 5 Intake	56	62	75	63	83	76	95
ADC 5 Exhaust	56	62	75	63	83	76	95
ADC 5 ADC-XF1	70	80	90	90	107	107	115
ADC 5 ADC-XF0	70	80	90	90	107	107	115
ADC 6 Intake	56	62	75	63	83	76	95
ADC 6 Exhaust	56	62	75	63	83	76	95
ADC 6 ADC-XF1	70	80	90	90	107	107	115
ADC 6 ADC-XF0	70	80	90	90	107	107	115
ADC 7 Intake	56	62	75	63	83	76	95
ADC 7 Exhaust	56	62	75	63	83	76	95
ADC 7 ADC-XF1	70	80	90	90	107	107	115
ADC 7 ADC-XF0	70	80	90	90	107	107	115
ADC 8 Intake	56	62	75	63	83	76	95
ADC 8 Exhaust	56	62	75	63	83	76	95
ADC 8 ADC-XF1	70	80	90	90	107	107	115
ADC 8 ADC-XF0	70	80	90	90	107	107	115
ADC 9 Intake	56	62	75	63	83	76	95
ADC 9 Exhaust	56	62	75	63	83	76	95
ADC 9 ADC-XF1	70	80	90	90	107	107	115
ADC 9 ADC-XF0	70	80	90	90	107	107	115

show chassis temperature-thresholds (MX2020 Router)

```
user@host> show chassis temperature-thresholds
```

	Fan speed		Yellow alarm		Red alarm		Fire Shutdown
	(degrees C)		(degrees C)		(degrees C)		(degrees C)
Item	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Routing Engine 0	70	80	95	95	110	110	112
Routing Engine 1	70	80	95	95	110	110	112
CB 0 IntakeA-Zone0	60	65	78	75	85	80	95
CB 0 IntakeB-Zone1	60	65	78	75	85	80	95
CB 0 IntakeC-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 0 TCBC-Zone0	60	65	78	75	85	80	95
CB 1 IntakeA-Zone0	60	65	78	75	85	80	95
CB 1 IntakeB-Zone1	60	65	78	75	85	80	95
CB 1 IntakeC-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 1 TCBC-Zone0	60	65	78	75	85	80	95
SPMB 0 Intake	56	62	75	63	83	76	95
SPMB 1 Intake	56	62	75	63	83	76	95
SFB 0 Intake-Zone0	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 0 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 0 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 0 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 0 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 0 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 1 Intake-Zone0	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 1 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 1 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 1 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 1 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 1 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 2 Intake-Zone0	56	62	75	63	82	70	87

SFB 2 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 2 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 2 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 2 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 2 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 2 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 2 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 3 Intake-Zone0	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 3 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 3 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 3 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 3 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 3 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 4 Intake-Zone0	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 4 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 4 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 4 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 4 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 4 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 5 Intake-Zone0	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 5 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 5 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 5 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 5 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 5 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 6 Intake-Zone0	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 6 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 6 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 6 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 6 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 6 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 7 Intake-Zone0	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 7 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 7 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 7 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 7 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 7 SFB-XF0-Zone0	70	80	90	90	107	107	115
FPC 0	55	60	75	65	90	80	95
FPC 1	55	60	75	65	90	80	95
FPC 2	55	60	75	65	90	80	95
FPC 3	55	60	75	65	90	80	95
FPC 4	55	60	75	65	90	80	95
FPC 5	55	60	75	65	90	80	95
FPC 6	55	60	75	65	90	80	95
FPC 7	55	60	75	65	90	80	95
FPC 8	55	60	75	65	90	80	95
FPC 9	55	60	75	65	90	80	95
FPC 10	55	60	75	65	90	80	95
FPC 11	55	60	75	65	90	80	95
FPC 12	55	60	75	65	90	80	95
FPC 13	55	60	75	65	90	80	95

FPC 14	55	60	75	65	90	80	95
FPC 15	55	60	75	65	90	80	95
FPC 16	55	60	75	65	90	80	95
FPC 17	55	60	75	65	90	80	95
FPC 18	55	60	75	65	90	80	95
FPC 19	55	60	75	65	90	80	95
ADC 0 Intake	56	62	75	63	83	76	95
ADC 0 Exhaust	56	62	75	63	83	76	95
ADC 0 ADC-XF1	70	80	90	90	107	107	115
ADC 0 ADC-XF0	70	80	90	90	107	107	115
ADC 1 Intake	56	62	75	63	83	76	95
ADC 1 Exhaust	56	62	75	63	83	76	95
ADC 1 ADC-XF1	70	80	90	90	107	107	115
ADC 1 ADC-XF0	70	80	90	90	107	107	115
ADC 2 Intake	56	62	75	63	83	76	95
ADC 2 Exhaust	56	62	75	63	83	76	95
ADC 2 ADC-XF1	70	80	90	90	107	107	115
ADC 2 ADC-XF0	70	80	90	90	107	107	115
ADC 3 Intake	56	62	75	63	83	76	95
ADC 3 Exhaust	56	62	75	63	83	76	95
ADC 3 ADC-XF1	70	80	90	90	107	107	115
ADC 3 ADC-XF0	70	80	90	90	107	107	115
ADC 4 Intake	56	62	75	63	83	76	95
ADC 4 Exhaust	56	62	75	63	83	76	95
ADC 4 ADC-XF1	70	80	90	90	107	107	115
ADC 4 ADC-XF0	70	80	90	90	107	107	115
ADC 5 Intake	56	62	75	63	83	76	95
ADC 5 Exhaust	56	62	75	63	83	76	95
ADC 5 ADC-XF1	70	80	90	90	107	107	115
ADC 5 ADC-XF0	70	80	90	90	107	107	115
ADC 6 Intake	56	62	75	63	83	76	95
ADC 6 Exhaust	56	62	75	63	83	76	95
ADC 6 ADC-XF1	70	80	90	90	107	107	115
ADC 6 ADC-XF0	70	80	90	90	107	107	115
ADC 7 Intake	56	62	75	63	83	76	95
ADC 7 Exhaust	56	62	75	63	83	76	95
ADC 7 ADC-XF1	70	80	90	90	107	107	115
ADC 7 ADC-XF0	70	80	90	90	107	107	115
ADC 8 Intake	56	62	75	63	83	76	95
ADC 8 Exhaust	56	62	75	63	83	76	95
ADC 8 ADC-XF1	70	80	90	90	107	107	115
ADC 8 ADC-XF0	70	80	90	90	107	107	115
ADC 9 Intake	56	62	75	63	83	76	95
ADC 9 Exhaust	56	62	75	63	83	76	95
ADC 9 ADC-XF1	70	80	90	90	107	107	115
ADC 9 ADC-XF0	70	80	90	90	107	107	115
ADC 10 Intake	56	62	75	63	83	76	95
ADC 10 Exhaust	56	62	75	63	83	76	95
ADC 10 ADC-XF1	70	80	90	90	107	107	115
ADC 10 ADC-XF0	70	80	90	90	107	107	115
ADC 11 Intake	56	62	75	63	83	76	95
ADC 11 Exhaust	56	62	75	63	83	76	95
ADC 11 ADC-XF1	70	80	90	90	107	107	115
ADC 11 ADC-XF0	70	80	90	90	107	107	115
ADC 12 Intake	56	62	75	63	83	76	95
ADC 12 Exhaust	56	62	75	63	83	76	95
ADC 12 ADC-XF1	70	80	90	90	107	107	115
ADC 12 ADC-XF0	70	80	90	90	107	107	115
ADC 13 Intake	56	62	75	63	83	76	95
ADC 13 Exhaust	56	62	75	63	83	76	95
ADC 13 ADC-XF1	70	80	90	90	107	107	115

ADC 13 ADC-XF0	70	80	90	90	107	107	115
ADC 14 Intake	56	62	75	63	83	76	95
ADC 14 Exhaust	56	62	75	63	83	76	95
ADC 14 ADC-XF1	70	80	90	90	107	107	115
ADC 14 ADC-XF0	70	80	90	90	107	107	115
ADC 15 Intake	56	62	75	63	83	76	95
ADC 15 Exhaust	56	62	75	63	83	76	95
ADC 15 ADC-XF1	70	80	90	90	107	107	115
ADC 15 ADC-XF0	70	80	90	90	107	107	115
ADC 16 Intake	56	62	75	63	83	76	95
ADC 16 Exhaust	56	62	75	63	83	76	95
ADC 16 ADC-XF1	70	80	90	90	107	107	115
ADC 16 ADC-XF0	70	80	90	90	107	107	115
ADC 17 Intake	56	62	75	63	83	76	95
ADC 17 Exhaust	56	62	75	63	83	76	95
ADC 17 ADC-XF1	70	80	90	90	107	107	115
ADC 17 ADC-XF0	70	80	90	90	107	107	115
ADC 18 Intake	56	62	75	63	83	76	95
ADC 18 Exhaust	56	62	75	63	83	76	95
ADC 18 ADC-XF1	70	80	90	90	107	107	115
ADC 18 ADC-XF0	70	80	90	90	107	107	115
ADC 19 Intake	56	62	75	63	83	76	95
ADC 19 Exhaust	56	62	75	63	83	76	95
ADC 19 ADC-XF1	70	80	90	90	107	107	115
ADC 19 ADC-XF0	70	80	90	90	107	107	115

show chassis temperature-thresholds (T4000 Core Routers)

```
user@host> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)
	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Chassis default	48	54	65	55	75	65	100
Routing Engine 0	55	65	85	85	100	100	102
Routing Engine 1	55	65	85	85	100	100	102
FPC 0	63	68	75	70	90	83	95
FPC 3	63	68	75	70	90	83	95
FPC 5	56	62	75	63	83	76	95
FPC 6	63	68	75	70	90	83	95
SIB 0	64	70	76	72	87	84	95
SIB 1	64	70	76	72	87	84	95
SIB 2	64	70	76	72	87	84	95
SIB 3	64	70	76	72	87	84	95
SIB 4	64	70	76	72	87	84	95

show chassis temperature-thresholds (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds
sfc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
SIB F13 0	64	70	76	72	90	84
SIB F13 3	64	70	76	72	90	84
SIB F13 6	64	70	76	72	90	84
SIB F13 8	64	70	76	72	90	84

SIB F13 11	64	70	76	72	90	84
SIB F13 12	64	70	76	72	90	84
SIB F2S 16	64	70	76	72	90	84
SIB F2S 17	64	70	76	72	90	84
SIB F2S 18	64	70	76	72	90	84
SIB F2S 19	64	70	76	72	90	84
SIB F2S 20	64	70	76	72	90	84
SIB F2S 21	64	70	76	72	90	84
SIB F2S 22	64	70	76	72	90	84
SIB F2S 23	64	70	76	72	90	84
SIB F2S 24	64	70	76	72	90	84
SIB F2S 25	64	70	76	72	90	84
SIB F2S 26	64	70	76	72	90	84
SIB F2S 27	64	70	76	72	90	84
SIB F2S 28	64	70	76	72	90	84
SIB F2S 29	64	70	76	72	90	84
SIB F2S 30	64	70	76	72	90	84
SIB F2S 31	64	70	76	72	90	84
SIB F2S 32	64	70	76	72	90	84
SIB F2S 33	64	70	76	72	90	84
SIB F2S 34	64	70	76	72	90	84
SIB F2S 35	64	70	76	72	90	84

```
lcc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
FPC 7	56	62	75	63	83	76
SIB 0	48	54	65	60	80	75
SIB 1	48	54	65	60	80	75
SIB 2	48	54	65	60	80	75
SIB 3	48	54	65	60	80	75
SIB 4	48	54	65	60	80	75

```
lcc1-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
...						

show chassis temperature-thresholds lcc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds lcc 1
lcc1-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
SIB 0	48	54	65	60	80	75
SIB 1	48	54	65	60	80	75
SIB 2	48	54	65	60	80	75
SIB 3	48	54	65	60	80	75
SIB 4	48	54	65	60	80	75

show chassis temperature-thresholds sfc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds sfc 0
sfc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
SIB F13 0	64	70	76	72	90	84
SIB F13 3	64	70	76	72	90	84
SIB F13 6	64	70	76	72	90	84
SIB F13 8	64	70	76	72	90	84
SIB F13 11	64	70	76	72	90	84
SIB F13 12	64	70	76	72	90	84
SIB F2S 16	64	70	76	72	90	84
SIB F2S 17	64	70	76	72	90	84
SIB F2S 18	64	70	76	72	90	84
SIB F2S 19	64	70	76	72	90	84
SIB F2S 20	64	70	76	72	90	84
SIB F2S 21	64	70	76	72	90	84
SIB F2S 22	64	70	76	72	90	84
SIB F2S 23	64	70	76	72	90	84
SIB F2S 24	64	70	76	72	90	84
SIB F2S 25	64	70	76	72	90	84
SIB F2S 26	64	70	76	72	90	84
SIB F2S 27	64	70	76	72	90	84
SIB F2S 28	64	70	76	72	90	84
SIB F2S 29	64	70	76	72	90	84
SIB F2S 30	64	70	76	72	90	84
SIB F2S 31	64	70	76	72	90	84
SIB F2S 32	64	70	76	72	90	84
SIB F2S 33	64	70	76	72	90	84
SIB F2S 34	64	70	76	72	90	84
SIB F2S 35	64	70	76	72	90	84

show chassis temperature-thresholds (TX Matrix Plus routers with 3D SIBs)

```
user@host> show chassis temperature-thresholds
sfc0-re0:
```

	Fan speed	Yellow alarm	Red alarm	Fire
Shutdown				

(degrees C) Item	(degrees C)		(degrees C)		(degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Normal						
Chassis default	48	54	65	55	75	65
100						
Routing Engine 0	70	75	90	87	102	97
115						
Routing Engine 1	70	75	90	87	102	97
115						
SIB F13 0 Board	60	65	78	75	85	80
95						
SIB F13 0 XF Junction	70	75	82	74	105	100
107						
SIB F13 4 Board	60	65	78	75	85	80
95						
SIB F13 4 XF Junction	70	75	82	74	105	100
107						
SIB F13 6 Board	60	65	78	75	85	80
95						
SIB F13 6 XF Junction	70	75	82	74	105	100
107						
SIB F2S 16 Board	60	65	78	75	85	80
95						
SIB F2S 16 XF Junction	70	75	82	74	105	100
107						
SIB F2S 17 Board	60	65	78	75	85	80
95						
SIB F2S 17 XF Junction	70	75	82	74	105	100
107						
SIB F2S 18 Board	60	65	78	75	85	80
95						
SIB F2S 18 XF Junction	70	75	82	74	105	100
107						
SIB F2S 19 Board	60	65	78	75	85	80
95						
SIB F2S 19 XF Junction	70	75	82	74	105	100
107						
SIB F2S 24 Board	60	65	78	75	85	80
95						
SIB F2S 24 XF Junction	70	75	82	74	105	100
107						
SIB F2S 25 Board	60	65	78	75	85	80
95						
SIB F2S 25 XF Junction	70	75	82	74	105	100
107						
SIB F2S 26 Board	60	65	78	75	85	80
95						
SIB F2S 26 XF Junction	70	75	82	74	105	100
107						
SIB F2S 27 Board	60	65	78	75	85	80
95						
SIB F2S 27 XF Junction	70	75	82	74	105	100
107						
1cc0-re0:						

Shutdown	Fan speed		Yellow alarm		Red alarm	Fire
(degrees C)	(degrees C)		(degrees C)		(degrees C)	

Item	Normal	High	Normal	Bad fan	Normal	Bad fan
Normal						
Chassis default	48	54	65	55	75	65
100						
Routing Engine 0	55	65	85	85	100	100
102						
FPC 0	63	68	75	70	90	83
95						
FPC 1	56	62	75	63	83	76
95						
FPC 7	56	62	75	63	83	76
95						
SIB 0	64	70	76	72	87	84
95						
SIB 0 ASIC Junction	63	68	75	70	105	100
107						
SIB 2	64	70	76	72	87	84
95						
SIB 2 ASIC Junction	63	68	75	70	105	100
107						
SIB 3	64	70	76	72	87	84
95						
SIB 3 ASIC Junction	63	68	75	70	105	100
107						

show chassis temperature-thresholds (QFX3500 Switch and QFX3600)

```
user@switch> show chassis temperature-thresholds
```

Item	Fan speed		Yellow alarm		Red alarm	
	(degrees C)		(degrees C)		(degrees C)	
Normal	High	Normal	Bad fan	Normal	Bad fan	
Normal						
FPC Sensor TopLeft I	48	56	53	43	56	46
FPC Sensor TopRight I	46	54	51	41	54	44
FPC Sensor TopLeft E	58	65	62	52	65	55
FPC Sensor TopRight E	56	64	61	51	64	54
FPC Sensor TopMiddle I	58	64	61	51	64	54
FPC Sensor TopMiddle E	67	74	71	61	74	64
FPC Sensor Bottom I	59	67	64	54	67	57
FPC Sensor Bottom E	66	73	70	60	73	63
FPC Sensor Die Temp	69	75	72	62	75	65
FPC Sensor Mgmt Brd I	46	54	51	41	54	44
FPC Sensor Switch I	56	63	60	50	63	53

show chassis temperature-thresholds interconnect-device (QFabric System)

```
user@switch> show chassis temperature-thresholds interconnect-device interconnect1
temperature-thresholds interconnect-device interconnect1
```

Item	Fan speed		Yellow alarm		Red alarm	
	(degrees C)		(degrees C)		(degrees C)	
Normal	High	Normal	Bad fan	Normal	Bad fan	
Chassis default	48	54	65	55	75	65

show chassis temperature-thresholds (PTX5000 Packet Transport Switch)

```
user@switch> show chassis temperature-thresholds
user@switch> show chassis temperature-thresholds
```

Item	Fan speed		Yellow alarm		Red alarm		Fire Shutdown
	(degrees C)		(degrees C)		(degrees C)		(degrees C)
Normal	High	Normal	Bad fan	Normal	Bad fan	Normal	
Routing Engine 0	70	75	90	87	102	97	115
Routing Engine 1	70	75	90	87	102	97	115
CB 0 Exhaust A	60	65	78	75	85	80	95

CB 0 Exhaust B	60	65	78	75	85	80	95
CB 1 Exhaust A	60	65	78	75	85	80	95
CB 1 Exhaust B	20	25	65	60	80	75	100
FPC 1 Exhaust A	60	65	78	75	85	80	95
FPC 1 Exhaust B	60	65	78	75	85	80	95
FPC 1 TL0	70	75	90	87	102	97	115
FPC 1 TQ0	70	75	90	87	102	97	115
FPC 1 TL1	70	75	90	87	102	97	115
FPC 1 TQ1	70	75	90	87	102	97	115
FPC 1 TL2	70	75	90	87	102	97	115
FPC 1 TQ2	70	75	90	87	102	97	115
FPC 1 TL3	70	75	90	87	102	97	115
FPC 1 TQ3	70	75	90	87	102	97	115
FPC 2 Exhaust A	60	65	78	75	85	80	95
FPC 2 Exhaust B	60	65	78	75	85	80	95
FPC 2 TL0	70	75	90	87	102	97	115
FPC 2 TQ0	70	75	90	87	102	97	115
FPC 2 TL1	70	75	90	87	102	97	115
FPC 2 TQ1	70	75	90	87	102	97	115
FPC 2 TL2	70	75	90	87	102	97	115
FPC 2 TQ2	70	75	90	87	102	97	115
FPC 2 TL3	70	75	90	87	102	97	115
FPC 2 TQ3	70	75	90	87	102	97	115
PIC 2/0 Ambient	60	65	78	75	85	80	95
PIC 2/0 cfp-2/0/1	60	65	70	67	75	72	85
PIC 2/1 Ambient	60	65	78	75	85	80	95
SIB 0 Exhaust	60	65	78	75	85	80	95
SIB 0 Junction	70	75	90	87	102	97	115
SIB 1 Exhaust	60	65	78	75	85	80	95
SIB 1 Junction	70	75	90	87	102	97	115
SIB 2 Exhaust	60	65	78	75	85	80	95
SIB 2 Junction	70	75	90	87	102	97	115
SIB 3 Exhaust	60	65	78	75	85	80	95
SIB 3 Junction	70	75	90	87	102	97	115
SIB 4 Exhaust	60	65	78	75	85	80	95
SIB 4 Junction	70	75	90	87	102	97	115
SIB 5 Exhaust	60	65	78	75	85	80	95
SIB 5 Junction	70	75	90	87	102	97	115
SIB 6 Exhaust	60	65	78	75	85	80	95
SIB 6 Junction	70	75	90	87	102	97	115
SIB 7 Exhaust	60	65	78	75	85	80	95
SIB 7 Junction	70	75	90	87	102	97	115
SIB 8 Exhaust	60	65	78	75	85	80	95
SIB 8 Junction	70	75	90	87	102	97	115

show chassis temperature-thresholds (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis temperature-thresholds
```

Fan speed	Yellow alarm		Red alarm		Fire Shutdown	
(degrees C)	(degrees C)		(degrees C)		(degrees C)	
Item	Normal	High	Normal	Bad fan	Normal	Bad fan
Normal						
Chassis default	48	54	65	55	75	65
100						
Routing Engine 0	70	80	95	95	110	110
112						
Routing Engine 1	70	80	95	95	110	110
112						
FPC 0	55	60	75	65	90	80

95						
FPC 1	55	60	75	65	90	80
95						
FPC 2	55	60	75	65	90	80
95						
FPC 4	55	60	75	65	90	80
95						
FPC 5	55	60	75	65	90	80
95						

show chassis zones

List of Syntax	Syntax on page 868 Syntax (QFX Series) on page 868
Syntax	show chassis zones <detail>
Syntax (QFX Series)	show chassis zones <detail> <interconnect-device <i>name</i> >
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	(QFabric systems only) Display the status of the two cooling system zones on the Interconnect device. Zone 1 consists of eight (0 – 7) front cards, which are cooled by two fan trays. Zone 2 consists of two control boards and eight rear cards, which are cooled by eight (0 – 7) fan trays. On MX2010 and MX2020 routers, display the status of the cooling system zones of the chassis. Zone 0 consists of the Control Board, ten (0–9) FPCs, and their respective PICs, Switch Fabric Boards, and Adapter Cards. Zone 1 consists of the Routing Engine, Control Board, and Switch Processor Mezzanine Boards.
Options	detail —(MX2010 and MX2020 routers only) (Optional) Display detailed status of the cooling system zones. detail <i>device-name</i> — (QFabric systems only) (Optional) Display detailed status of the two cooling systems on the Interconnect device. interconnect-device <i>name</i> — (QFabric systems only) (Optional) Display the status of the cooling zones on the Interconnect device.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• request chassis beacon on page 358• show chassis fan on page 618• show chassis temperature-thresholds on page 852
List of Sample Output	show chassis zones interconnect-device (QFabric System) on page 869 show chassis zones (MX2010 Router) on page 869 show chassis zones detail (MX2010 Router) on page 870 show chassis zones (MX2020 Router) on page 871 show chassis zones detail (MX2020 Router) on page 871 show chassis beacon interconnect-device (QFabric System) on page 872 show chassis beacon interconnect-device fpc (QFabric System) on page 873 show chassis beacon node-device (QFabric System) on page 873 show chassis beacon node-device fpc (QFabric System) on page 873

Output Fields Table 51 on page 464 lists the output fields for the **show chassis zones** command. Output fields are listed in the approximate order in which they appear.

Table 72: show chassis zones Output Fields

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.
Beacon State	Status of the beacon state: <ul style="list-style-type: none"> Off—The beacon is OFF. On—The beacon is ON.
show chassis zones command output fields for MX2020 and MX2010 routers:	
Driving FRU	Field replacable unit (FRU).
Temperature	Temperature of the specified FRU in degrees Celsius and degrees Fahrenheit.
Condition	Condition of the specified FRU. Condition can be HIGH TEMP , WARM TEMP , OK , and Offline .
Num Fans Missing	Number of fans or fan trays missing.
Num Fans Failed	Number of fans or fan trays that have failed.
Fan Duty Cycle	Fan duty cycle value.
show chassis zones detail command output fields for MX2020 and MX2010 routers:	
Item	Chassis component: <ul style="list-style-type: none"> Information about the chassis, Routing Engines, Control Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).
Measurement	Fan tray speed utilization in percentage.
Status	Status of the specified item. Status can be OK , Absent , or Offline .

Sample Output

show chassis zones interconnect-device (QFabric System)

```
user@switch> show chassis zones interconnect-device interconnect1
Slot          Beacon State
FPC           0          OFF
```

show chassis zones (MX2010 Router)

```
user@host> show chassis zones
```

```

ZONE 0 Status
  Driving FRU          FPC 6
  Temperature          81 degrees C / 177 degrees F
  Condition            HIGH TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          SFB 0 Exhaust-Zone1
  Temperature          71 degrees C / 159 degrees F
  Condition            WARM TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

```

show chassis zones detail (MX2010 Router)

```

user@host > show chassis zones
ZONE 0 Status
Item              Status              Measurement
CB 0              WARM TEMP
CB 1              WARM TEMP
FPC 0             HIGH TEMP
FPC 1             HIGH TEMP
FPC 2             WARM TEMP
FPC 3             HIGH TEMP
FPC 4             HIGH TEMP
FPC 5             HIGH TEMP
FPC 6             HIGH TEMP
FPC 7             HIGH TEMP
FPC 8             HIGH TEMP
FPC 9             HIGH TEMP
ADC 0             WARM TEMP
ADC 1             WARM TEMP
ADC 2             WARM TEMP
ADC 3             WARM TEMP
ADC 4             WARM TEMP
ADC 5             WARM TEMP
ADC 6             WARM TEMP
ADC 7             WARM TEMP
ADC 8             WARM TEMP
ADC 9             WARM TEMP
SFB 0             WARM TEMP
SFB 1             WARM TEMP
SFB 2             WARM TEMP
SFB 3             Offline
SFB 4             HIGH TEMP
SFB 5             WARM TEMP
SFB 6             HIGH TEMP
SFB 7             WARM TEMP
Fan Tray 0       OK                  Spinning at 98% fan tray speed
Fan Tray 1       OK                  Spinning at 98% fan tray speed

ZONE 1 Status
Item              Status              Measurement
CB 0              WARM TEMP
CB 1              WARM TEMP
Routing Engine 0  OK
Routing Engine 1  OK
SFB 0             WARM TEMP

```

SFB 1	WARM TEMP	
SFB 2	WARM TEMP	
SFB 3	Offline	
SFB 4	HIGH TEMP	
SFB 5	WARM TEMP	
SFB 6	HIGH TEMP	
SFB 7	WARM TEMP	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 64% fan tray speed
Fan Tray 3	OK	Spinning at 64% fan tray speed

show chassis zones (MX2020 Router)

```
user@host> show chassis zones
ZONE 0 Status
  Driving FRU          FPC 0
  Temperature          31 degrees C / 87 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          FPC 19
  Temperature          32 degrees C / 89 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30
```

show chassis zones detail (MX2020 Router)

```
user@host> show chassis zones detail
ZONE 0 Status
Item                Status                Measurement
CB 0                OK
CB 1                OK
FPC 0               OK
FPC 1               OK
FPC 2               OK
FPC 3               OK
FPC 4               OK
FPC 5               OK
FPC 6               OK
FPC 7               OK
FPC 8               OK
FPC 9               OK
ADC 0               OK
ADC 1               OK
ADC 2               OK
ADC 3               OK
ADC 4               OK
ADC 5               OK
ADC 6               OK
ADC 7               OK
ADC 8               OK
ADC 9               OK
SFB 0               OK
SFB 1               OK
SFB 2               OK
```

SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
Fan Tray 0	OK	Spinning at 38% fan tray speed
Fan Tray 1	OK	Spinning at 37% fan tray speed

ZONE 1 Status

Item	Status	Measurement
CB 0	OK	
CB 1	OK	
Routing Engine 0	OK	
Routing Engine 1	OK	
FPC 10	OK	
FPC 11	OK	
FPC 12	OK	
FPC 13	OK	
FPC 14	OK	
FPC 15	OK	
FPC 16	OK	
FPC 17	OK	
FPC 18	OK	
FPC 19	OK	
ADC 10	OK	
ADC 11	OK	
ADC 12	OK	
ADC 13	OK	
ADC 14	OK	
ADC 15	OK	
ADC 16	OK	
ADC 17	OK	
ADC 18	OK	
ADC 19	OK	
SFB 0	OK	
SFB 1	OK	
SFB 2	OK	
SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 38% fan tray speed
Fan Tray 3	OK	Spinning at 38% fan tray speed

show chassis beacon interconnect-device (QFabric System)

```

user@switch> show chassis beacon interconnect-device interconnect1
Chassis          OFF
CB 0             OFF
CB 1             OFF
FC 0 FPC 0       OFF
FC 1 FPC 1       OFF
RC 0 FPC 8       OFF
RC 1 FPC 9       OFF

```

show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                                ON
```

show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device node1
node1                                ON
```

show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                                ON
```

show cli

List of Syntax	Syntax on page 874 Syntax (QFX Series) on page 874
Syntax	show cli
Syntax (QFX Series)	show cli <authorization> <directory> <history <i>count</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display configured CLI settings.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli on page 875
Output Fields	Table 73 on page 874 lists the output fields for the show cli command. Output fields are listed in the approximate order in which they appear.

Table 73: show cli Output Fields

Field Name	Field Description
CLI complete-on-space	Capability to complete a partial command entry when you type a space or a tab: on or off .
CLI idle-timeout	Maximum time that an individual session can be idle before the user is logged out from the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is disabled .
CLI restart-on-upgrade	CLI is set to prompt you to restart the router or switch after upgrading the software: on or off .
CLI screen-length	Number of lines of text that the terminal screen displays.
CLI screen-width	Number of characters in a line on the terminal screen.
CLI terminal	Terminal type.
CLI is operating in	Mode: enhanced .
CLI timestamp	Date and time format for the timestamp. If the timestamp is not set, the state is disabled .
CLI working directory	Pathname of the working directory.

Sample Output

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI timestamp disabled
CLI working directory is '/var/home/regress'
```

show cli authorization

Syntax	show cli authorization
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the permissions for the current user.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli authorization on page 878
Output Fields	Table 74 on page 876 lists the output fields for the show cli authorization command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

Table 74: show cli authorization Output Fields

Field Name	Field Description
access	Can view access configuration information.
access-control	Can modify access configuration.
admin	Can view user account information.
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.

Table 74: show cli authorization Output Fields (*continued*)

Field Name	Field Description
flow-tap-control	Can configure flow-tap configuration information.
idp-profiler-operation	Can configure Profiler data.
interface	Can view interface configuration information.
interface-control	Can modify interface configuration information.
maintenance	Can perform system maintenance.
network	Can access the network by entering the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view Packet Gateway Control Protocol session mirroring configuration.
pgcp-session-mirroring-control	Can modify Packet Gateway Control Protocol session mirroring configuration all-control.
reset	Can reset or restart interfaces and system processes.
rollback	Can roll back to previous configurations.
routing	Can view routing configuration information.
routing-control	Can modify routing configuration information.
secret	Can view passwords and authentication keys in the configuration.
secret-control	Can modify passwords and authentication keys in the configuration.
security	Can view security configuration information.
security-control	Can modify security configuration information.
shell	Can start a local shell.
snmp	Can view SNMP configuration information.
snmp-control	Can modify SNMP configuration information.
system	Can view system configuration information.
system-control	Can modify system configuration information.
trace	Can view trace file settings information.

Table 74: show cli authorization Output Fields (*continued*)

Field Name	Field Description
trace-control	Can modify trace file settings information.
view	Can view current values and statistics.
view-configuration	Can view all configuration information (not including secrets).

Sample Output

show cli authorization

```

user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
  admin      -- Can view user accounts
  admin-control-- Can modify user accounts
  clear      -- Can clear learned network information
  configure  -- Can enter configuration mode
  control    -- Can modify any configuration
  edit       -- Can edit full files
  field      -- Special for field (debug) support
  floppy     -- Can read and write from the floppy
  interface  -- Can view interface configuration
  interface-control-- Can modify interface configuration
  network    -- Can access the network
  reset      -- Can reset/restart interfaces and daemons
  routing    -- Can view routing configuration
  routing-control-- Can modify routing configuration
  shell      -- Can start a local shell
  snmp       -- Can view SNMP configuration
  snmp-control-- Can modify SNMP configuration
  system     -- Can view system configuration
  system-control-- Can modify system configuration
  trace      -- Can view trace file settings
  trace-control-- Can modify trace file settings
  view       -- Can view current values and statistics
  maintenance -- Can become the super-user
  firewall   -- Can view firewall configuration
  firewall-control-- Can modify firewall configuration
  secret     -- Can view secret configuration
  secret-control-- Can modify secret configuration
  rollback   -- Can rollback to previous configurations
  security   -- Can view security configuration
  security-control-- Can modify security configuration
  access     -- Can view access configuration
  access-control-- Can modify access configuration
  view-configuration-- Can view all configuration (not including secrets)
  flow-tap   -- Can view flow-tap configuration
  flow-tap-control-- Can configure flow-tap service
Individual command authorization:
  Allow regular expression: none
  Deny regular expression: none
  Allow configuration regular expression: none
  Deny configuration regular expression: none

```


show cli directory

Syntax	show cli directory
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the current working directory.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli directory on page 880
Output Fields	Table 75 on page 880 lists the output fields for the show cli directory command. Output fields are listed in the approximate order in which they appear.

Table 75: show cli directory Output Fields

Field Name	Field Description
Current directory	Pathname of the current working directory.

Sample Output

show cli directory

```
user@host> show cli directory
Current directory: /var/home/regress
```

show cli history

Syntax	<code>show cli history</code> <code><count></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of previous CLI commands.
Options	none —Display all previous CLI commands. count —(Optional) Maximum number of commands to display.
Required Privilege Level	view
List of Sample Output	show cli history on page 881
Output Fields	Table 76 on page 881 lists the output fields for the show cli history command. Output fields are listed in the approximate order in which they appear.

Table 76: show cli history Output Fields

Field Name	Field Description
<i>timestamp</i>	Time at which the command was entered.
<i>command-syntax</i>	Command that was entered.

Sample Output

show cli history

```
user@host> show cli history
11:14:14 -- show arp
11:22:10 -- show cli authorization
11:27:12 -- show cli history
```

show host

Syntax	<code>show host <i>hostname</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Domain Name System (DNS) hostname information.
Options	<i>hostname</i> —Hostname or address.
Additional Information	The show host command displays the raw data received from the DNS server.
Required Privilege Level	view
List of Sample Output	show host on page 882

Sample Output

show host

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254

user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

show interfaces diagnostics optics

Syntax	<code>show interfaces diagnostics optics <i>interface-name</i></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Display diagnostics data and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, and QSFP+ optical transceivers installed in a QFX Series product. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p>
Options	<i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 317 • Installing a Transceiver in a QFX Series Device • Removing a Transceiver from a QFX Series Device • Junos® OS Network Interfaces
List of Sample Output	show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver) on page 887 show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver) on page 888
Output Fields	lists the output fields for the <code>show interfaces diagnostics optics</code> command. Output fields are listed in the approximate order in which they appear.

Table 77: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in volts.
(Not available for XFP transceivers)	

Table 77: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power (Not available for SFP and SFP+ transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Receiver signal average optical power (Not available for XFP transceivers)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is On or Off .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is On or Off .
Laser bias current high warning	Displays whether the laser bias power setting high warning is On or Off .
Laser bias current low warning	Displays whether the laser bias power setting low warning is On or Off .
Laser output power high alarm	Displays whether the laser output power high alarm is On or Off .
Laser output power low alarm	Displays whether the laser output power low alarm is On or Off .
Laser output power high warning	Displays whether the laser output power high warning is On or Off .
Laser output power low warning	Displays whether the laser output power low warning is On or Off .
Module temperature high alarm	Displays whether the module temperature high alarm is On or Off .
Module temperature low alarm	Displays whether the module temperature low alarm is On or Off .
Module temperature high warning	Displays whether the module temperature high warning is On or Off .
Module temperature low warning	Displays whether the module temperature low warning is On or Off .
Module voltage high alarm (Not available for XFP transceivers)	Displays whether the module voltage high alarm is On or Off .
Module voltage low alarm (Not available for XFP transceivers)	Displays whether the module voltage low alarm is On or Off .
Module voltage high warning (Not available for XFP transceivers)	Displays whether the module voltage high warning is On or Off .
Module voltage low warning (Not available for XFP transceivers)	Displays whether the module voltage low warning is On or Off .
Laser rx power high alarm	Displays whether the receive laser power high alarm is On or Off .

Table 77: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low alarm	Displays whether the receive laser power low alarm is On or Off .
Laser rx power high warning	Displays whether the receive laser power high warning is On or Off .
Laser rx power low warning	Displays whether the receive laser power low warning is On or Off .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Module not ready alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module not ready alarm is On or Off . When the output is On , the module has an operational fault.
Module power down alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module power down alarm is On or Off . When the output is On , the module is in a limited power mode, low for normal operation.
Tx data not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is On or Off .
Tx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is On or Off .
Tx laser fault alarm (Not available for SFP and SFP+ transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is On or Off .
Tx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is On or Off .
Rx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is On or Off .
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When on , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is On or Off .
Rx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is On or Off .
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.

Table 77: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser Rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser Rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser Rx power high warning.

Table 77: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser Rx power low warning.

Sample Output

show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver)

```

user@host> show interfaces diagnostics optics xe-0/0/1
Physical interface: xe-0/0/1
  Laser bias current           : 4.968 mA
  Laser output power          : 0.4940 mW / -3.06 dBm
  Module temperature          : 27 degrees C / 81 degrees F
  Module voltage              : 3.2310 V
  Receiver signal average optical power : 0.0000
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm     : Off
  Module voltage low alarm      : Off
  Module voltage high warning   : Off
  Module voltage low warning    : Off
  Laser rx power high alarm     : Off
  Laser rx power low alarm      : On
  Laser rx power high warning   : Off
  Laser rx power low warning    : On
  Laser bias current high alarm threshold : 10.500 mA
  Laser bias current low alarm threshold  : 2.000 mA
  Laser bias current high warning threshold : 9.000 mA
  Laser bias current low warning threshold : 2.500 mA
  Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
  Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
  Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
  Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold     : 3.630 V
  Module voltage low alarm threshold      : 2.970 V
  Module voltage high warning threshold   : 3.465 V
  Module voltage low warning threshold    : 3.135 V
  Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
  Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
  Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
  Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver)

```

user@host> show interfaces diagnostics optics node1:xe-0/0/1
Physical interface: node1:xe-0/0/1
  Laser bias current           : 4.968 mA
  Laser output power          : 0.4940 mW / -3.06 dBm
  Module temperature          : 27 degrees C / 81 degrees F
  Module voltage              : 3.2310 V
  Receiver signal average optical power : 0.0000
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm     : Off
  Module voltage low alarm      : Off
  Module voltage high warning   : Off
  Module voltage low warning    : Off
  Laser rx power high alarm     : Off
  Laser rx power low alarm      : On
  Laser rx power high warning   : Off
  Laser rx power low warning    : On
  Laser bias current high alarm threshold : 10.500 mA
  Laser bias current low alarm threshold : 2.000 mA
  Laser bias current high warning threshold : 9.000 mA
  Laser bias current low warning threshold : 2.500 mA
  Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
  Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
  Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
  Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold : 3.630 V
  Module voltage low alarm threshold : 2.970 V
  Module voltage high warning threshold : 3.465 V
  Module voltage low warning threshold : 3.135 V
  Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
  Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
  Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
  Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

show log

List of Syntax	Syntax on page 889 Syntax (QFabric System) on page 889 Syntax (TX Matrix Routers) on page 889
Syntax	<pre>show log <filename user <username>></pre>
Syntax (QFabric System)	<pre>show log filename <device-type (device-id device-alias)></pre>
Syntax (TX Matrix Routers)	<pre>show log <all-lcc lcc number scc> <filename user <username>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p>
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.
Options	<p>none—List all log files.</p> <p><all-lcc lcc number scc>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p>device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p> <ul style="list-style-type: none"> • director-device—Display logs for Director devices. • infrastructure-device—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup). • interconnect-device—Display logs for Interconnect devices. • node-device—Display logs for Node devices.



NOTE: If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

Required Privilege Level trace

List of Sample Output [show log on page 890](#)
[show log filename on page 890](#)
[show log filename \(QFabric System\) on page 891](#)
[show log user on page 891](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

show log user

```

user@host> show log user
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

show ntp associations

Syntax	<code>show ntp associations</code> <code><no-resolve></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Network Time Protocol (NTP) peers and their state.
Options	none —Display NTP peers and their state. no-resolve —(Optional) Suppress symbolic addressing.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show ntp status on page 894
List of Sample Output	show ntp associations on page 893
Output Fields	Table 78 on page 892 describes the output fields for the show ntp associations command. Output fields are listed in the approximate order in which they appear.

Table 78: show ntp associations Output Fields

Field Name	Field Description
remote	Address or name of the remote NTP peer.
refid	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0 .
st	Stratum of the remote peer.
t	Type of peer: b (broadcast), l (local), m (multicast), or u (unicast).
when	When the last packet from the peer was received.
poll	Polling interval, in seconds.
reach	Reachability register, in octal.
delay	Current estimated delay of the peer, in milliseconds.
offset	Current estimated offset of the peer, in milliseconds.
disp	Current estimated dispersion of the peer, in milliseconds.

Table 78: show ntp associations Output Fields (*continued*)

Field Name	Field Description
<i>peer-name</i>	<p>Peer name and status of the peer in the clock selection process:</p> <ul style="list-style-type: none"> • space—Discarded because of a high stratum value or failed sanity checks. • x—Designated "falseticker" by the intersection algorithm. • .—Culled from the end of the candidate list. • — —Discarded by the clustering algorithm. • +—Included in the final selection set. • #—Selected for synchronization, but the distance exceeds the maximum. • *—Selected for synchronization. • o—Selected for synchronization, but the packets-per-second (pps) signal is in use.

Sample Output

show ntp associations

```

user@host> show ntp associations
      remote           refid      st t when poll reach  delay  offset   disp
=====
*wolfe-gw.junipe tick.ucla.edu    2 u  43   64  377    1.86   0.319   0.08

```

show ntp status

Syntax	<code>show ntp status</code> <code><no-resolve></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the values of internal variables returned by Network Time Protocol (NTP) peers.
Options	none —Display the values of internal variables returned by NTP peers. no-resolve —(Optional) Suppress symbolic addressing.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ntp associations on page 892
List of Sample Output	show ntp status on page 895
Output Fields	Table 79 on page 894 describes the output fields for the show ntp status command. Output fields are listed in the approximate order in which they appear.

Table 79: show ntp status Output Fields

Field Name	Field Description
status	System status word, a code representing the status items listed.
leap_none	Indicates a normal synchronized state with no leap seconds imminent. Other options could be leap_add_sec , leap_del_sec , or leap_alarm , indicating a leap second will be added, deleted, or a leap second requirement is upcoming.
sync_ntp	Indicates the current synchronization source, in this case, an NTP server. Other options include sync_alarm and sync_unspec , both indicating that the router has not been synched.
x events	Indicates the number of events that have occurred since that last code change. An event is often the receipt of an NTP polling message.
event_peer/strat_chg	Describes the most recent event, in this case, the stratum of the peer server changed.
version	A detailed description of the version of NTP being used.
processor	Indicates the current hardware platform and version of the processor.
system	Detailed description of the name and version of the operating system in use.
leap	The number of leap seconds in use.

Table 79: show ntp status Output Fields (*continued*)

Field Name	Field Description
stratum	The stratum of the peer server. Anything greater than 1 is a secondary reference source, and the number roughly represents the number of hops away from the stratum 1 server.. Stratum 1 is a primary reference, such as an atomic clock.
precision	The precision of the peer clock, how precisely the frequency and time can be maintained with this particular timekeeping system.
rootdelay	The total roundtrip delay to the primary reference source, in seconds.
rootdispersion	The maximum error relative to the primary reference source, in seconds.
peer	An identification number of the peer in use.
refid	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0.
reftime	The local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
poll	The NTP broadcast message polling interval, in seconds.
clock	The current time on the local router clock.
state	The current mode of NTP operation, where 1 is symmetric active, 2 is symmetric passive, 3 is client, 4 is server, and 5 is broadcast.
offset	Current estimated offset of the peer, in milliseconds. Indicates the time difference between the reference clock and the local clock.
frequency	The frequency of the clock.
jitter	Indicates the magnitude of jitter, in milliseconds, between several time queries.
stability	A measure of how well this clock can maintain a constant frequency.

Sample Output

show ntp status

```

user@host> show ntp status
assID=0 status=0544 leap_none, sync_local_proto, 4 events, event_peer/strat_chg,
version="ntpd 4.2.2p1@1.1570-o Tue May 19 13:57:55 UTC 2009 (1)",
processor="x86_64", system="Linux/2.6.18-164.el5", leap=00, stratum=4,
precision=-10, rootdelay=0.000, rootdispersion=11.974, peer=59475,
refid=LOCAL(0),
reftime=d495c32c.0e71eaf2 Mon, Jan 7 2013 13:57:00.056, poll=10,
clock=d495c32c.cebd43bd Mon, Jan 7 2013 13:57:00.807, state=4,
offset=0.000, frequency=0.000, jitter=0.977, noise=0.977,
stability=0.000, tai=0

```


show subscribers

Syntax show subscribers
 <detail | extensive | terse>
 <aci-interface-set-name *aci-interface-set-name*>
 <address *address*>
 <agent-circuit-identifier *agent-circuit-identifier-substring*>
 <client-type *client-type*>
 <count>
 <interface *interface*>
 <logical-system *logical-system*>
 <mac-address *mac-address*>
 <physical-interface *physical-interface-name*>
 <profile-name *profile-name*>
 <routing-instance *routing-instance*>
 <stacked-vlan-id *stacked-vlan-id*>
 <subscriber-state *subscriber-state*>
 <user-name *user-name*>
 <vlan-id *vlan-id*>

Release Information Command introduced in Junos OS Release 9.3.
 Command introduced in Junos OS Release 9.3 for EX Series switches.
client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
 The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.

Description Display information for active subscribers.

Options **detail | extensive | terse**—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.168.17.1). If you specify the IP address as a prefix with a netmask (for example, 192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.



NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Related Documentation

- *show subscribers summary*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

[show subscribers \(IPv4\) on page 902](#)
[show subscribers \(IPv6\) on page 902](#)

[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 903](#)
[show subscribers \(LNS on MX Series Routers\) on page 903](#)
[show subscribers client-type dhcp detail on page 903](#)
[show subscribers count on page 903](#)
[show subscribers address detail \(IPv6\) on page 903](#)
[show subscribers detail \(IPv4\) on page 904](#)
[show subscribers detail \(IPv6\) on page 904](#)
[show subscribers detail \(IPv6 Static Demux Interface\) on page 905](#)
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 905](#)
[show subscribers detail \(Tunneled Subscriber\) on page 905](#)
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 905](#)
[show subscribers detail \(ACI Interface Set Session\) on page 906](#)
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 906](#)
[show subscribers extensive on page 907](#)
[show subscribers extensive \(RPF Check Fail Filter\) on page 907](#)
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 907](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 908](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack over PPP with Shared Dynamic Profile\) on page 909](#)
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 910](#)
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 910](#)
[show subscribers interface extensive on page 911](#)
[show subscribers logical-system terse on page 912](#)
[show subscribers physical-interface count on page 912](#)
[show subscribers routing-instance inst1 count on page 912](#)
[show subscribers stacked-vlan-id detail on page 912](#)
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 912](#)
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 912](#)
[show subscribers user-name detail on page 912](#)
[show subscribers vlan-id on page 913](#)
[show subscribers vlan-id detail on page 913](#)

Output Fields [Table 80 on page 899](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 80: show subscribers Output Fields

Field Name	Field Description
User Name	Name of subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.

Table 80: show subscribers Output Fields (*continued*)

Field Name	Field Description
Primary DNS Address	IP address of primary DNS server.
Secondary DNS Address	IP address of secondary DNS server.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
Interface Type	Whether the subscriber interface is Static or Dynamic .
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.
Interface Set Type	Interface type of the ACI interface set: Dynamic . This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.

Table 80: show subscribers Output Fields (*continued*)

Field Name	Field Description
State	Current state of the subscriber session (Init, Configured, Active, Terminating, Tunneled).
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
Agent Remote ID	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
Login Time	Date and time at which the subscriber logged in.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.

Table 80: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
Family	Family that is active for the subscriber session—innet, inet6 or both inet and inet6.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).

Sample Output

show subscribers (IPv4)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/3/0.1073741824  100                WHOLESALER-CLIENT default:default
demux0.1073741824    100.0.0.10         RETAILER1-CLIENT test1:retailer1
demux0.1073741825    101.0.0.3          RETAILER1-CLIENT test1:retailer1
demux0.1073741826    102.0.0.3          RETAILER2-CLIENT test1:retailer2

```

show subscribers (IPv6)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/0/0.0      2001::c0:0:0:0/74  WHOLESALER-CLIENT default:default
*               2002::1/128        subscriber-25   default:default

```

show subscribers (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers
Interface          IP Address/VLAN ID      User Name
LS:RI
demux0.1073741834  0x8100.1002 0x8100.1
default:default
demux0.1073741835  0x8100.1001 0x8100.1
default:default
pp0.1073741836     61.1.1.1              dualstackuser1@ISP1.com
default:ASP-1
*                  2041:1:1::/48
*                  2061:1:1:1::/64
pp0.1073741837     23.1.1.3              dualstackuser2@ISP1.com
default:ASP-1
*                  2001:1:2:5::/64

```

show subscribers (LNS on MX Series Routers)

```

user@host> show subscribers
Interface          IP Address/VLAN ID      User Name      LS:RI
si-4/0/0.1         192.168.4.1            xyz@example.com default:default

```

show subscribers client-type dhcp detail

```

user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT

```

show subscribers count

```

user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188

```

show subscribers address detail (IPv6)

```

user@host> show subscribers address 100.16.12.137 detail

```

```
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 100.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Session Timeout (seconds): 3600
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
```

```

Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (IPv6 Static Demux Interface)

```

user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT

```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```

user@host> show subscribers detail
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST

```

show subscribers detail (Tunneled Subscriber)

```

user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512

```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active

```

Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
```

```

IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST

```

show subscribers extensive

```

user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48

```

show subscribers extensive (RPF Check Fail Filter)

```

user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...

```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```

user@host> show subscribers extensive

```

```
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
```



```

MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

show subscribers extensive (IPv4 and IPv6 Dual Stack over PPP with Shared Dynamic Profile)

```

user@host> show subscribers extensive
Type: PPPoE
User Name: DEFAULTUSER
IP Address: 100.16.0.1
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:6::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073741832
Interface type: Dynamic
Underlying Interface: demux0.2
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:02:02:00:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: 27480
Session ID: 27480
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2013-02-12 13:56:03 PST
Service Sessions: 1
IP Address Pool: v4-pool-0
IPv6 Address Pool: v6RAPrefixPool1
IPv6 Interface Address: 1016:0:0:6::1/64
IPv6 Framed Interface Id: 202:2ff:fe00:1

Service Session ID: 27481
Service Session Name: l3-v4v6-mixed
State: Active
Family: inet, inet6
IPv4 Input Filter Name: v4downstrm-filter-pp0.1073741832-in
IPv4 Output Filter Name: v4upstrm-filter-pp0.1073741832-out
IPv6 Input Filter Name: v6-dn-filter-pp0.1073741832-in
IPv6 Output Filter Name: v6-up-filter-pp0.1073741832-out

Type: DHCP
IPv6 Prefix: 2016::/40
Logical System: default
Routing Instance: default
Interface: pp0.1073741832
Interface type: Static
Underlying Interface: demux0.2
MAC Address: 00:02:02:00:00:01
State: Active
Radius Accounting ID: 27482

```

```
Session ID: 27482
Underlying Session ID: 27480
Login Time: 2013-02-12 13:56:06 PST
DHCP Options: len 42
00 08 00 02 00 c8 00 01 00 0a 00 03 00 01 00 02 02 00 00 01
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Address Pool: v6DelegatedPdPool1
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
```

```

User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers interface extensive

```

user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out

```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface          IP Address/VLAN ID  User Name          LS:RI
demux0.1073741825  101.0.0.3           RETAILER1-CLIENT  test1:retailer1
demux0.1073741826  102.0.0.3           RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
```

```

Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c

```

show subscribers vlan-id

```

user@host> show subscribers vlan-id 100
Interface          IP Address          User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825

```

show subscribers vlan-id detail

```

user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

```

show system alarms

Syntax	show system alarms
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display active system alarms.
Options	This command has no options.
Additional Information	<p>System alarms are preset. They include a configuration alarm that appears when no rescue configuration alarm is set and a license alarm that appears when a software feature is configured and no valid license is configured for the feature. For more information about system alarms, see the <i>Junos OS System Basics Configuration Guide</i>.</p> <p>In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.</p>
Required Privilege Level	admin
List of Sample Output	show system alarms on page 914 show system alarms (Fan Tray) on page 914 show system alarms (QFX Series) on page 914

Sample Output

show system alarms

```
user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent
```

show system alarms (Fan Tray)

```
user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure
```

show system alarms (QFX Series)

```
user@switches> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent
```

show system audit

List of Syntax	Syntax on page 915 Syntax (EX Series Switch and MX Series Router) on page 915 Syntax (TX Matrix Router) on page 915 Syntax (TX Matrix Plus Router) on page 915 Syntax (QFX Series) on page 915
Syntax	show system audit <root-only>
Syntax (EX Series Switch and MX Series Router)	show system audit <all-members> <local> <member <i>member-id</i> > <root-only>
Syntax (TX Matrix Router)	show system audit <all-lcc lcc <i>number</i> scc> <root-only>
Syntax (TX Matrix Plus Router)	show system audit <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <root-only>
Syntax (QFX Series)	show system audit <infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i> root-only>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display the state and checksum values for file systems.
Options	<p>none—Display the state and checksum values for all file systems.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for all of the chassis.</p> <p>all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for all T1600 or T4000 routers connected to the TX Matrix Plus router.</p> <p>all-members—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for a specific T640 router</p>

that is connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

infrastructure *name*—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Interconnect device.

local—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

member *member-id*—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Node group

root-only—(Optional) Check only the root (/) file system. On a QFabric system, you can check the root (/) file system on the infrastructure (fabric manager Routing Engine and fabric control Routing Engine), Interconnect device, or Node group.

scc—(TX Matrix routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information To redirect the output to a file, issue the following command:

ssh device-name 'show system audit root-only' > output-file

If you save the output of the **show system audit root-only** command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

By default, when you issue the **show system audit** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level admin

List of Sample Output [show system audit root-only on page 917](#)
[show system audit lcc \(TX Matrix Router\) on page 918](#)
[show system audit lcc \(TX Matrix Plus Router\) on page 919](#)
[show system audit root-only \(QFX3500 Switch\) on page 921](#)

Sample Output

show system audit root-only

```
user@host> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
           md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
           mode=0744 size=1934552 time=944688902.0 \
           md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
           md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
           md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \
           md5digest=93d722493ed38477338a1405d7dcbb40
boot.help  uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
           md5digest=9b7126385734bcae753f4179ab59d8e5
compat     type=link mode=0777 size=11 time=915149058.0 \
           link=/usr/compat
kernel     mode=0444 size=1947607 time=950230892.0 \
           md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr uid=1112 size=1947642 time=950252597.0 \
           md5digest=82e1637682d58ec28964dfce7fccb62e
kernel.config \
           mode=0644 size=0 time=915149058.0 \
           md5digest=d41d8cd98f00b204e9800998ecf8427e
sys        type=link mode=0777 size=11 time=915149029.0 \
           link=/usr/src/sys
```

show system audit lcc (TX Matrix Router)

```

user@host> show system audit lcc 2
lcc2-re0:
-----
#       user: root
#       machine: rodin-lcc2
#       tree: /
#       date: Mon Sep 13 11:55:33 2004

# .
/set type=file uid=0 gid=0 mode=0555 nlink=1 flags=none
.      type=dir nlink=20 size=512 time=1094982121.0
  COPYRIGHT mode=0644 size=4735 time=986012708.0 \
    md5digest=78396df1404ad742e6eb1be28f0cd63b
  kernel type=link mode=0700 size=17 time=1090266262.0 \
    link=/packages/jkernel

# ./altconfig
altconfig type=dir nlink=2 size=512 time=1089801320.0
# ./altconfig
..

# ./altroot
altroot type=dir nlink=2 size=512 time=1089801320.0
# ./altroot
..

# ./b
b type=dir mode=0755 nlink=2 size=512 time=1093961429.0
# ./b
..

# ./bin
/set type=file uid=0 gid=0 mode=0700 nlink=1 flags=none
bin type=dir mode=0755 nlink=2 size=512 time=1089843059.0
  [ type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cat
  chmod type=link size=29 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1090266270.0 \

```

```

kill      link=/packages/mnt/jbase/bin/hostname
          type=link size=28 time=1090266270.0 \
ln        link=/packages/mnt/jbase/bin/kill
          type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ln
ls        link=/packages/mnt/jbase/bin/ls
          type=link size=26 time=1090266270.0 \
mkdir    link=/packages/mnt/jbase/bin/mkdir
          type=link size=29 time=1090266270.0 \
mv        link=/packages/mnt/jbase/bin/mv
          type=link size=26 time=1090266270.0 \
ps        link=/packages/mnt/jbase/bin/ps
          type=link size=26 time=1090266270.0 \
pwd       link=/packages/mnt/jbase/bin/pwd
          type=link size=27 time=1090266270.0 \
rcp       link=/packages/mnt/jbase/bin/rcp
          type=link size=27 time=1090266270.0 \
red       link=/packages/mnt/jbase/bin/red
          type=link size=26 time=1090266270.0 \
rm        link=/packages/mnt/jbase/bin/rm
          type=link size=26 time=1090266270.0 \
rmdir    link=/packages/mnt/jbase/bin/rmdir
          type=link size=29 time=1090266270.0 \
sh        link=/packages/mnt/jbase/bin/sh
          type=link size=26 time=1090266270.0 \
sleep    link=/packages/mnt/jbase/bin/sleep
          type=link size=29 time=1090266270.0 \
stty     link=/packages/mnt/jbase/bin/stty
          type=link size=28 time=1090266270.0 \
sync     link=/packages/mnt/jbase/bin/sync
          type=link size=28 time=1090266270.0 \
tcsh     link=/packages/mnt/jbase/bin/csh
          type=link size=27 time=1090266270.0 \
test     link=/packages/mnt/jbase/bin/test
          type=link size=28 time=1090266270.0 \

# ./bin
..

# ./boot
/set type=file uid=0 gid=0 mode=0444 nlink=1 flags=none
boot      type=dir mode=0555 nlink=3 size=512 time=1095069935.0
          size=512 time=1094978286.0 \
          md5digest=6f780822dd4ae482a20462b66e542cca
boot0     mode=0555 size=512 time=1094978294.0 \
          md5digest=8d112b09df342cd0b60fdb9bdcde8e07
boot1     mode=0555 size=7680 time=1094978294.0 \
          md5digest=28eb58c4068c6b85717e1484f9e028e4
boot2     mode=0555 size=165888 time=1094978298.0 \
          md5digest=1474c6b800dfc82ba552d7c36116d07d
cdboot    size=5996 time=1094982121.0 \
          md5digest=c53dc948eb07e2ea4eb0413e4c4634a3
kgzldr.o  mode=0555 size=163840 time=1094978298.0 \
          md5digest=82d9dc2d31033476bfb61bb7264c4fed
loader    size=9237 time=986013631.0 \
          md5digest=43144391465ad50267d31e0a320be1de
loader.4th
...

```

show system audit lcc (TX Matrix Plus Router)

```
user@host> show system audit all-chassis
```

```

sfc0-re0:
-----
#          user: root
#          machine: finalfive
#          tree: /
#          date: Mon May 18 00:13:16 2009

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1 flags=none
.      type=dir nlink=23 size=512 time=1242347096.0
  COPYRIGHT mode=0644 size=6196 time=1168587741.0 \
    md5digest=bbad415e1c29bbdd9b383537100412c
    kernel type=link size=17 time=1242347011.0 link=/packages/jkernel
    staging type=link mode=0777 size=8 time=1242346935.0 link=/var/tmp

# ./snap
.snap type=dir mode=0775 nlink=2 size=512 time=1242346922.0
# ./snap
..

# ./altconfig
altconfig type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altconfig
..

# ./altroot
altroot type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altroot
..

# ./bin
bin type=dir nlink=2 size=512 time=1242346944.0
  \133 type=link size=28 time=1242346942.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cat
  chflags type=link size=31 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chflags
  chmod type=link size=29 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1242346941.0 \

```

```

kill      link=/packages/mnt/jbase/bin/hostname
          type=link size=28 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mv
pax       type=link size=27 time=1242346944.0 \
          link=/packages/mnt/jbase/bin/pax
ps        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/test
# ./bin
...

```

show system audit root-only (QFX3500 Switch)

```

user@switch> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
          md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
          mode=0744 size=1934552 time=944688902.0 \
          md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
          md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
          md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \

```

```
boot.help      md5digest=93d722493ed38477338a1405d7dcbb40
                uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
                md5digest=9b7126385734bcae753f4179ab59d8e5
compat         type=link mode=0777 size=11 time=915149058.0 \
                link=/usr/compat
kernel         mode=0444 size=1947607 time=950230892.0 \
                md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr     uid=1112 size=1947642 time=950252597.0 \
                md5digest=82e1637682d58ec28964dfee7fccb62e
kernel.config \
                mode=0644 size=0 time=915149058.0 \
                md5digest=d41d8cd98f00b204e9800998ecf8427e
sys            type=link mode=0777 size=11 time=915149029.0 \
                link=usr/src/sys
```

show system boot-messages

List of Syntax	Syntax on page 923 Syntax (EX Series Switches) on page 923 Syntax (TX Matrix Router) on page 923 Syntax (TX Matrix Plus Router) on page 923 Syntax (MX Series Router) on page 923 Syntax (QFX Series) on page 923
Syntax	show system boot-messages
Syntax (EX Series Switches)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system boot-messages infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .
Options	none —Display all boot time messages. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all connected T1600 or T4000 LCCs. all-members —(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines.

interconnect-device *name*—(QFabric systems only) (Optional) Display boot time messages on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for a specific router connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display boot time messages on the Node group.

scc—(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system boot-messages** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

- [show system boot-messages \(TX Matrix Router\) on page 925](#)
- [show system boot-messages lcc \(TX Matrix Router\) on page 926](#)
- [show system boot-messages \(TX Matrix Plus Router\) on page 927](#)
- [show system boot-messages \(QFX3500 Switch\) on page 927](#)

Sample Output

show system boot-messages (TX Matrix Router)

```

user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x66a Stepping=10
    Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on

pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300

```

```

ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

show system boot-messages lcc (TX Matrix Router)

```

user@host> show system boot-messages lcc 2
lcc2-re0:
-----
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@benten.juniper.net:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
Timecounter "TSC" frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x387f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR,SSE>
real memory = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
Pentium Pro MTRR support enabled
md0: Malloc disk

```

```

DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...

```

show system boot-messages (TX Matrix Plus Router)

```

user@host> show system boot-messages
sfc0-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/JUNIPER
MPTable: Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R)
CPU          L5238 @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin =
"GenuineIntel" Id = 0x1067a Stepping = 10   Features=0xbfebfbff
...
lcc1-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU          @ 1.86GHz (1862.01-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x1067a Stepping = 10
Features=0xbfebfbff
...

```

show system boot-messages (QFX3500 Switch)

```

user@switch> show sytem boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
-----
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
-----

Copyright (c) 1996-2010, Juniper Networks, Inc.

```

All rights reserved.

Copyright (c) 1992-2006 The FreeBSD Project.

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994

The Regents of the University of California. All rights reserved.

JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
```

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, bt1b_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
  L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
  L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard
Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource. tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
```

```

pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a

```

show system buffers

List of Syntax	Syntax on page 930 Syntax (EX Series) on page 930 Syntax (TX Matrix Router) on page 930 Syntax (TX Matrix Plus Router) on page 930 Syntax (MX Series Router) on page 930 Syntax (QFX Series) on page 930
Syntax	show system buffers
Syntax (EX Series)	show system buffers <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system buffers <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system buffers <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system buffers <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system buffers <infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i> root-only (infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).
Options	none —Show all buffer statistics. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for all routers connected to the TX Matrix Plus router. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show buffer statistics for all of the chassis.

all-members—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Show buffer statistics for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Show buffer statistics for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for a specific router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Show buffer statistics for the Node group

sfc—(TX Matrix Plus routers only) (Optional) Show buffer statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the **show system buffers** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	show system buffers on page 933 show system buffers scc (TX Matrix Router) on page 934 show system buffers sfc (TX Matrix Plus Router) on page 934 show system buffers all-chassis (TX Matrix Plus Router) on page 934 show system buffers node-group (QFabric System) on page 935
Output Fields	Table 81 on page 933 describes the output fields for the show system buffers command. Output fields are listed in the approximate order in which they appear.

Table 81: show system buffers Output Fields

Field Name	Field Description
mbufs in use	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
mbufs allocated to packet headers	Number of memory buffers currently holding packet headers
mbufs allocated to control blocks	Number of memory buffers currently holding the state for sockets.
mbufs allocated to send data	Number of memory buffers currently holding socket send data.
mbufs allocated to pfe refill data	Number of memory buffers currently holding Packet Forwarding Engine refill data.
mbufs allocated to fxp data	Number of memory buffers currently holding fxp data.
mbufs allocated to socket names and addresses	Number of memory buffers currently holding addresses for sockets.
mbuf clusters in use	Allocation statistics for memory buffer clusters.
allocated to network	Total amount of memory in use by the networking and interprocess communication (IPC) code.
requests for memory denied	Number of times a memory allocation request within the IPC and networking code failed.
requests for memory delayed	Number of times a memory allocation request within the IPC and networking code was postponed.
calls to protocol drain routines	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

Sample Output

show system buffers

```

user@host> show system buffers
397/893/1290 mbufs in use (current/cache/total)
395/331/726/30000 mbuf clusters in use (current/cache/total/max)
384/256 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
889K/885K/1774K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/5/1024 sfbufs in use (current/peak/max)

```

```
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

show system buffers scc (TX Matrix Router)

```
user@host> show system buffers scc
213 mbufs in use:
    11 mbufs allocated to packet headers
    26 mbufs allocated to socket names and addresses
    2 mbufs allocated to socket options
    17 mbufs allocated to socket send data
    2 mbufs allocated to pfe data
    155 mbufs allocated to fxp data (rx)
    511 mbufs allocated to <mbuf type 86>
    256 mbufs allocated to <mbuf type 92>
924/1162 mbuf clusters in use
2788 Kbytes allocated to network (75% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

show system buffers sfc (TX Matrix Plus Router)

```
user@host> show system buffers sfc 0

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

show system buffers all-chassis (TX Matrix Plus Router)

```
user@host> show system buffers all-chassis

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
```

```
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc0-re0:
```

```
-----
772/2558/3330 mbufs in use (current/cache/total)
772/598/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1737K/1835K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc1-re0:
```

```
-----
773/2437/3210 mbufs in use (current/cache/total)
773/453/1226/30000 mbuf clusters in use (current/cache/total/max)
768/384 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1739K/1515K/3254K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/7/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc2-re0:
```

```
-----
816/2514/3330 mbufs in use (current/cache/total)
816/554/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1836K/1736K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
```

show system buffers node-group (QFabric System)

```
user@switch> show system buffers node-group node1
node-group node1:
```

```
-----
2/2698/2700 mbufs in use (current/cache/total)
2/1520/1522/30000 mbuf clusters in use (current/cache/total/max)
0/1280 mbuf+clusters out of packet secondary zone in use (current/cache)
```

```
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
4K/3714K/3719K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/6/6656 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

re0:

```
-----
516/639/1155 mbufs in use (current/cache/total)
515/147/662/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1159K/453K/1612K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

re1:

```
-----
519/771/1290 mbufs in use (current/cache/total)
518/176/694/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1165K/544K/1710K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

show system certificate

Syntax	<code>show system certificate</code> <code><certificate-id></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	(Encryption interface on M Series, T Series routers, and QFX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority.
Options	none —Display all installed certificates signed by the Juniper Networks certificate authority. certificate-id —(Optional) Display the details of a particular certificate.
Required Privilege Level	maintenance
List of Sample Output	show system certificate on page 938 show system certificate (QFX Series) on page 938
Output Fields	Table 82 on page 937 lists the output fields for the show system certificate command. Output fields are listed in the approximate order in which they appear.

Table 82: show system certificate Output Fields

Field Name	Field Description
Certificate identifier	Unique identifier associated with a certificate. The certificate identifier is the common name of the subject.
Issuer Subject	Information about the certificate issuer and the distinguished name (DN) of the issuer, respectively: <ul style="list-style-type: none"> • Organization—Name of the owner's organization. • Organizational unit—Name of the owner's department. • Country—Two-character country code in which the owner's system is located. • State—State in the USA in which the owner is using the certificate. • Locality—City in which the owner's system is located. • Common name—Name of the owner of the certificate. • E-mail address—E-mail address of the owner of the certificate.
Validity	When a certificate is valid.
Signature algorithm	Encryption algorithm applied to the installed certificate.
Public key algorithm	Encryption algorithm applied to the public key.

Sample Output

show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system commit

Syntax	show system commit
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the system commit history and any pending commit operation.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear system commit on page 331
List of Sample Output	show system commit on page 940 show system commit (At a Particular Time) on page 940 show system commit (At the Next Reboot) on page 940 show system commit (Rollback Pending) on page 940 show system commit (QFX Series) on page 940
Output Fields	Table 83 on page 939 describes the output fields for the show system commit command. Output fields are listed in the approximate order in which they appear.

Table 83: show system commit Output Fields

Field Name	Field Description
Junos XML protocol	Displays the last 50 commit operations listed, most recent to first. The identifier Junos XML protocol designates a configuration created for recovery using the request system configuration rescue save command.
Junos XML protocol	Date and time of the commit operation.
Junos XML protocol	User who executed the commit operation.
Junos XML protocol	Method used to execute the commit operation: <ul style="list-style-type: none"> • Junos XML protocol—CLI interactive user performed the commit operation. • Junos XML protocol—Junos XML protocol client performed the commit operation. • synchronize—The commit synchronize command was performed on the other Routing Engine. • snmp—An SNMP set request caused the commit operation. • button—A button on the router or switch was pressed to commit a rescue configuration for recovery. • autoinstall—A configuration obtained through autoinstallation was committed. • other—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.

Sample Output

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

show system configuration archival

Syntax	show system configuration archival
Release Information	Introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display directory and number of files queued for archival transfer.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options	This command has no options.
Required Privilege Level	maintenance
List of Sample Output	show system configuration archival on page 941

Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

show system configuration rescue

Syntax	show system configuration rescue
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a rescue configuration, if one exists.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• show system configuration archival on page 941
List of Sample Output	show system configuration rescue on page 942

Sample Output

show system configuration rescue

```
user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
  services {
```

```
        ftp;  
        rlogin;  
        rsh;  
        telnet;  
    }  
}  
.....
```

show system connections

List of Syntax	Syntax on page 944 Syntax (EX Series) on page 944 Syntax (TX Matrix Router) on page 944 Syntax (TX Matrix Plus Router) on page 944 Syntax (MX Series Router) on page 944 Syntax (QFX Series) on page 944
Syntax	<code>show system connections</code> <code><extensive></code> <code><all-chassis all-lcc lcc <i>number</i> scc></code> <code><inet inet6></code> <code><show-routing-instances></code>
Syntax (EX Series)	<code>show system connections</code> <code><extensive></code> <code><all-members></code> <code><inet inet6></code> <code><local></code> <code><member <i>member-id</i>></code> <code><show-routing-instances></code>
Syntax (TX Matrix Router)	<code>show system connections</code> <code><extensive></code> <code><all-chassis all-lcc lcc <i>number</i> scc></code> <code><inet inet6></code> <code><show-routing-instances></code>
Syntax (TX Matrix Plus Router)	<code>show system connections</code> <code><extensive></code> <code><all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></code> <code><inet inet6></code> <code><show-routing-instances></code>
Syntax (MX Series Router)	<code>show system connections</code> <code><extensive></code> <code><all-members></code> <code><inet inet6></code> <code><local></code> <code><member <i>member-id</i>></code> <code><show-routing-instances></code>
Syntax (QFX Series)	<code>show system connections</code> <code><extensive></code> <code><inet></code> <code><infrastructure <i>name</i>></code> <code><interconnect-device <i>name</i>></code> <code><node-group <i>name</i>></code> <code><show-routing-instances></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

Options **none**—Display information about all active IP sockets on the Routing Engine.

extensive—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system connection activity for all the routers in the chassis.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for all connected T1600 or T4000 LCCs

all-members—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for all members of the Virtual Chassis configuration.

inet | inet6—(Optional) Display IPv4 connections or IPv6 connections, respectively.

infrastructure name—(QFabric systems only) (Optional) Display system connection activity for the fabric control Routing Engines or fabric manager Routing Engines.

interconnect-device name—(QFabric systems only) (Optional) Display system connection activity for the Interconnect device.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display system connection activity for the Node group.

scc—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix Plus router.

show-routing-instances—(Optional) Display routing instances.

Additional Information By default, when you issue the **show system connections** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system connections on page 947](#)
[show system connections extensive on page 947](#)
[show system connections lcc \(TX Matrix Router\) on page 949](#)
[show system connections show-routing-instances on page 949](#)
[show system connections \(TX Matrix Plus Router\) on page 950](#)
[show system connections sfc \(TX Matrix Plus Router\) on page 953](#)
[show system connections show-routing-instances \(TX Matrix Plus Router\) on page 955](#)
[show system connections \(QFX3500 Switch\) on page 960](#)

Output Fields [Table 84 on page 946](#) describes the output fields for the **show system connections** command. Output fields are listed in the approximate order in which they appear.

Table 84: show system connections Output Fields

Field Name	Field Description
Proto	Protocol of the socket: IP , TCP , or UDP for IPv4 or IPv6.
Recv-Q	Number of input packets received by the protocol and waiting to be processed by the application.
Send-Q	Number of output packets sent by the application and waiting to be processed by the protocol.

Table 84: show system connections Output Fields (*continued*)

Field Name	Field Description
Local Address	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
Foreign Address	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.
Routing Instance	(Displayed only when the show-routing-instance option is used.) Routing instances associated with active IP sockets on the Routing Engine.
(state)	For TCP, the protocol state of the socket.

Sample Output

show system connections

```

user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp      0      2 192.168.4.16.513       208.197.169.254.894    ESTABLISHED
tcp      0      0 192.168.4.16.513       208.197.169.195.945    ESTABLISHED
tcp      0      0 *.23                   *.*                     LISTEN
tcp      0      0 *.22                   *.*                     LISTEN
tcp      0      0 *.513                  *.*                     LISTEN
tcp00 *.514             *.*                     LISTEN
tcp 0 0*.21                   *.*                     LISTEN
tcp00 *.79              *.*                     LISTEN
tcp 00 *.1023                *.*                     LISTEN
tcp 00 *.111                 *.*                     LISTEN
udp00192.168.4.16.1634   208.197.169.249.2049
udp00192.168.4.16.1627   208.197.169.254.2049
udp00192.168.4.16.1371   208.197.169.195.2049
udp00*. *              *.*
udp00*.9999             *.*
udp00 *.161             *.*
udp00192.168.4.16.1039   192.168.4.16.1023
udp00192.168.4.16.1038   192.168.4.16.1023
udp 00 192.168.4.16.1037     192.168.4.16.1023
udp00192.168.4.16.1036   192.168.4.16.1023
udp00*.1022             *.*
udp00*.1023             *.*
udp00*.111              *.*
udp00*. *               *.*

```

show system connections extensive

```

user@host> show system connections extensive

Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp4      0      6 192.168.187.15.23

```

```

172.27.133.138.3013 ESTABLISHED
sndsbcc: 6 sndsbmbcnt: 256 sndsbmbmax: 272000
sndsblowat: 2048 sndsbhiwat: 34000
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 533120
rcvsblowat: 1 rcvsbhiwat: 66640
proc id: 0 proc name:
iss: 2566994072 sndup: 2566994491
snduna: 2566994491 sndnxt: 2566994494 sndwnd: 64094
sndmax: 2566994494 sndcwnd: 6589 sndsssthresh: 2720
irs: 236981199 rcvup: 236981325
rcvnxt: 236981327 rcvadv: 237046862 rcvwnd: 66640
rtt: 140058623 srtt: 15519 rttv: 908
rxtcur: 1200 rxtshift: 0 rtseq: 2566994491
rttmin: 1000 mss: 1360
flags: SACK_PERMIT [0x2000200]
tcp4 0 0 10.255.165.93.179
10.255.165.203.65141 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 2555995917 sndup: 2555995917
snduna: 2555995917 sndnxt: 2555995917 sndwnd: 16384
sndmax: 2555995917 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2123825753 rcvup: 2123860681
rcvnxt: 2123860681 rcvadv: 2123877065 rcvwnd: 16384
rtt: 0 srtt: 3309 rttv: 72
rxtcur: 1200 rxtshift: 0 rtseq: 2555995898
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
tcp4 0 0 10.255.165.203.65141
10.255.165.93.179 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 5022 proc name: rpd
iss: 2123825753 sndup: 2123860662
snduna: 2123860681 sndnxt: 2123860681 sndwnd: 16384
sndmax: 2123860681 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2555995917 rcvup: 2555995917
rcvnxt: 2555995917 rcvadv: 2556012301 rcvwnd: 16384
rtt: 0 srtt: 3279 rttv: 22
rxtcur: 1200 rxtshift: 0 rtseq: 2123860662
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x100003e0]
tcp4 0 0 10.255.165.203.179
10.255.165.113.52404 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 1109297190 sndup: 1109332099
snduna: 1109332118 sndnxt: 1109332118 sndwnd: 16384
sndmax: 1109332118 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 1476831634 rcvup: 1476866449
rcvnxt: 1476866449 rcvadv: 1476882833 rcvwnd: 16384
rtt: 0 srtt: 3235 rttv: 18
rxtcur: 1200 rxtshift: 0 rtseq: 1109332099

```



```

rttmin:      1000  mss:      500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]

```

show system connections lcc (TX Matrix Router)

```
user@host> show system connections lcc 2
```

```
lcc2-re0:
```

```
-----
Active Internet connections (including servers)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	(state)
tcp4	0	0	192.168.66.131.1342	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2059	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.4571	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2496	192.168.66.130.23	ESTABLISHED
tcp4	0	0	*.3221	*.*	LISTEN
tcp4	0	0	*.23	*.*	LISTEN
tcp4	0	0	*.22	*.*	LISTEN
tcp4	0	0	*.514	*.*	LISTEN
tcp4	0	0	*.513	*.*	LISTEN
tcp4	0	0	*.21	*.*	LISTEN
tcp4	0	0	*.79	*.*	LISTEN
tcp4	0	0	*.6234	*.*	LISTEN
udp4	0	0	*.514	*.*	
udp4	0	0	*.6333	*.*	

show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
```

```
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address      Foreign Address     Routing Instance
(state)
```

tcp4	0	0	192.168.69.204.23	172.17.28.19.4267	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.58540	10.209.7.138.23	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.23	172.17.28.19.1098	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.57668	192.168.9.1.179	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.179	192.168.8.1.49209	default
			ESTABLISHED		
tcp4	0	0	128.0.0.1.6234	128.0.3.17.1024	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.9000	128.0.0.4.59103	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.59103	128.0.0.4.9000	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	*.32012	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.9000	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.33007	*.*	
__juniper_private2__			LISTEN		
tcp46	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.6154	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.6153	*.*	

```

__juniper_private1__ LISTEN
tcp4      0      0 *.7000    *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.6152    *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.6156    *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.33005   *. *
__juniper_private2__ LISTEN
tcp4      0      0 *.31343   *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.31341   *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.32003   *. *
__juniper_private2__ LISTEN
tcp4      0      0 *.666     *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.38      *. *
__juniper_private1__ LISTEN
tcp4      0      0 *.3221    *. *
LISTEN                                          default

```

show system connections (TX Matrix Plus Router)

```

user@host> show system connections
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address                               Foreign Address
                               (state)
tcp4      0      3 192.168.178.11.23
172.17.28.19.3565      ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.17.28.204.62719    ESTABLISHED
tcp4      0      0 192.168.178.11.23
192.168.69.199.51255   ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.24.26.227.42860    ESTABLISHED
tcp4      0      0 *.6156    *. *
LISTEN
tcp4      0      0 162.0.0.4.32012
ESTABLISHED            162.0.0.5.58935
tcp4      0      0 *.32012    *. *
LISTEN
tcp4      0      0 *.33007    *. *
LISTEN
tcp4      0      0 *.666     *. *
LISTEN
tcp4      0      0 162.0.0.4.6161
ESTABLISHED            162.0.0.5.62026
tcp4      0      0 *.33005    *. *
LISTEN
tcp4      0      0 162.0.0.4.9000
ESTABLISHED            162.0.0.4.51611
tcp4      0      0 162.0.0.4.51611
ESTABLISHED            162.0.0.4.9000
tcp4      0      0 *.6151    *. *
LISTEN
tcp4      0      0 *.6154    *. *
LISTEN
tcp4      0      0 *.6153    *. *

```

```

tcp4      0      0 *.31343      LISTEN      *. *
tcp4      0      0 *.31341      LISTEN      *. *
tcp4      0      0 *.9000       LISTEN      *. *
tcp4      0      0 *.6152       LISTEN      *. *
tcp4      0      0 *.32003      LISTEN      *. *
tcp4      0      0 *.33009      LISTEN      *. *
tcp4      0      0 *.3221       LISTEN      *. *
tcp4      0      0 *.23         LISTEN      *. *
tcp4      0      0 *.22         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.21         LISTEN      *. *
tcp4      0      0 *.79         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.6234       LISTEN      *. *
udp4      0      0 127.0.0.1.123 LISTEN      *. *
udp4      0      0 10.255.178.11.123 LISTEN      *. *
udp4      0      0 *.123        LISTEN      *. *
udp46     0      0 *.514        LISTEN      *. *
udp4      0      0 *.514        LISTEN      *. *
udp46     0      0 *.62027      LISTEN      *. *
udp4      0      0 *.59363      LISTEN      *. *
udp4      0      0 *.31342      LISTEN      *. *
udp46     0      0 *.161        LISTEN      *. *
udp4      0      0 *.161        LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp46     0      0 *.49152      LISTEN      *. *
udp46     0      0 *.4784       LISTEN      *. *
udp46     0      0 *.3784       LISTEN      *. *
udp4      0      0 *.49152      LISTEN      *. *
udp4      0      0 *.4784       LISTEN      *. *
udp4      0      0 *.3784       LISTEN      *. *
udp4      0      0 *.6333       LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *

```

```
lcc0-re0:
```

```
-----
Active Internet connections (including servers)
```

```

Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 192.168.178.3.23

```

```

172.24.26.227.50399
tcp4      0      0 *.6234          ESTABLISHED      *.*
          0      0 *.7000          LISTEN            *.*
          0      0 *.9000          LISTEN            *.*
          0      0 *.33009         LISTEN            *.*
          0      0 *.3221          LISTEN            *.*
          0      0 *.23            LISTEN            *.*
          0      0 *.22            LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.513           LISTEN            *.*
          0      0 *.21            LISTEN            *.*
          0      0 *.79            LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.513           LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.59924         LISTEN            *.*
          0      0 *.59412         LISTEN            *.*
          0      0 *.161           LISTEN            *.*
          0      0 *.161           LISTEN            *.*
          0      0 *.31342         LISTEN            *.*
          0      0 *.6333          LISTEN            *.*

```

lcc1-re0:

Active Internet connections (including servers)

Proto	Recv-Q	Send-Q	Local Address (state)	Foreign Address
tcp4	0	0	*.6234 LISTEN	*.*
tcp4	0	0	*.7000 LISTEN	*.*
tcp4	0	0	*.9000 LISTEN	*.*
tcp4	0	0	*.3221 LISTEN	*.*
tcp4	0	0	*.23 LISTEN	*.*
tcp4	0	0	*.22 LISTEN	*.*
tcp4	0	0	*.514 LISTEN	*.*
tcp4	0	0	*.513 LISTEN	*.*
tcp4	0	0	*.21 LISTEN	*.*
tcp4	0	0	*.79 LISTEN	*.*

```

tcp4      0      0 *.514          *.*
          LISTEN
tcp4      0      0 *.513          *.*
          LISTEN
tcp4      0      0 *.33009        *.*
          LISTEN
udp46     0      0 *.514          *.*
udp4      0      0 *.514          *.*
udp46     0      0 *.59924        *.*
udp4      0      0 *.59412        *.*
udp4      0      0 *.31342        *.*
udp46     0      0 *.161          *.*
udp4      0      0 *.161          *.*
udp4      0      0 *.6333         *.*

```

lcc2-re0:

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 *.6234            *.*
          LISTEN
tcp4      0      0 *.7000            *.*
          LISTEN
tcp4      0      0 *.9000            *.*
          LISTEN
tcp4      0      0 *.33009           *.*
          LISTEN
tcp4      0      0 *.3221            *.*
          LISTEN
tcp4      0      0 *.23              *.*
          LISTEN
tcp4      0      0 *.22              *.*
          LISTEN
tcp4      0      0 *.514             *.*
...

```

show system connections sfc (TX Matrix Plus Router)

```

user@host> show system connections sfc 0
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 162.0.0.4.514        132.0.0.4.952
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        131.0.0.4.694
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        130.0.0.4.860
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        129.0.0.4.716
          TIME_WAIT
tcp4      0      0 162.0.0.4.996        132.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.798        131.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.995        130.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.895        129.0.0.4.514
          TIME_WAIT

```

tcp4	0	0	192.168.178.11.21		
172.17.28.204.64662				TIME_WAIT	
tcp4	0	0	192.168.178.11.21		
172.17.28.204.51612				TIME_WAIT	
tcp4	0	0	*,6156		*,*
			LISTEN		
tcp4	0	0	*,9000		*,*
			LISTEN		
tcp4	0	0	*,666		*,*
			LISTEN		
tcp4	0	2	192.168.178.11.23		
172.17.28.19.3565				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.17.28.204.62719				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
192.168.69.199.51255				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.24.26.227.42860				ESTABLISHED	
tcp4	0	0	162.0.0.4.32012		162.0.0.5.58935
			ESTABLISHED		
tcp4	0	0	*,32012		*,*
			LISTEN		
tcp4	0	0	*,33007		*,*
			LISTEN		
tcp4	0	1432	162.0.0.4.6161		162.0.0.5.62026
			ESTABLISHED		
tcp4	0	0	*,33005		*,*
			LISTEN		
tcp4	0	0	162.0.0.4.9000		162.0.0.4.51611
			FIN_WAIT_2		
tcp4	0	0	162.0.0.4.51611		162.0.0.4.9000
			CLOSE_WAIT		
tcp4	0	0	*,6151		*,*
			LISTEN		
tcp4	0	0	*,6154		*,*
			LISTEN		
tcp4	0	0	*,6153		*,*
			LISTEN		
tcp4	0	0	*,31343		*,*
			LISTEN		
tcp4	0	0	*,31341		*,*
			LISTEN		
tcp4	0	0	*,6152		*,*
			LISTEN		
tcp4	0	0	*,32003		*,*
			LISTEN		
tcp4	0	0	*,33009		*,*
			LISTEN		
tcp4	0	0	*,3221		*,*
			LISTEN		
tcp4	0	0	*,23		*,*
			LISTEN		
tcp4	0	0	*,22		*,*
			LISTEN		
tcp4	0	0	*,514		*,*
			LISTEN		
tcp4	0	0	*,513		*,*
			LISTEN		
tcp4	0	0	*,21		*,*
			LISTEN		
tcp4	0	0	*,79		*,*

```

                                LISTEN
tcp4      0      0 *.514                                *.*
                                LISTEN
tcp4      0      0 *.513                                *.*
                                LISTEN
tcp4      0      0 *.6234                               *.*
                                LISTEN
udp4      0      0 127.0.0.1.123                       *.*
udp4      0      0 10.255.178.11.123                   *.*
udp4      0      0 *.123                                *.*
udp46     0      0 *.514                                *.*
udp4      0      0 *.514                                *.*
udp46     0      0 *.50895                              *.*
udp4      0      0 *.50794                              *.*
udp4      0      0 *.31342                              *.*
udp46     0      0 *.161                                *.*
udp4      0      0 *.161                                *.*
udp4      0      0 *.31340                              *.*
udp4      0      0 *.31340                              *.*
udp46     0      0 *.49152                              *.*
udp46     0      0 *.4784                               *.*
udp46     0      0 *.3784                               *.*
udp4      0      0 *.49152                              *.*
udp4      0      0 *.4784                               *.*
udp4      0      0 *.3784                               *.*
udp4      0      0 *.6333                               *.*
ip4       104    0 *.*                                  *.*
ip4       0      0 *.*                                  *.*
ip4       0      0 *.*                                  *.*

```

show system connections show-routing-instances (TX Matrix Plus Router)

```

user@host> show system connections show-routing-instances
sfc0-re0:
-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
-----
                                Routing Instance      (state)
tcp4      0      0 *.6156                   __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.9000                   __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.666                   __juniper_private1__    LISTEN      *.*
tcp4      0      2 192.168.178.11.23        default                 ESTABLISHED 172.17.28.19.3565
tcp4      0      0 192.168.178.11.23        default                 ESTABLISHED 172.17.28.204.62719
tcp4      0      0 192.168.178.11.23        default                 ESTABLISHED 192.168.69.199.51255
tcp4      0      0 192.168.178.11.23        default                 ESTABLISHED 172.24.26.227.42860
tcp4      0      0 162.0.0.4.32012          __juniper_private1__    ESTABLISHED 162.0.0.5.58935
tcp4      0      0 *.32012                  __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.33007                  __juniper_private2__    LISTEN      *.*
tcp4      0      0 162.0.0.4.6161          __juniper_private1__    ESTABLISHED 162.0.0.5.62026
tcp4      0      0 *.33005

```

tcp4	0	0	162.0.0.4.9000	__juniper_private2__	LISTEN	162.0.0.4.51611
tcp4	0	0	162.0.0.4.51611	__juniper_private1__	FIN_WAIT_2	162.0.0.4.9000
tcp4	0	0	*.6151	__juniper_private1__	CLOSE_WAIT	*.*
tcp4	0	0	*.6154	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6153	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31343	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31341	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6152	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.32003	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
udp4	0	0	127.0.0.1.123	__juniper_private1__	LISTEN	*.*
udp4	0	0	10.255.178.11.123	default		*.*
udp4	0	0	*.123	default		*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp46	0	0	*.50895	default		*.*
udp4	0	0	*.50794	default		*.*
udp4	0	0	*.31342	default		*.*
udp46	0	0	*.161	__juniper_private1__		*.*
udp4	0	0	*.161	default		*.*
				default		


```

udp4      0      0 *.31340      __juniper_private2__      *.*
udp4      0      0 *.31340      __juniper_private1__      *.*
udp46     0      0 *.49152      default                    *.*
udp46     0      0 *.4784       default                    *.*
udp46     0      0 *.3784       default                    *.*
udp4      0      0 *.49152      default                    *.*
udp4      0      0 *.4784       default                    *.*
udp4      0      0 *.3784       default                    *.*
udp4      0      0 *.6333       __juniper_private1__      *.*
ip4       0      0 *.*         default                    *.*
ip4       0      0 *.*         default                    *.*
ip4       0      0 *.*         default                    *.*

```

```
lcc0-re0:
```

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000              __juniper_private1__ LISTEN        *.*
tcp4      0      0 192.168.178.3.23    default              ESTABLISHED  172.24.26.227.50399
tcp4      0      0 *.6234              __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.9000              __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.33009             __juniper_private2__ LISTEN        *.*
tcp4      0      0 *.3221              default              LISTEN        *.*
tcp4      0      0 *.23                default              LISTEN        *.*
tcp4      0      0 *.22                default              LISTEN        *.*
tcp4      0      0 *.514               default              LISTEN        *.*
tcp4      0      0 *.513               default              LISTEN        *.*
tcp4      0      0 *.21                default              LISTEN        *.*
tcp4      0      0 *.79                default              LISTEN        *.*
tcp4      0      0 *.514               __juniper_private1__ LISTEN        *.*
tcp4      0      0 *.513               __juniper_private1__ LISTEN        *.*
udp46     0      0 *.514               default                    *.*
udp4      0      0 *.514

```

```

udp46      0      0 *.59924    default          *.
udp4        0      0 *.59412    default          *.
udp46      0      0 *.161      default          *.
udp4        0      0 *.161      default          *.
udp4        0      0 *.31342    default          *.
udp4        0      0 *.6333     __juniper_private1__
                        __juniper_private1__

```

```
lcc1-re0:
```

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000         __juniper_private1__ LISTEN         *.
tcp4      0      0 *.6234         __juniper_private1__ LISTEN         *.
tcp4      0      0 *.9000         __juniper_private1__ LISTEN         *.
tcp4      0      0 *.3221         __juniper_private1__ LISTEN         *.
tcp4      0      0 *.23           default             LISTEN         *.
tcp4      0      0 *.22           default             LISTEN         *.
tcp4      0      0 *.514          default             LISTEN         *.
tcp4      0      0 *.513          default             LISTEN         *.
tcp4      0      0 *.21           default             LISTEN         *.
tcp4      0      0 *.79           default             LISTEN         *.
tcp4      0      0 *.514          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.513          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.33009        __juniper_private2__ LISTEN         *.
udp46     0      0 *.514          default             *.
udp4       0      0 *.514          default             *.
udp46     0      0 *.59924        default             *.
udp4       0      0 *.59412        default             *.
udp4       0      0 *.31342        default             *.
udp46     0      0 *.161          __juniper_private1__ *.
udp4       0      0 *.161          default             *.
udp4       0      0 *.6333          default             *.
                        __juniper_private1__

```

lcc2-re0:

Active Internet connections (including servers) (including routing-instances)						
Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	default	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp4	0	0	*.31342	__juniper_private1__		*.*
udp46	0	0	*.62103	default		*.*
udp4	0	0	*.59924	default		*.*
udp46	0	0	*.161	default		*.*
udp4	0	0	*.161	default		*.*
udp4	0	0	*.6333	__juniper_private1__		*.*

lcc3-re0:

Active Internet connections (including servers) (including routing-instances)						
Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private1__		*.*

tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*,*
				default	LISTEN	*,*
tcp4	0	0	*.23	default	LISTEN	*,*
tcp4	0	0	*.22	default	LISTEN	*,*
tcp4	0	0	*.514	default	LISTEN	*,*
tcp4	0	0	*.513	default	LISTEN	*,*
tcp4	0	0	*.21	default	LISTEN	*,*
tcp4	0	0	*.79	default	LISTEN	*,*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*,*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*,*
udp46	0	0	*.514	default		*,*
udp4	0	0	*.514	default		*,*
udp46	0	0	*.62103	default		*,*
udp4	0	0	*.59924	default		*,*
udp4	0	0	*.31342	__juniper_private1__		*,*
udp46	0	0	*.161	default		*,*
udp4	0	0	*.161	default		*,*
udp4	0	0	*.6333	__juniper_private1__		*,*

show system connections (QFX3500 Switch)

```

user@switch> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
          (state)
tcp4      0      0 10.94.204.110.23        172.17.28.19.1308    ESTABLISHED
tcp4      0      0 128.0.0.1.6234          128.0.0.1.65142     ESTABLISHED
tcp4      0      0 128.0.0.1.65142         128.0.0.1.6234     ESTABLISHED
tcp4      0      0 128.0.0.1.33003         128.0.0.1.61441     ESTABLISHED
tcp4      0      0 128.0.0.1.61441         128.0.0.1.33003     ESTABLISHED
tcp46     0      0 *.179                   *.*                  LISTEN
tcp4      0      0 *.179                   *.*                  LISTEN
tcp4      0      0 128.0.0.16.9000         128.0.0.16.50970    ESTABLISHED
tcp4      0      0 128.0.0.16.50970        128.0.0.16.9000     ESTABLISHED
tcp4      0      0 *.38                     *.*                  LISTEN

```

			LISTEN	
tcp4	0	0 *.3491		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.59437
			ESTABLISHED	
tcp4	0	0 128.0.0.1.59437		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33023		128.0.0.1.63605
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63605		128.0.0.1.33023
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.63830
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63830		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 *.667		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.7000		128.0.0.1.51580
			ESTABLISHED	
tcp4	0	0 128.0.0.1.51580		128.0.0.1.7000
			ESTABLISHED	
tcp4	0	0 128.0.0.1.6234		128.0.0.1.53646
			ESTABLISHED	
tcp4	0	0 *.33001		*.*
			LISTEN	
tcp4	0	0 *.33003		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.53646		128.0.0.1.6234
			ESTABLISHED	
tcp4	0	0 128.0.0.16.9000		128.0.0.16.63454
			ESTABLISHED	
tcp4	0	0 128.0.0.16.63454		128.0.0.16.9000
			ESTABLISHED	
tcp4	0	0 *.666		*.*
			LISTEN	
tcp4	0	0 *.7000		*.*
			LISTEN	
tcp4	0	0 *.51627		*.*
			LISTEN	
tcp4	0	0 *.3492		*.*
			LISTEN	
tcp4	0	0 *.33023		*.*
			LISTEN	
tcp4	0	0 *.33013		*.*
			LISTEN	
tcp4	0	0 *.7202		*.*
			LISTEN	
tcp4	0	0 *.6151		*.*
			LISTEN	
tcp4	0	0 *.9000		*.*
			LISTEN	
tcp4	0	0 *.6161		*.*
			LISTEN	
tcp4	0	0 *.6011		*.*
			LISTEN	
tcp4	0	0 *.3221		*.*
			LISTEN	

tcp4	0	0 *.23		*. *
			LISTEN	
tcp4	0	0 *.22		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.21		*. *
			LISTEN	
tcp4	0	0 *.79		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.1127		*. *
			LISTEN	
tcp4	0	0 *.1129		*. *
			LISTEN	
tcp4	0	0 *.1128		*. *
			LISTEN	
tcp4	0	0 *.6234		*. *
			LISTEN	
udp46	0	0 *.514		*. *
udp4	0	0 *.514		*. *
udp4	0	0 128.0.0.1.123		*. *
udp46	0	0 *.53344		*. *
udp4	0	0 *.54261		*. *
udp46	0	0 *.161		*. *
udp4	0	0 *.161		*. *
udp4	0	0 *.31342		*. *
udp4	0	0 *.59137		*. *
udp4	0	0 *. *		*. *
udp46	0	0 *.49152		*. *
udp46	0	0 *.4784		*. *
udp46	0	0 *.3784		*. *
udp4	0	0 *.49152		*. *
udp4	0	0 *.4784		*. *
udp4	0	0 *.3784		*. *
udp4	0	0 10.255.204.110.123		*. *
udp4	0	0 *.123		*. *
udp4	0	0 *.67		*. *
udp4	0	0 *.6333		*. *
udp4	0	0 *.2293		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *

show system core-dumps

List of Syntax	Syntax on page 963 Syntax (EX Series Switches) on page 963 Syntax (TX Matrix Router) on page 963 Syntax (TX Matrix Plus Router) on page 963 Syntax (QFX Series) on page 963
Syntax	<pre>show system core-dumps <brief detail> <core-filename> <core-file-info></pre>
Syntax (EX Series Switches)	<pre>show system core-dumps <all-members> <brief detail> <core-filename> <core-file-info> <local> <member member-id></pre>
Syntax (TX Matrix Router)	<pre>show system core-dumps <all-chassis all-lcc lcc number scc> <brief detail> <core-filename> <core-file-info></pre>
Syntax (TX Matrix Plus Router)	<pre>show system core-dumps <all-chassis all-lcc lcc number sfc number> <brief detail> <core-filename> <core-file-info></pre>
Syntax (QFX Series)	<pre>show system core-dumps <brief detail> <component (UUID serial number all)> <core-file-info component (UUID serial number) core-file-name> <display-period (hours minutes seconds)> <display-order> <kernel-crashinfo component (UUID serial number)> <repository (core log)></pre>
Release Information	<p>Command introduced before Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Show core files on all routers or switches running Junos OS. You can use the show system core-dumps command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and</p>

path and filename. On a QFabric system, you can view core-dump files on individual QFabric system devices as well as on the entire QFabric system.

You can use the option **core-filename** and its options **core-file-info**, **brief**, and **detail** to display more information about the specified core-dump files.

Options **none**—Display a list of all existing core-dump files.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on a TX Matrix router, display system core files for the TX Matrix router switch-card chassis [SCC] and all the T640 routers [LCCs] connected to the TX Matrix router.

On a routing matrix based on a TX Matrix Plus router, display system core files for the TX Matrix Plus router (switch-fabric chassis [SFC]) and all the T1600 routers [LCCs] connected to the TX Matrix Plus router.

<all-lcc | lcc number>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on the TX Matrix router, display core dump files for all T640 routers (line-card chassis [LCCs]) or a specific T640 router [LCC] connected to the TX Matrix router.

On a routing matrix based on the TX Matrix Plus router, display logging information for all T1600 routers (line-card chassis [LCCs]) or a specific T1600 router (LCC) connected to the TX Matrix Plus router. When using the **lcc number** option, replace **number** with a value from 0 through 3.



NOTE: The **all-chassis** option displays system core files for the SCC or SFC and the LCCs connected to the SCC or SFC in the routing matrix while the **all-lcc** option only displays system core files for the LCCs in the routing matrix.

all-members—(EX4200 switches) (Optional) Display system core files on all members of the Virtual Chassis configuration.

brief—(Optional) View details of a binary file.

component (UUID | serial number | all)—(QFabric systems only) (Optional) Display a list of core-dump files located on individual QFabric system device or on the entire QFabric system.

core-file-info—(Optional) Display the stack trace of a core file.

core-filename—(Optional) Name of a specific core file to display.

detail—(Optional) View stack trace with details of the binary file.

display-order (timestamp-sort | alphanumeric-sort)—(QFabric systems only) (Optional) Display list of debug artifacts generated within the specified period—for example,

within the last hour, within the last 20 minutes, or within the last 32 seconds—or according to their filename.

display-period (*hours | minutes | seconds*)—(QFabric systems only) (Optional) Display core-dump files generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds.

kernel-crashinfo component (*UUID | serial number*)—(QFabric systems only) (Optional) Display kernel crash information from the EEPROM on a QFabric system device.

local—(EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

repository (*core | log*)—(QFabric systems only) (Optional) Specify either the core or log repository in which to view core-dump files.

scc—(TX Matrix routers only) (Optional) Display system core files on the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display system core files on the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

view

List of Sample Output

[show system core-dumps on page 967](#)
[show system core-dumps on page 967](#)
[show system core-dumps \(TX Matrix Plus Router\) on page 967](#)
[show system core-dumps \(QFX3500 Switch\) on page 969](#)
[show system core-dumps \(QFabric Systems\) on page 969](#)
[show system core-dumps core-file-info component serial number core-file-name \(QFabric Systems\) on page 970](#)
[show system core-dumps component serial number display-order alphanumeric-sort repository core \(QFabric Systems\) on page 970](#)
[show system core-dumps display-period \(QFabric Systems\) on page 970](#)
[show system core-dumps kernel-crashinfo component serial number \(QFabric Systems\) on page 972](#)
[show system core-dumps repository core \(QFabric Systems\) on page 974](#)
[show system core-dumps repository log \(QFabric Systems\) on page 974](#)

Output Fields

[Table 85 on page 965](#) describes the output fields for the **show system core-dumps** command. Output fields are listed in the approximate order in which they appear.

Table 85: show system core-dumps Output Fields

Field Name	Field Description
<i>Permissions</i>	Read/write permissions for the file named.

Table 85: show system core-dumps Output Fields (*continued*)

Field Name	Field Description
<i>Links</i>	Number of links to the file.
<i>Owner</i>	Name of the file owner.
<i>Group</i>	Name of the group with file access.
<i>File size</i>	File size in bytes.
<i>Modified</i>	Last file modification date and time.
<i>Path/filename</i>	File path where the file resides and the filename.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the core repository, and the log files are located in the log repository. The default Repository scope is shared since both the core and log repositories are shared by all of the QFabric system devices.
Repository head:	Path to the top-level repository location.
Repository name:	Name of the repository: core or log .
List of nodes for core repository:	List of core-dump files associated with a particular QFabric system device located in the core repository.
Node Group	Name of the QFabric system device.
Node Identifier	UUID or serial number of the QFabric system device.
Num	Number of core-dump and log files.
Model	Model number of the QFabric system device.
Usage	Usage of the repository in megabytes.
Total usage of core repository:	Total usage of core-dump files associated with a particular QFabric system device located in the core repository. Usage is specified in megabytes and as a percentage.
Total usage of log repository:	Total usage of log files associated with a particular QFabric system device located in the log repository. Usage is specified in megabytes and as a percentage.
List of nodes for core repository:	List of core-dump files associated with a particular QFabric system device located in the core repository.
List of nodes for log repository:	List of log files associated with a particular QFabric system device located in the log repository.

Table 85: show system core-dumps Output Fields (*continued*)

Field Name	Field Description
Filename	Name of the core-dump file.
Date	Last core-dump file modification date and time.
Size	Size of the core-dump file.
Core filename	Filename of the core-dump file.
Process name	Name of the process that is generating a core-dump file or log file.
Release	Junos OS release.
Build server	Junos OS build server.
Build date	Junos OS build date.
Stack trace	Stack trace of the core-dump file.

Sample Output

show system core-dumps

This example shows the command output if core files exist.

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

show system core-dumps

This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

show system core-dumps (TX Matrix Plus Router)

```
user@host> show system core-dumps
sfc0-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 1627592
-rw-r--r-- 1 root field 535346090 May 15 07:36
```

```

rpd.core-tarball.0.090515.0736.tgz
-rw-r--r-- 1 root field 105632057 May 15 07:37
rpd.core-tarball.1.090515.0737.tgz
-rw-r--r-- 1 root field 101981681 May 15 07:38
rpd.core-tarball.2.090515.0738.tgz
-rw-r--r-- 1 root field 85854573 May 15 07:40
rpd.core-tarball.3.090515.0740.tgz
-rw-r--r-- 1 root field 4157845 May 15 08:18
rpd.core-tarball.4.090515.0818.tgz

lcc0-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 12

lcc1-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 10024
-rw-r--r-- 1 root field 1875794 Apr 22 15:47
chassisd.core-tarball.0.090422.1547.tgz
-rw-r--r-- 1 root field 1894183 Apr 22 19:02
chassisd.core-tarball.0.090422.1902.tgz
-rw-r--r-- 1 root field 1290240 Apr 26 16:01 ksyncd_1558.core.0.090426.1601

lcc2-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 21124008
-rw-r--r-- 1 root wheel 1022376528 May 2 06:43
core-LCC2-EGFPC7.core.0.090502.0643
-rw-r--r-- 1 root wheel 1022376528 May 2 08:13
core-LCC2-EGFPC7.core.0.090502.0813
-rw-r--r-- 1 root wheel 1022376544 May 5 06:15
core-LCC2-EGFPC7.core.0.090505.0615
-rw-r--r-- 1 root wheel 1022376544 May 6 10:59
core-LCC2-EGFPC7.core.0.090506.1059
-rw-r--r-- 1 root wheel 1022376528 May 2 06:58
core-LCC2-EGFPC7.core.1.090502.0658
-rw-r--r-- 1 root wheel 754271232 May 5 06:33
core-LCC2-EGFPC7.core.1.090505.0633
-rw-r--r-- 1 root wheel 264897536 May 6 11:12
core-LCC2-EGFPC7.core.1.090506.1112
-rw-r--r-- 1 root wheel 1022376528 May 2 07:22
core-LCC2-EGFPC7.core.2.090502.0722
-rw-r--r-- 1 root wheel 163633152 May 5 06:52
```

```

core-LCC2-EGFPC7.core.2.090505.0652
-rw-r--r-- 1 root wheel 171312128 May 6 12:13
core-LCC2-EGFPC7.core.2.090506.1213
-rw-r--r-- 1 root wheel 1022376528 May 2 07:39
core-LCC2-EGFPC7.core.3.090502.0739
-rw-r--r-- 1 root wheel 1022376528 May 2 07:55
core-LCC2-EGFPC7.core.4.090502.0755
-rw-r--r-- 1 root wheel 427277312 May 7 04:47
core-LCC2-STFPC4.core.0.090507.0447
-rw-r--r-- 1 root wheel 419609600 May 7 04:47
core-LCC2-STFPC5.core.0.090507.0447
-rw-r--r-- 1 root wheel 432356352 May 7 04:47
core-LCC2-STFPC6.core.0.090507.0447

/var/tmp/cores:
total 2568
-rw-r--r-- 1 root field 1290240 May 14 14:26 ksyncd_1540.core.0.090514.1426
...

```

show system core-dumps (QFX3500 Switch)

```

user@switch> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw---- 1 root field 1545143 Jun 4 2012 /var/tmp/pafxpc.core.0.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.1.gz
-rw-rw---- 1 root field 1545141 Jun 4 2012 /var/tmp/pafxpc.core.2.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.3.gz
-rw-rw---- 1 root field 1545142 Jun 5 2012 /var/tmp/pafxpc.core.4.gz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 5

```

show system core-dumps (QFabric Systems)

```

user@switch> show system core-dumps
Repository scope: shared
Repository head: /pbdata/export
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	0M
NW-NG-0	BBAK0394	0	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	0M
NW-NG-0	d0afdale-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	0M
IC-WS001	WS001/YN5999	0	qfxc08-3008	0M
node-device1	BBAK0372	0	qfx3500	0M
node-device1	EE3093	0	qfx3500	0M

```

Total usage of core repository: 0M of 70000M (0.0%)

List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	0M

DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	OM
NW-NG-0	BBAK0394	1	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	OM
IC-WS001	WS001/YW3803	1	qfxc08-3008	OM
node-device1	BBAK0372	1	qfx3500	OM
node-device1	EE3093	1	qfx3500	OM

Total usage of log repository:OM of 70000M (0.0%)

show system core-dumps core-file-info component serial number core-file-name (QFabric Systems)

```

user@switch> show system core-dumps core-file-info component
e8ff4b3e-7d92-11e0-be5d-00e081c1fe0e cosd.core.0.1519.05162011131846.gz
Repository scope: shared
Repository head: /pbstorage
Repository name: core
Core filename: /pbstorage/rdumps/e8ff4b3e-7d92-11e0-be5d-
00e081c1fe0e/5658.cosd.core.0.1519.05162011131846
Process name: cosd
Release: 11.3I0
Build server: /c/ssengupta/dfx_ha_v1/obj-i386-dcp/dcp/usr.sbin/cosd
Build date: 2011-05-14 01:11:44 UTC
Stack trace:
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#1 0x887d4a45 in pselect () from /usr/lib/libc.so.6
#2 0x88774719 in pselect () from /usr/lib/libthr.so.2
#3 0x885de5db in __evGetNext () from /usr/lib/libisc.so.2
#4 0x885debf0 in __evMainLoop () from /usr/lib/libisc.so.2
#5 0x081125b2 in cosd_loop ()
#6 0x0812e19a in main ()

```

show system core-dumps component serial number display-order alphanumeric-sort repository core (QFabric Systems)

```

user@switch> show system core-dumps component BBAK8891 display-order alphanumeric-sort
repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of core dumps for component BBAK8891
Repository location: /pbdata/export/rdumps/BBAK8891

```

Filename	Date	Size
eswd.core.0.1361.11172011214257.gz	Nov 17 21:43:10 2011	4779553
eswd.core.1.80267.11172011214514.gz	Nov 17 21:45:19 2011	3541648
eswd.core.2.80682.11172011214535.gz	Nov 17 21:45:43 2011	2156683
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617

Number of core dumps in repository:4

show system core-dumps display-period (QFabric Systems)

```

user@switch> show system core-dumps display-period 24h
show system core-dumps display-period 24h
Repository scope: shared
Repository head: /pbdata/export
List of core dumps at repository: /pbdata/export/rdumps

```

Delta timespec: Last 24h

Component: BBAK8273

Filename	Size	Date
----------	------	------

vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375794
-------------------------------------	----------------------	--------

Component: cedb7b0e-0025-11e1-9a5f-00e081c52990

Filename	Size	Date
----------	------	------

vccpd.core.0.1461.11182011151131.gz	Nov 18 15:11:31 2011	120951
-------------------------------------	----------------------	--------

Component: ee19c4f8-0025-11e1-ae6f-00e081c52990

Filename	Size	Date
----------	------	------

vccpd.core.0.1462.11182011151131.gz	Nov 18 15:11:31 2011	109420
-------------------------------------	----------------------	--------

Component: BBAK8281

Filename	Size	Date
----------	------	------

vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	375373
-------------------------------------	----------------------	--------

Component: BBAK8891

Filename	Size	Date
----------	------	------

vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617
-------------------------------------	----------------------	--------

Component: BBAK8276

Filename	Size	Date
----------	------	------

vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:35 2011	375350
-------------------------------------	----------------------	--------

Component: BBAK8868

Filename	Size	Date
----------	------	------

vccpd.core.0.1196.11182011151130.gz	Nov 18 15:11:34 2011	376211
-------------------------------------	----------------------	--------

Component: BBAK8835

Filename	Size	Date
----------	------	------

vccpd.core.0.1195.11182011151130.gz	Nov 18 15:11:35 2011	375700
-------------------------------------	----------------------	--------

Component: BBAK8283

Filename	Size	Date
----------	------	------

vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:36 2011	368298
-------------------------------------	----------------------	--------

Component: YW3781/YW3781

Filename	Size	Date
----------	------	------

vccpd.core.0.1220.11182011151131.gz	Nov 18 15:11:38 2011	380002
-------------------------------------	----------------------	--------

Component: 09726be2-0026-11e1-82d9-00e081c52990

Filename	Size	Date
----------	------	------

vccpd.core.0.1461.11182011151130.gz	Nov 18 15:11:31 2011	119965
-------------------------------------	----------------------	--------

Component: BBAK8309

Filename	Size	Date
----------	------	------

vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	378930
-------------------------------------	----------------------	--------

Component: 303d476a-0026-11e1-abf4-00e081c52990

Filename	Size	Date
----------	------	------

vccpd.core.0.1460.11182011151131.gz	Nov 18 15:11:31 2011	118385
-------------------------------------	----------------------	--------

Component: YW3798/YW3798

Filename	Size	Date
----------	------	------

vccpd.core.0.1219.11182011151131.gz	Nov 18 15:11:36 2011	380455
-------------------------------------	----------------------	--------

List of log dumps at repository: /pbdata/export/rlogs

Delta timespec: Last 24h

Component: BBAK8273

Filename	Size	Date
----------	------	------

vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:39 2011	20415
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151131.tgz	Nov 18 15:11:33 2011	19651
Component: ee19c4f8-0025-11e1-aef6-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1462.11182011151133.tgz	Nov 18 15:11:36 2011	24650
Component: BBAK8281		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	19445
Component: BBAK8891		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:41 2011	21916
Component: BBAK8276		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:39 2011	20461
Component: BBAK8868		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	21924
Component: BBAK8835		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151137.tgz	Nov 18 15:11:39 2011	19424
Component: BBAK8283		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:42 2011	31186
Component: YW3781/YW3781		
Filename	Size	Date
vccpd.tarball.0.1220.11182011151141.tgz	Nov 18 15:11:45 2011	27565
Component: 09726be2-0026-11e1-82d9-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151130.tgz	Nov 18 15:11:34 2011	19613
Component: BBAK8309		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151138.tgz	Nov 18 15:11:46 2011	50362
Component: 303d476a-0026-11e1-abf4-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1460.11182011151133.tgz	Nov 18 15:11:33 2011	19360
Component: YW3798/YW3798		
Filename	Size	Date
vccpd.tarball.0.1219.11182011151140.tgz	Nov 18 15:11:49 2011	24473

show system core-dumps kernel-crashinfo component serial number (QFabric Systems)

```
user@switch> show system core-dumps kernel-crashinfo component A0001/YA0197
Node: A0001/YA0197
```

Information about previous kernel crash:

-- Kernel panic data --

Panic string: kdb_sysctl_panic

System uptime: 3 day 20 hr 59 min 40 sec Kernel crash time: 2011-11-15 Wed 15:25:17

Kernel build linkstamp: JUNOS 11.3I #0: 2011-11-10 20:42:27 UTC

-- Stacktrace of panicing context --

Processor 1 (crash monarch):

savectx+0x0 (c9552800,80214efc,802a7fbc,c88ad05c) ra 801b93a8 sz 0

kdm_kcore_save_crashinfo+0x254 (c9552800,0,802a7fbc,c88ad05c) ra 801b9f44 sz 784

kdm_kcore_kern_panic_event_handler+0x4b0 (c9552800,0,802a7fbc,c88ad05c) ra 8022a9b8 sz 88

panic+0x1d0 (c9552800,0,4,77fed534) ra 802540c0 sz 56

kdb_sysctl_panic+0x70 (c9552800,0,4,77fed534) ra 80237e58 sz 40 sysctl_root+0x12c

(c9552800,0,4,e8bc5cf8) ra 80238e50 sz 48

userland_sysctl+0x164 (c9552800,0,4,e8bc5cf8) ra 8023956c sz 104

__sysctl+0xe4 (c9552800,0,4,e8bc5cf8) ra 806d62e8 sz 160

trap+0xe1c (c9552800,0,4,e8bc5cf8) ra 80896e68 sz 128

MipsUserGenException+0x1a4 (c9552800,0,4,405cd12c) ra 0 sz 0

pid 82340, process: sysctl

Processor 0:

restoreintr+0x14 (1,81bca820,3,0) ra 806cdc3c sz 0

spinlock_exit+0x30 (1,81bca820,3,0) ra 8025d354 sz 24

sleepq_release+0x64 (1,81bca820,3,0) ra 8025e670 sz 24

sleepq_timeout+0x224 (1,81bca820,3,0) ra 80240294 sz 48

softclock+0x434 (1,81bca820,3,0) ra 802067f8 sz 80

ithread_loop+0x244 (1,81bca820,3,0) ra 80200e28 sz 64 fork_exit+0xc0

(1,81bca820,3,0) ra 80897c28 sz 48

MipsNMIException+0x34 (1,81bca820,3,0) ra 0 sz 0

pid 82340, process: sysctl

Processor 2:

cpu_idle+0x20 (80960000,51bbc,2031df,81bca1b8) ra 80204948 sz 24 idle_proc+0x130

(80960000,51bbc,2031df,81bca1b8) ra 80200e28 sz 56 fork_exit+0xc0

(80960000,51bbc,2031df,81bca1b8) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2031df,81bca1b8) ra 0 sz 0

pid 82340, process: sysctl

Processor 3:

cpu_idle+0x20 (80960000,51bbc,2038df,81bca300) ra 80204948 sz 24 idle_proc+0x130

(80960000,51bbc,2038df,81bca300) ra 80200e28 sz 56 fork_exit+0xc0

(80960000,51bbc,2038df,81bca300) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2038df,81bca300) ra 0 sz 0

pid 82340, process: sysctl

Processor 4:

cpu_idle+0x20 (80960000,51bbc,2037df,81bca448) ra 80204948 sz 24 idle_proc+0x130

(80960000,51bbc,2037df,81bca448) ra 80200e28 sz 56 fork_exit+0xc0

(80960000,51bbc,2037df,81bca448) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2037df,81bca448) ra 0 sz 0

pid 82340, process: sysctl

Processor 5:

restoreintr+0x14 (1,51bbc,203edf,81bca590) ra 806cdc3c sz 0

spinlock_exit+0x30 (1,51bbc,203edf,81bca590) ra 80204a34 sz 24 idle_proc+0x21c

(1,51bbc,203edf,81bca590) ra 80200e28 sz 56 fork_exit+0xc0

(1,51bbc,203edf,81bca590) ra 80897c28 sz 48

MipsNMIException+0x34 (1,51bbc,203edf,81bca590) ra 0 sz 0

pid 82340, process: sysctl

```

Processor 6:
cpu_idle+0x20 (80960000,51bbc,205cdf,81bca6d8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,205cdf,81bca6d8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,205cdf,81bca6d8) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,205cdf,81bca6d8) ra 0 sz 0
pid 82340, process: sysctl

Processor 7:
lockmgr+0x5ac (c97e8484,c8dd9800,0,c8dd9800) ra 8c11c81c sz 48
sal_sem_take+0x134 (c97e8484,c8dd9800,0,c8dd9800) ra 8c351108 sz 56
_bcm_esw_linkscan_thread+0x45c (c97e8484,c8dd9800,0,c8dd9800) ra 8c11cdb4 sz 104
sal_thread_start_wrap+0x74 (c97e8484,c8dd9800,0,c8dd9800) ra 80200e28 sz 32
fork_exit+0xc0 (c97e8484,c8dd9800,0,c8dd9800) ra 80897c28 sz 48
MipsNMIException+0x34 (c97e8484,c8dd9800,0,c8dd9800) ra 0 sz 0
pid 82340, process: sysctl
-- End of stacktrace --

```

show system core-dumps repository core (QFabric Systems)

```

user@switch> show system core-dumps repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	0M
NW-NG-0	BBAK0394	0	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	0M
IC-WS001	WS001/YN5999	0	qfxc08-3008	0M
node-device1	BBAK0372	0	qfx3500	0M
node-device1	EE3093	0	qfx3500	0M

Total usage of core repository: 0M of 70000M (0.0%)

show system core-dumps repository log (QFabric Systems)

```

user@switch> show system core-dumps repository log
Repository scope: shared
Repository head: /pbdata/export
Repository name: log
List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	0M
NW-NG-0	BBAK0394	1	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	0M

IC-WS001	WS001/YW3803	1	qfxc08-3008	0M
node-device1	BBAK0372	1	qfx3500	0M
node-device1	EE3093	1	qfx3500	0M
Total usage of log repository:0M of 70000M (0.0%)				

show system directory-usage

List of Syntax	Syntax on page 976 Syntax (EX Series) on page 976 Syntax (TX Matrix Router) on page 976 Syntax (TX Matrix Plus Router) on page 976 Syntax (MX Series Router) on page 976 Syntax (QFX Series) on page 976
Syntax	show system directory-usage <depth <i>number</i> > <path>
Syntax (EX Series)	show system directory-usage <all-members> <depth <i>number</i> > <local> <member <i>member-id</i> > <path>
Syntax (TX Matrix Router)	show system directory-usage <all-chassis all-lcc lcc <i>number</i> scc> <depth <i>number</i> > <path>
Syntax (TX Matrix Plus Router)	show system directory-usage <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <depth <i>number</i> > <path>
Syntax (MX Series Router)	show system directory-usage <all-members> <depth <i>number</i> > <local> <member <i>member-id</i> > <path>
Syntax (QFX Series)	show system directory-usage <depth <i>number</i> > <path> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display directory usage information.
Options	none—Display all directory usage information.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display directory usage information about all the T640 routers (in a routing matrix based on a TX Matrix router). Display directory usage information about all the T1600 or T4000 routers (in a routing matrix based on a TX Matrix Plus router) in the chassis.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

depth *number*—(Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

infrastructure *name*— (QFabric systems only) (Optional) Display directory information for the fabric control Routing Engines and fabric manager Routing Engines.

interconnect-device *name*— (QFabric systems only) (Optional) Display directory information for the Interconnect device.

node-group *name*— (QFabric systems only) (Optional) Display directory information for the Node group.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

path—(Optional) Path or root directory to traverse.

scc—(TX Matrix router only) (Optional) Display directory information for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display directory information for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system directory-usage scc \(TX Matrix Router\) on page 979](#)
[show system directory-usage sfc \(TX Matrix Plus Router\) on page 979](#)
[show system directory-usage \(QFX3500 Switch\) on page 979](#)

Output Fields [Table 86 on page 978](#) describes the output fields for the **show system directory-usage** command. Output fields are listed in the approximate order in which they appear.

Table 86: show system directory-usage Output Fields

Field Name	Field Description
<i>bytes</i>	Number of bytes used by files in a directory.
<i>directory-name</i>	Name of the directory.

Sample Output

show system directory-usage scc (TX Matrix Router)

```

user@host> show system directory-usage /var/tmp scc
/var/tmp
1.0K    /var/tmp/vi.recover
2.0K    /var/tmp/instmp.tPMk8u
1.0K    /var/tmp/install
        /var/tmp/instmp.GUMpur
4.8M    /var/tmp/instmp.GUMpur/packages
6.4M    /var/tmp/troy1
297M    /var/tmp/dsw
        /var/tmp/pkg_tmp.2073
83K     /var/tmp/pkg_tmp.2073/bin
        /var/tmp/instmp.oMIDb1
89K     /var/tmp/instmp.oMIDb1/bin
        /var/tmp/instmp.byhMjR
4.6M    /var/tmp/instmp.byhMjR/packages
        /var/tmp/instmp.6fqHf3
1.7M    /var/tmp/instmp.6fqHf3/packages
        /var/tmp/instmp.mljECe
4.6M    /var/tmp/instmp.mljECe/packages

```

show system directory-usage sfc (TX Matrix Plus Router)

```

user@switch> show system directory-usage /var/tmp sfc 0
sfc0-re0:
-----
/var/tmp
46K     /var/tmp/gres-tp
        /var/tmp/sec-download
2.0K    /var/tmp/sec-download/sub-download
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
795M    /var/tmp/cores
766K    /var/tmp/pr440594

```

show system directory-usage (QFX3500 Switch)

```

user@switch> show system directory-usage
/var/tmp
30K     /var/tmp/gres-tp
2.0K    /var/tmp/rtbdb
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
2.0K    /var/tmp/pics

```

show system license

Syntax	show system license <installed keys usage>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display licenses and information about how they are used.
Options	<p>none—Display all license information.</p> <p>installed—(Optional) Display installed licenses only.</p> <p>keys—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.</p> <p>usage—(Optional) Display the state of licensed features.</p>
Required Privilege Level	maintenance
List of Sample Output	show system license on page 981 show system license installed on page 981 show system license keys on page 982 show system license usage on page 982 show system license (QFX Series) on page 982
Output Fields	Table 87 on page 980 lists the output fields for the show system license command. Output fields are listed in the approximate order in which they appear.

Table 87: show system license Output Fields

Field Name	Field Description
Feature name	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.
Licenses used	<p>Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.</p> <p>NOTE: In Junos OS Release 10.1 and later, the Licenses used column displays the actual usage count based on the number of active sessions or connections as reported by the corresponding feature daemons. This is applicable for scalable license-based features such as Subscriber Access (scale-subscriber), L2TP (scale-l2tp), Mobile IP (scale-mobile-ip), and so on.</p>

Table 87: show system license Output Fields (*continued*)

Field Name	Field Description
Licenses installed	<p>Information about the installed license key:</p> <ul style="list-style-type: none"> • License identifier—Identifier associated with a license key. • State—State of the license key: valid or invalid. An invalid state indicates that the key was entered incorrectly or is not valid for the specific device. • License version—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. • Valid for device—Device that can use a license key. • Group defined—Group membership of a device. • Features—Feature associated with a license, such as data link switching (DLSw).
Licenses needed	Number of licenses required for features being used but not yet properly licensed.
Expiry	Amount of time left within the grace period before a license is required for a feature being used.

Sample Output

show system license

```
user@host> show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

Licenses installed:

License identifier: XXXXXXXXXX

License version: 2

Features:

```
subscriber-accounting - Per Subscriber Radius Accounting
permanent
subscriber-authentication - Per Subscriber Radius Authentication
permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
subscriber-ip - Dynamic and Static IP
permanent
```

show system license installed

```
user@host> show system license installed
```

License identifier: XXXXXXXXXX

License version: 2

Features:

- subscriber-accounting - Per Subscriber Radius Accounting
permanent
- subscriber-authentication - Per Subscriber Radius Authentication
permanent
- subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
- subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
- subscriber-ip - Dynamic and Static IP
permanent

show system license keys

```
user@host> show system license keys
XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxx
```

show system license usage

```
user@host> show system license usage
License usage:
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

show system license (QFX Series)

```
user@switch> show system license
License usage:
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
qfx-edge-fab	1	1	1	permanent

```
Licenses installed:
License identifier: JUNOS417988
License version: 1
Features:
  qfx-edge-fab - QFX3000 Series QF/Node feature license
                permanent
```

show system processes

List of Syntax	Syntax on page 983 Syntax (EX Series Switches) on page 983 Syntax (MX Series Routers) on page 983 Syntax (QFX Series) on page 983 Syntax (TX Matrix Routers) on page 983 Syntax (TX Matrix Plus Router) on page 983
Syntax	<pre>show system processes <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (EX Series Switches)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (MX Series Routers)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (QFX Series)	<pre>show system processes <brief detail extensive summary > <health (pid <i>process-identifer</i> process-name <i>process-name</i>)> <interconnect-device <i>name</i>> <node-group <i>name</i>> <providers> <resource-limits> <wide></pre>
Syntax (TX Matrix Routers)	<pre>show system processes <brief detail extensive summary> <all-chassis all-lcc lcc <i>number</i> scc> <wide></pre>
Syntax (TX Matrix Plus Router)	<pre>show system processes <brief detail extensive summary> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></pre>

<wide>

Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about software processes that are running on the router or switch and that have controlling terminals.
Options	<p>none—Display standard information about system processes.</p> <p>brief detail extensive summary—(Optional) Display the specified level of detail.</p> <p>adaptive-services—(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.</p> <p>alarm-control—(Optional) Display the process to configure the system alarm.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 or T4000 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis.</p> <p>all-lcc—(TX Matrix routers and TX Matrix Plus router only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all connected T1600 or T4000 LCCs.</p> <p>all-members—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.</p> <p>ancpd-service—Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.</p> <p>application-identification—Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.</p> <p>audit-process—(Optional) Display the RADIUS accounting process.</p> <p>auto-configuration—Display the Interface Auto-Configuration process.</p> <p>bootp—Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.</p> <p>captive-portal-content-delivery—Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.</p>

- ce-l2tp-service**—(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
- cfm**—Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.
- chassis-control**—(Optional) Display the chassis management process.
- class-of-service**—(Optional) Display the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
- clksyncd-service**—Display the external clock synchronization process, which uses synchronous Ethernet (SyncE).
- craft-control**—Display the process for the I/O of the craft interface.
- database-replication**—(EX Series switches and MX Series routers only) (Optional) Display the database replication process.
- datapath-trace-service**—Display the packet path tracing process.
- dhcp-service**—(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
- diameter-service**—(Optional) Display the diameter process.
- disk-monitoring**—(Optional) Display the disk monitoring process, which checks the health of the hard disk drive on the Routing Engine.
- dynamic-flow-capture**—(Optional) Display the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.
- ecc-error-logging**—(Optional) Display the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
- ethernet-connectivity-fault-management**—Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Display the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- event-processing**—(Optional) Display the event process (eventd).
- firewall**—(Optional) Display the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only)
(Optional) Display the general authentication process.

health (pid *process-identifier* | process-name *process-name*)—(Optional) Display process health information, either by process id (PID) or by process name.

iccp-service—Display the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—Display the intrusion detection and prevention (IDP) protocol process.

ilmi—Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.

inet-process—Display the IP multicast family process.

init—Display the process that initializes the USB modem.

interface-control—(Optional) Display the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

kernel-replication—(Optional) Display the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Display the Layer 2 address flooding and learning process.

l2cpd-service—Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

lACP—(Optional) Display the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the local Virtual Chassis member.

local-policy-decision-function—Display the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

logical-system-mux—Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

mac-validation—Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

mib-process—(Optional) Display the MIB II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Display the Mobile IP process, which configures Junos OS Mobile IP features.

moundd-service—(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.

mpls-traceroute—(Optional) Display the MPLS Periodic Traceroute process.

mspd—(Optional) Display the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Display the multicast snooping process, which makes Layer 2 devices such as VLAN switches aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Display the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

neighbor-liveness—Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.

nfsd-service—(Optional) Display the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

ntp—Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.

packet-triggered-subscribers—Display the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Display the Peer Selection Service process.

periodic-packet-services—Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

pfe—Display the Packet Forwarding Engine management process.

pgcp-service—(Optional) Display the pgcpd service process running on the Routing Engine.

pgm—Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.

pic-services-logging—(Optional) Display the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

ppp—(Optional) Display the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Display the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

process-monitor—Display the process health monitor process (pmond).

providers—(Optional) Display provider processes.

redundancy-interface-process—(Optional) Display the ASP redundancy process.

remote-operations—(Optional) Display the remote operations process, which provides the ping and traceroute MIBs.

resource-cleanup—Display the resource cleanup process.

resource-limits (brief | detail) process-name—(Optional) Display process resource limits.

routing—(Optional) Display the routing protocol process.

sampling—(Optional) Display the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbcc-configuration-process—Display the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).

sdk-service—Display the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

send—(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.

service-deployment—(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

sfc number—(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router. Replace *number* with 0.

snmp—Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

sonet-aps—Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.

static-subscribers—(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

tunnel-oamd—(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

vrrp—(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

watchdog—Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.

wide—(Optional) Display process information that might be wider than 80 columns.

Additional Information By default, when you issue the **show system processes** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise,

if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [List of Junos OS Processes](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

- [show system processes on page 992](#)
- [show system processes brief on page 992](#)
- [show system processes detail on page 993](#)
- [show system processes extensive on page 993](#)
- [show system processes lcc wide \(TX Matrix Routing Matrix\) on page 994](#)
- [show system processes summary on page 995](#)
- [show system processes \(TX Matrix Plus Router\) on page 995](#)
- [show system processes sfc \(TX Matrix Plus Router\) on page 1002](#)
- [show system processes lcc wide \(TX Matrix Plus Routing Matrix\) on page 1005](#)
- [show system processes \(QFX Series\) on page 1006](#)

Output Fields [Table 88 on page 990](#) describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

Table 88: show system processes Output Fields

Field Name	Field Description	Level of Output
last pid	Last process identifier assigned to the process.	brief extensive summary
load averages	Three load averages followed by the current time.	brief extensive summary
processes	Number of existing processes and the number of processes in each state (sleeping , running , starting , zombies , and stopped).	brief extensive summary
Mem	Information about physical and virtual memory allocation.	brief extensive summary
Swap	Information about physical and virtual memory allocation.	brief extensive summary
PID	Process identifier.	detail extensive summary
TT	Control terminal name.	none detail

Table 88: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
STAT	<p>Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process:</p> <ul style="list-style-type: none"> • D—In disk or other short-term, uninterruptible wait • I—Idle (sleeping longer than about 20 seconds) • R—Runnable • S—Sleeping for less than 20 seconds • T—Stopped • Z—Dead (zombie) • + —The process is in the foreground process group of its control terminal. • <—The process has raised CPU scheduling priority. • >—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped. • A—The process requested random page replacement. • E—The process is trying to exit. • L—The process has pages locked in core. • N—The process has reduced CPU scheduling priority. • S—The process requested first-in, first-out (FIFO) page replacement. • s—The process is a session leader. • V—The process is temporarily suspended. • W—The process is swapped out. • X—The process is being traced or debugged. 	none detail
UID	User identifier.	detail
USERNAME	Process owner.	extensive summary
PPID	Parent process identifier.	detail
CPU	<p>(D)—Short-term CPU usage.</p> <p>(E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.</p>	detail extensive summary
RSS	Resident set size.	detail
WCHAN	Symbolic name of the wait channel.	detail
STARTED	Local time when the process started running.	detail
PRI	Current priority of the process. A lower number indicates a higher priority.	detail extensive summary
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	detail extensive summary
SIZE	Total size of the process (text, data, and stack), in kilobytes.	extensive summary

Table 88: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
RES	Current amount of resident memory, in kilobytes.	extensive summary
STATE	Current state of the process (for example, sleep , wait , run , idle , zombie , or stop).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used. (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running.	detail extensive summary

Sample Output

show system processes

```

user@host> show system processes
PID  TT  STAT      TIME  COMMAND
  0  ??  DLs      0:00.70  (swapper)
  1  ??  Is       0:00.35  /sbin/init --
  2  ??  DL       0:00.00  (pagedaemon)
  3  ??  DL       0:00.00  (vmdaemon)
  4  ??  DL       0:42.37  (update)
  5  ??  DL       0:00.00  (if_jnx)
 80  ??  Ss       0:14.66  syslogd -s
 96  ??  Is       0:00.01  portmap
128  ??  Is       0:02.70  cron
173  ??  Is       0:02.24  /usr/local/sbin/sshd (sshd1)
189  ??  S        0:03.80  /sbin/watchdog -t180
190  ??  I        0:00.03  /usr/sbin/tnetd -N
191  ??  S        2:24.76  /sbin/ifd -N
192  ??  S<       0:55.44  /usr/sbin/xntpd -N
195  ??  S        0:53.11  /usr/sbin/snmpd -N
196  ??  S        1:15.73  /usr/sbin/mib2d -N
198  ??  I        0:00.75  /usr/sbin/inetd -N
2677 ??  I        0:00.01  /usr/sbin/mgd -N
2712 ??  Ss       0:00.24  rlogind
2735 ??  R        0:00.00  /bin/ps -ax
1985 p0-  S        0:07.41  ./rpd -N
2713 p0  Is       0:00.24  -tcsh (tcsh)
2726 p0  S+       0:00.07  cli

```

show system processes brief

```

user@host> show system processes brief
last pid:  543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

show system processes detail

```

user@host> show system processes detail
PID  UID  PPID CPU PRI NI  RSS WCHAN  STARTED  TT  STAT  TIME COMMAND
3151  1049  3129  2  28  0  672 -      1:13PM  p0  R+    0:00.00 ps -ax -r
   1    0    0  0  10  0  376 wait   1:51PM  ??  Is    0:00.29 /sbin/ini
   2    0    0  0 -18  0   12 psleep  1:51PM  ??  DL    0:00.00 (pagedae
   3    0    0  0  28  0   12 psleep  1:51PM  ??  DL    0:00.00 (vmdaemon
   4    0    0  0  28  0   12 update  1:51PM  ??  DL    0:07.15 (update)
   5    0    0  0   2  0   12 pfesel   1:51PM  ??  IL    0:02.90 (if_pfe)
  27    0    1  0  10  0 17936 mfsidl   1:51PM  ??  Is    0:00.46 mfs /dev/
  81    0    1  0   2  0   496 select  1:52PM  ??  Ss    0:31.21 syslogd -
 119    1    1  0   2  0   492 select  1:52PM  ??  Is    0:00.00 portmap
 134    0    1  0   2  0   580 select  1:52PM  ??  S     0:02.95 amd -p -a
 151    0    1  0  18  0   532 pause   1:52PM  ??  Is    0:00.34 cron
 183    0    1  0   2  0   420 select  1:52PM  ??  Ss    0:00.07 /usr/loca
 206    0    1  0  18  0    72 pause   1:52PM  ??  S     0:00.51 /sbin/wat
 207    0    1  0   2  0   520 select  1:52PM  ??  I     0:00.16 /usr/sbin
 208    0    1  0   2  0   536 select  1:52PM  ??  S     0:08.21 /sbin/dcd
 210    0    1 255  2 -12   740 select  1:52PM  ??  S<    0:05.83 /usr/sbin
 211    0    1  0   2  0   376 select  1:52PM  ??  S     0:00.03 /usr/sbin
 215    0    1  0   2  0   548 select  1:52PM  ??  I     0:00.50 /usr/sbin
 219    0    1  0   3  0   540 ttyin   1:52PM  v0  Is+   0:00.02 /usr/libe
 220    0    1  0   3  0   540 ttyin   1:52PM  v1  Is+   0:00.01 /usr/libe
 221    0    1  0   3  0   540 ttyin   1:52PM  v2  Is+   0:00.01 /usr/libe
 222    0    1  0   3  0   540 ttyin   1:52PM  v3  Is+   0:00.01 /usr/libe
 735    0    1  0   2  0   468 select  2:47PM  ??  S     0:19.14 /usr/sbin
 736    0    1  0   2  0   212 select  2:47PM  ??  S     0:14.13 /usr/sbin
1380    0    1  0   3  0   888 ttyin    7:32PM  d0  Is+   0:00.46 bash
3019    0   207  0   2  0   636 select 10:49AM  ??  Ss    0:02.93 tnp.chass
3122    0  1380  0   2  0  1764 select 12:33PM  d0  S     0:00.77 ./rpd -N
3128    0   215  0   2  0   580 select 12:45PM  ??  Ss    0:00.12 rlogind
3129  1049  3128  0  18  0   944 pause 12:45PM  p0  Ss    0:00.14 -tcsh (tc
   0    0    0  0 -18  0    0 sched  1:51PM  ??  DLs   0:00.10 (swapper

```

show system processes extensive

```

user@host> show system processes extensive
last pid: 544; load averages: 0.00, 0.00, 0.00 18:30:33
37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND
544 root 30 0 604K 768K RUN 0:00 0.00% 0.00% top
   3 root 28 0 0K 12K psleep 0:00 0.00% 0.00% vm Daemon
   4 root 28 0 0K 12K update 0:03 0.00% 0.00% update
 528 aviva 18 0 660K 948K pause 0:00 0.00% 0.00% tcsh
 204 root 18 0 300K 544K pause 0:00 0.00% 0.00% csh
 131 root 18 0 332K 532K pause 0:00 0.00% 0.00% cron
 186 root 18 0 196K 68K pause 0:00 0.00% 0.00% watchdog
 27 root 10 0 512M 16288K mfsidl 0:00 0.00% 0.00% mount_mfs
   1 root 10 0 620K 344K wait 0:00 0.00% 0.00% init
 304 root 3 0 884K 900K ttyin 0:00 0.00% 0.00% bash
 200 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 203 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 202 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 201 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 194 root 2 0 2248K 1640K select 0:11 0.00% 0.00% rpd
 205 root 2 0 964K 800K select 0:12 0.00% 0.00% tnp.chassisd
 189 root 2 -12 352K 740K select 0:03 0.00% 0.00% xntpd

```

```

114 root      2   0   296K   612K select  0:00  0.00%  0.00% amd
188 root      2   0   780K   600K select  0:00  0.00%  0.00% dcd
527 root      2   0   176K   580K select  0:00  0.00%  0.00% rlogind
195 root      2   0   212K   552K select  0:00  0.00%  0.00% inetd
187 root      2   0   192K   532K select  0:00  0.00%  0.00% tnetd
 83 root      2   0   188K   520K select  0:00  0.00%  0.00% syslogd
538 root      2   0  1324K   516K select  0:00  0.00%  0.00% mgd
 99 daemon    2   0   176K   492K select  0:00  0.00%  0.00% portmap
163 root      2   0   572K   420K select  0:00  0.00%  0.00% nsrexecd
192 root      2   0   560K   400K select  0:10  0.00%  0.00% snmpd
191 root      2   0  1284K   376K select  0:00  0.00%  0.00% mgd
537 aviva     2   0   636K   364K select  0:00  0.00%  0.00% cli
193 root      2   0   312K   204K select  0:07  0.00%  0.00% mib2d
  5 root      2   0      0K    12K pfesel  0:00  0.00%  0.00% if_pfe
  2 root     -18   0      0K    12K psleep  0:00  0.00%  0.00% pagedaemon
  0 root     -18   0      0K     0K sched   0:00  0.00%  0.00% swapper

```

show system processes lcc wide (TX Matrix Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

```

-----
PID TT  STAT      TIME COMMAND
  0 ??  DLs      0:00.00 (swapper)
  1 ??  ILs      0:00.10 /sbin/preinit -- (init)
  2 ??  DL       0:00.00 (pagedaemon)
  3 ??  DL       0:00.00 (vmdaemon)
  4 ??  DL       0:00.00 (bufdaemon)
  5 ??  DL       0:00.04 (syncer)
  6 ??  DL       0:00.00 (netdaemon)
  7 ??  IL       0:00.00 (if_pic_listen)
  8 ??  IL       0:00.00 (scs_housekeeping)
  9 ??  IL       0:00.00 (if_pfe_listen)
 10 ??  DL       0:00.00 (vmuncachedaemon)
 11 ??  SL       0:00.02 (cb_poll)
 172 ??  ILs      0:00.21 mfs -o noauto /dev/ad1s1b /tmp (newfs)
2909 ??  Is       0:00.00 pccardd
2932 ??  Ss       0:00.07 syslogd -r -s
3039 ??  Is       0:00.00 cron
3217 ??  I        0:00.00 /sbin/watchdog -d
3218 ??  I        0:00.02 /usr/sbin/tnetd -N
3221 ??  S        0:00.11 /usr/sbin/alarmd -N
3222 ??  S        0:00.85 /usr/sbin/craftd -N
3223 ??  S        0:00.05 /usr/sbin/mgd -N
3224 ??  I        0:00.02 /usr/sbin/inetd -N
3225 ??  I        0:00.00 /usr/sbin/tnp.sntpd -N
3226 ??  I        0:00.01 /usr/sbin/tnp.sntpc -N
3228 ??  I        0:00.01 /usr/sbin/smartd -N
3231 ??  I        0:00.01 /usr/sbin/eccd -N
3425 ??  S        0:00.09 /usr/sbin/dfwd -N
3426 ??  S        0:00.19 /sbin/dcd -N
3427 ??  I        0:00.04 /usr/sbin/pfed -N
3430 ??  S        0:00.10 /usr/sbin/ksyncd -N
3482 ??  S        1:53.63 /usr/sbin/chassisd -N
4285 ??  SL       0:00.01 (peer proxy)
4286 ??  SL       0:00.00 (peer proxy)
4303 ??  Ss       0:00.00 mgd: (mgd) (root) (mgd)
4304 ??  R        0:00.00 /bin/ps -ax -ww
3270 d0  Is+      0:00.00 /usr/libexec/getty std.9600 ttyd0

```

show system processes summary

```
user@host> show system processes summary
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping
```

```
Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
527	root	2	0	176K	580K	select	0:00	0.04%	0.04%	rlogind
543	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top

show system processes (TX Matrix Plus Router)

```
user@host> show system processes
sfc0-re0:
```

```
-----
PID TT STAT TIME COMMAND
0 ?? Wls 0:00.00 [swapper]
1 ?? ILs 0:00.18 /packages/mnt/jbase/sbin/init --
2 ?? DL 0:00.20 [g_event]
3 ?? DL 0:00.39 [g_up]
4 ?? DL 0:00.32 [g_down]
5 ?? DL 0:00.00 [thread taskq]
6 ?? DL 0:00.09 [kqueue taskq]
7 ?? DL 0:00.01 [pagedaemon]
8 ?? DL 0:00.00 [vmdaemon]
9 ?? DL 0:06.63 [pagezero]
10 ?? DL 0:00.00 [ktrace]
11 ?? RL 310:52.98 [idle]
12 ?? WL 0:11.03 [swi2: net]
13 ?? WL 0:27.58 [swi7: clock sio]
14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:03.02 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:11.41 [irq16: uhci0 uhci*]
21 ?? DL 0:00.00 [usb0]
22 ?? DL 0:00.00 [usbtask]
23 ?? WL 0:39.51 [irq17: uhci1 uhci*]
24 ?? DL 0:00.00 [usb1]
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.83 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
```

```

42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.34 [bcmTX]
1342 ?? SL 0:01.68 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.40 [bcmLINK.0]
1345 ?? SL 0:33.83 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? S 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.19 /usr/sbin/mgd -N
1512 ?? I 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.10 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd
23926 ?? Ss 0:00.01 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:25.07 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.38 /usr/sbin/l2ald -N
45730 ?? S< 0:00.12 /usr/sbin/apsd -N

```



```

45731 ?? SN      0:00.10 /usr/sbin/sampled -N
45732 ?? S       0:00.03 /usr/sbin/ilmid -N
45733 ?? S       0:00.09 /usr/sbin/rmopd -N
45734 ?? S       0:00.30 /usr/sbin/cosd
45735 ?? I       0:00.00 /usr/sbin/rtspd -N
45736 ?? S       0:00.06 /usr/sbin/fsad -N
45737 ?? S       0:00.05 /usr/sbin/rdd -N
45738 ?? S       0:00.10 /usr/sbin/pppd -N
45739 ?? S       0:00.05 /usr/sbin/dfcd -N
45740 ?? S       0:00.07 /usr/sbin/lfmd -N
45741 ?? S       0:00.01 /usr/sbin/implsoamd -N
45742 ?? I       0:00.01 /usr/sbin/sendd -N
45743 ?? S       0:00.08 /usr/sbin/appidd -N
45744 ?? S       0:00.05 /usr/sbin/mspd -N
45745 ?? S       0:00.25 /usr/sbin/jdiameterd -N
45746 ?? S       0:00.10 /usr/sbin/pfed -N
45747 ?? S       0:00.19 /usr/sbin/lpdfd -N
45748 ?? S       0:00.63 /sbin/dcd -N
45750 ?? S       0:00.45 /usr/sbin/mib2d -N
45751 ?? S       0:00.15 /usr/sbin/dfwd -N
45752 ?? S       0:00.15 /usr/sbin/irsd -N
45764 ?? S       0:20.59 /usr/sbin/snmpd -N
56479 ?? Ss      0:00.00 mgd: (mgd) (root) (mgd)
56480 ?? R       0:00.00 /bin/ps -ax
1142 d0- I       0:00.01 /usr/sbin/usbd -N
1160 d0- S       0:29.17 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+     0:00.00 /usr/libexec/getty std.9600 ttyd0
2392 p1 Is      0:00.00 login [pam] (login)
2393 p1 I        0:00.00 -csh (csh)
2394 p1 I        0:00.00 su -
2395 p1 I+       0:00.01 -su (csh)
23782 p2 Is      0:00.00 login [pam] (login)
23881 p2 I        0:00.00 -csh (csh)
23925 p2 S+      0:00.03 cli
7332 p3 Is      0:00.00 login [pam] (login)
7333 p3 I        0:00.00 -csh (csh)
23780 p3 S+      0:00.02 telnet aj

```

lcc0-re0:

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  WLS      0:00.00 [swapper]
  1  ??  ILs      0:00.16 /packages/mnt/jbase/sbin/init --
  2  ??  DL       0:00.01 [g_event]
  3  ??  DL       0:00.16 [g_up]
  4  ??  DL       0:00.11 [g_down]
  5  ??  DL       0:00.00 [thread taskq]
  6  ??  DL       0:00.00 [kqueue taskq]
  7  ??  DL       0:00.00 [pagedaemon]
  8  ??  DL       0:00.00 [vmdaemon]
  9  ??  DL       0:01.77 [pagezero]
 10  ??  DL       0:00.00 [ktrace]
 11  ??  RL      17:22.31 [idle]
 12  ??  WL       0:00.32 [swi2: net]
 13  ??  WL       0:01.21 [swi7: clock sio]
 14  ??  WL       0:00.00 [swi6: vm]
 15  ??  DL       0:00.10 [yarrow]
 16  ??  WL       0:00.00 [swi9: +]
 17  ??  WL       0:00.00 [swi8: +]
 18  ??  WL       0:00.00 [swi5: cambio]
 19  ??  WL       0:00.00 [swi9: task queue]

```

```

20 ?? WL 0:02.73 [irq10: bcm0 uhci1*]
21 ?? WL 0:00.02 [irq11: cb0 uhci0+*]
22 ?? DL 0:00.00 [usb0]
23 ?? DL 0:00.00 [usbtask]
24 ?? DL 0:00.00 [usb1]
25 ?? DL 0:00.05 [usb2]
26 ?? DL 0:00.00 [usb3]
27 ?? DL 0:00.00 [usb4]
28 ?? DL 0:00.00 [usb5]
29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vnlru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.01 [schedcpu]
55 ?? DL 0:00.73 [md0]
77 ?? DL 0:03.54 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1078 ?? DL 0:00.00 [jsr_kkcm]
1363 ?? SL 0:00.09 [bcmTX]
1364 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1365 ?? SL 0:03.08 [bcmLINK.0]
1370 ?? Is 0:00.00 /usr/sbin/cron
1522 ?? S 0:00.00 /sbin/watchdog -t-1
1523 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1524 ?? I 0:00.01 /usr/sbin/tinetd -N
1526 ?? S 0:04.98 /usr/sbin/chassisd -N
1527 ?? S 0:00.04 /usr/sbin/alarmd -N
1528 ?? I 0:00.40 /usr/sbin/craftd -N
1529 ?? S 0:00.08 /usr/sbin/mgd -N
1532 ?? I 0:00.04 /usr/sbin/inetd -N
1533 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1534 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1536 ?? S 0:00.01 /usr/sbin/smartd -N
1540 ?? I 0:00.07 /usr/sbin/jcsd -N

```

```

1541 ?? S      0:00.11 /usr/sbin/idpd -N
1542 ?? I      0:00.00 /usr/libexec/getty Pc ttyv0
2089 ?? DL     0:00.01 [peer proxy]
2090 ?? DL     0:00.01 [peer proxy]
2091 ?? DL     0:00.01 [peer proxy]
2657 ?? S      0:00.02 /usr/sbin/dfwd -N
2658 ?? S      0:00.02 /sbin/dcd -N
2659 ?? S      0:00.05 /usr/sbin/snmpd -N
2660 ?? S      0:00.01 /usr/sbin/mib2d -N
2661 ?? S      0:00.01 /usr/sbin/pfed -N
2662 ?? S      0:00.01 /usr/sbin/irsd -N
2667 ?? S      0:00.13 /usr/sbin/ksyncd -N
2690 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
2691 ?? R      0:00.00 /bin/ps -ax
1164 d0- S     0:00.00 /usr/sbin/usbd -N
1182 d0- S     0:00.34 /usr/sbin/eventd -N -r -s -A
1543 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

lcc1-re0:

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  Wls      0:00.00 [swapper]
  1  ??  ILs      0:00.17 /packages/mnt/jbase/sbin/init --
  2  ??  DL       0:00.01 [g_event]
  3  ??  DL       0:00.16 [g_up]
  4  ??  DL       0:00.11 [g_down]
  5  ??  DL       0:00.00 [thread taskq]
  6  ??  DL       0:00.00 [kqueue taskq]
  7  ??  DL       0:00.00 [pagedaemon]
  8  ??  DL       0:00.00 [vmdaemon]
  9  ??  DL       0:01.77 [pagezero]
 10  ??  DL       0:00.00 [ktrace]
 11  ??  RL      17:22.83 [idle]
 12  ??  WL       0:00.35 [swi2: net]
 13  ??  WL       0:01.20 [swi7: clock sio]
 14  ??  WL       0:00.00 [swi6: vm]
 15  ??  DL       0:00.10 [yarrow]
 16  ??  WL       0:00.00 [swi9: +]
 17  ??  WL       0:00.00 [swi8: +]
 18  ??  WL       0:00.00 [swi5: cambio]
 19  ??  WL       0:00.00 [swi9: task queue]
 20  ??  WL       0:02.87 [irq10: bcm0 uhci1*]
 21  ??  WL       0:00.02 [irq11: cb0 uhci0+*]
 22  ??  DL       0:00.00 [usb0]
 23  ??  DL       0:00.00 [usbtask]
 24  ??  DL       0:00.00 [usb1]
 25  ??  DL       0:00.05 [usb2]
 26  ??  DL       0:00.00 [usb3]
 27  ??  DL       0:00.00 [usb4]
 28  ??  DL       0:00.00 [usb5]
 29  ??  DL       0:00.04 [usb6]
 30  ??  DL       0:00.00 [usb7]
 31  ??  WL       0:00.00 [irq14: ata0]
 32  ??  WL       0:00.00 [irq15: ata1]
 33  ??  WL       0:00.00 [irq1: atkbd0]
 34  ??  WL       0:00.00 [swi0: sio]
 35  ??  WL       0:00.00 [swi3: ip6opt ipopt]
 36  ??  WL       0:00.00 [swi4: ip6mismatch+]
 37  ??  WL       0:00.00 [swi1: ipfwd]
 38  ??  DL       0:00.00 [bufdaemon]
 39  ??  DL       0:00.00 [vn1ru]

```

```

40 ?? DL      0:00.01 [syncer]
41 ?? DL      0:00.00 [softdepflush]
42 ?? DL      0:00.00 [netdaemon]
43 ?? DL      0:00.00 [vmuncachedaemon]
44 ?? DL      0:00.00 [if_pic_listen]
45 ?? DL      0:00.02 [vmkmemdaemon]
46 ?? DL      0:00.01 [cb_poll]
47 ?? DL      0:00.00 [if_pfe_listen]
48 ?? DL      0:00.00 [scs_housekeeping]
49 ?? IL      0:00.00 [kern_dump_proc]
50 ?? IL      0:00.00 [nfsiod 0]
51 ?? IL      0:00.00 [nfsiod 1]
52 ?? IL      0:00.00 [nfsiod 2]
53 ?? IL      0:00.00 [nfsiod 3]
54 ?? DL      0:00.02 [schedcpu]
55 ?? DL      0:00.75 [md0]
77 ?? DL      0:03.40 [md1]
98 ?? DL      0:00.37 [md2]
116 ?? DL     0:00.02 [md3]
137 ?? DL     0:00.56 [md4]
158 ?? DL     0:00.15 [md5]
179 ?? DL     0:00.00 [md6]
215 ?? DL     0:00.03 [md7]
225 ?? DL     0:00.03 [md8]
1052 ?? DL    0:00.00 [jsr_kkcm]
1337 ?? SL    0:00.09 [bcmTX]
1338 ?? SL    0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL    0:03.10 [bcmLINK.0]
1344 ?? Is    0:00.00 /usr/sbin/cron
1496 ?? S     0:00.00 /sbin/watchdog -t-1
1497 ?? S     0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? I     0:00.01 /usr/sbin/tnetd -N
1500 ?? S     0:04.97 /usr/sbin/chassisd -N
1501 ?? S     0:00.04 /usr/sbin/alarmd -N
1502 ?? I     0:00.40 /usr/sbin/craftd -N
1503 ?? S     0:00.08 /usr/sbin/mgd -N
1506 ?? I     0:00.04 /usr/sbin/inetd -N
1507 ?? I     0:00.00 /usr/sbin/tnp.snmpd -N
1508 ?? I     0:00.00 /usr/sbin/tnp.sntpd -N
1510 ?? S     0:00.01 /usr/sbin/smartd -N
1514 ?? I     0:00.07 /usr/sbin/jcsd -N
1515 ?? S     0:00.18 /usr/sbin/idpd -N
1516 ?? I     0:00.00 /usr/libexec/getty Pc ttyv0
2068 ?? DL    0:00.01 [peer proxy]
2069 ?? DL    0:00.01 [peer proxy]
2070 ?? DL    0:00.01 [peer proxy]
2666 ?? S     0:00.02 /sbin/dcd -N
2667 ?? S     0:00.01 /usr/sbin/irsd -N
2668 ?? S     0:00.01 /usr/sbin/pfed -N
2669 ?? S     0:00.05 /usr/sbin/snmpd -N
2670 ?? S     0:00.01 /usr/sbin/mib2d -N
2671 ?? S     0:00.02 /usr/sbin/dfwd -N
2675 ?? S     0:00.13 /usr/sbin/ksyncd -N
2699 ?? Ss    0:00.00 mgd: (mgd) (root) (mgd)
2700 ?? R     0:00.00 /bin/ps -ax
1138 d0- S     0:00.00 /usr/sbin/usbd -N
1156 d0- S     0:00.37 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+   0:00.00 /usr/libexec/getty std.9600 ttyd0

```

```
lcc2-re0:
```

PID	TT	STAT	TIME	COMMAND
0	??	WLs	0:00.00	[swapper]
1	??	ILs	0:00.18	/packages/mnt/jbase/sbin/init --
2	??	DL	0:00.01	[g_event]
3	??	DL	0:00.17	[g_up]
4	??	DL	0:00.12	[g_down]
5	??	DL	0:00.00	[thread taskq]
6	??	DL	0:00.00	[kqueue taskq]
7	??	DL	0:00.00	[pagedaemon]
8	??	DL	0:00.00	[vmdaemon]
9	??	DL	0:01.77	[pagezero]
10	??	DL	0:00.00	[ktrace]
11	??	RL	17:19.13	[idle]
12	??	WL	0:00.36	[swi2: net]
13	??	WL	0:01.20	[swi7: clock sio]
14	??	WL	0:00.00	[swi6: vm]
15	??	DL	0:00.13	[yarrow]
16	??	WL	0:00.00	[swi9: +]
17	??	WL	0:00.00	[swi8: +]
18	??	WL	0:00.00	[swi5: cambio]
19	??	WL	0:00.00	[swi9: task queue]
20	??	WL	0:03.03	[irq10: bcm0 uhci1*]
21	??	WL	0:00.02	[irq11: cb0 uhci0+*]
22	??	DL	0:00.00	[usb0]
23	??	DL	0:00.00	[usbtask]
24	??	DL	0:00.00	[usb1]
25	??	DL	0:00.05	[usb2]
26	??	DL	0:00.00	[usb3]
27	??	DL	0:00.00	[usb4]
28	??	DL	0:00.00	[usb5]
29	??	DL	0:00.04	[usb6]
30	??	DL	0:00.00	[usb7]
31	??	WL	0:00.00	[irq14: ata0]
32	??	WL	0:00.00	[irq15: ata1]
33	??	WL	0:00.00	[irq1: atkbd0]
34	??	WL	0:00.00	[swi0: sio]
35	??	WL	0:00.00	[swi3: ip6opt ipopt]
36	??	WL	0:00.00	[swi4: ip6mismatch+]
37	??	WL	0:00.00	[swi1: ipfwd]
38	??	DL	0:00.00	[bufdaemon]
39	??	DL	0:00.00	[vn1ru]
40	??	DL	0:00.01	[syncer]
41	??	DL	0:00.00	[softdepflush]
42	??	DL	0:00.00	[netdaemon]
43	??	DL	0:00.00	[vmuncachedaemon]
44	??	DL	0:00.00	[if_pic_listen]
45	??	DL	0:00.02	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.48	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.56	[md4]

```

158 ?? DL      0:00.15 [md5]
179 ?? DL      0:00.00 [md6]
215 ?? DL      0:00.03 [md7]
225 ?? DL      0:00.03 [md8]
1052 ?? DL     0:00.00 [jsr_kkcm]
1337 ?? SL     0:00.09 [bcmTX]
1338 ?? SL     0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL     0:03.22 [bcmLINK.0]
1344 ?? Is     0:00.00 /usr/sbin/cron
1496 ?? S      0:00.00 /sbin/watchdog -t-1
1497 ?? S      0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? S      0:00.01 /usr/sbin/tnetd -N
1500 ?? R      0:05.17 /usr/sbin/chassisd -N
1501 ?? S      0:00.04 /usr/sbin/alarmd -N
1502 ?? I      0:00.39 /usr/sbin/craftd -N
1503 ?? S      0:00.08 /usr/sbin/mgd -N
1506 ?? I      0:00.05 /usr/sbin/inetd -N
1507 ?? I      0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I      0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S      0:00.01 /usr/sbin/smardd -N
1514 ?? I      0:00.07 /usr/sbin/jcsd -N
1515 ?? S      0:00.17 /usr/sbin/idpd -N
1516 ?? I      0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL     0:00.01 [peer proxy]
2592 ?? DL     0:00.01 [peer proxy]
2593 ?? DL     0:00.01 [peer proxy]
2597 ?? DL     0:00.00 [peer proxy]
3192 ?? S      0:00.01 /usr/sbin/irsd -N
3193 ?? S      0:00.05 /usr/sbin/snmpd -N
3194 ?? S      0:00.02 /sbin/dcd -N
3195 ?? S      0:00.01 /usr/sbin/pfed -N
3196 ?? S      0:00.01 /usr/sbin/mib2d -N
3197 ?? S      0:00.02 /usr/sbin/dfwd -N
3198 ?? S      0:00.13 /usr/sbin/ksyncd -N
3228 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
3229 ?? R      0:00.00 /bin/ps -ax
1138 d0- S     0:00.00 /usr/sbin/usbd -N
1156 d0- S     0:00.42 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
...

```

show system processes sfc (TX Matrix Plus Router)

```

user@host> show system processes sfc 0
sfc0-re0:

```

```

-----
PID  TT  STAT      TIME COMMAND
 0  ??  Ws      0:00.00 [swapper]
 1  ??  SLs     0:00.18 /packages/mnt/jbase/sbin/init --
 2  ??  DL      0:00.20 [g_event]
 3  ??  DL      0:00.39 [g_up]
 4  ??  DL      0:00.32 [g_down]
 5  ??  DL      0:00.00 [thread taskq]
 6  ??  DL      0:00.09 [kqueue taskq]
 7  ??  DL      0:00.01 [pagedaemon]
 8  ??  DL      0:00.00 [vmdaemon]
 9  ??  DL      0:06.63 [pagezero]
10  ??  DL      0:00.00 [ktrace]
11  ??  RL      312:09.00 [idle]
12  ??  WL      0:11.07 [swi2: net]
13  ??  WL      0:27.70 [swi7: c'clock sio]

```

```

14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:03.03 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:11.46 [irq16: uhci0 uhci*]
21 ?? DL 0:00.00 [usb0]
22 ?? DL 0:00.00 [usbtask]
23 ?? WL 0:39.63 [irq17: uhci1 uhci*]
24 ?? DL 0:00.00 [usb1]
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.84 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.35 [bcmTX]
1342 ?? SL 0:01.69 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.57 [bcmLINK.0]
1345 ?? SL 0:33.97 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? I 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N

```

```

1508 ?? S      0:14.54 /usr/sbin/craftd -N
1509 ?? S      0:01.20 /usr/sbin/mgd -N
1512 ?? S      0:00.05 /usr/sbin/inetd -N
1513 ?? S      0:00.10 /usr/sbin/tnp.snmpd -N
1517 ?? S      0:00.11 /usr/sbin/smardd -N
1525 ?? S      0:01.11 /usr/sbin/idpd -N
1526 ?? S      0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I      0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL     0:00.30 [peer proxy]
1617 ?? DL     0:00.32 [peer proxy]
1618 ?? DL     0:00.34 [peer proxy]
1619 ?? DL     0:00.30 [peer proxy]
2391 ?? Is     0:00.01 telnetd
7331 ?? Ss     0:00.03 telnetd
9538 ?? DL     0:01.16 [jsr_kkcm]
9613 ?? DL     0:00.18 [peer proxy]
23781 ?? Ss     0:00.01 telnetd
23926 ?? Ss     0:00.03 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S      0:03.14 /usr/sbin/rpd -N
36874 ?? S      0:00.08 /usr/sbin/lmpd
36876 ?? S      0:00.17 /usr/sbin/lacpd -N
36877 ?? S      0:00.15 /usr/sbin/bfdd -N
36878 ?? S      0:05.05 /usr/sbin/ppmd -N
36907 ?? S      0:26.63 /usr/sbin/chassisd -N
37775 ?? S      0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S      0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S      0:00.40 /usr/sbin/l2ald -N
45730 ?? S<     0:00.13 /usr/sbin/apd -N
45731 ?? SN     0:00.10 /usr/sbin/sampled -N
45732 ?? S      0:00.03 /usr/sbin/ilmid -N
45733 ?? S      0:00.09 /usr/sbin/rmopd -N
45734 ?? S      0:00.31 /usr/sbin/cosd
45735 ?? I      0:00.00 /usr/sbin/rtsdpd -N
45736 ?? S      0:00.06 /usr/sbin/fsad -N
45737 ?? S      0:00.05 /usr/sbin/rdd -N
45738 ?? S      0:00.10 /usr/sbin/pppd -N
45739 ?? S      0:00.05 /usr/sbin/dfcd -N
45740 ?? S      0:00.08 /usr/sbin/lfmd -N
45741 ?? S      0:00.01 /usr/sbin/mpiioamd -N
45742 ?? I      0:00.01 /usr/sbin/sendd -N
45743 ?? S      0:00.08 /usr/sbin/appidd -N
45744 ?? S      0:00.05 /usr/sbin/mspd -N
45745 ?? S      0:00.27 /usr/sbin/jdiameterd -N
45746 ?? S      0:00.10 /usr/sbin/pfed -N
45747 ?? S      0:00.19 /usr/sbin/lpdfd -N
45748 ?? S      0:00.64 /sbin/dcd -N
45750 ?? S      0:00.46 /usr/sbin/mib2d -N
45751 ?? S      0:00.16 /usr/sbin/dfwd -N
45752 ?? S      0:00.15 /usr/sbin/irsd -N
45764 ?? S      0:20.60 /usr/sbin/snmpd -N
56481 ?? Ss     0:00.02 telnetd
56548 ?? Rs     0:00.19 mgd: (mgd) (regress)/dev/tty0 (mgd)
56577 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
56578 ?? R      0:00.00 /bin/ps -ax
1142 d0- S      0:00.01 /usr/sbin/usbd -N
1160 d0- S      0:29.71 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
56482 p0 Is     0:00.00 login [pam] (login)
56483 p0 S      0:00.01 -csh (csh)
56547 p0 S+     0:00.02 cli
2392 p1 Is     0:00.00 login [pam] (login)

```



```

2393 p1 I 0:00.00 -csh (csh)
2394 p1 I 0:00.00 su -
2395 p1 I+ 0:00.01 -su (csh)
23782 p2 Is 0:00.00 login [pam] (login)
23881 p2 I 0:00.00 -csh (csh)
23925 p2 S+ 0:00.03 cli
7332 p3 Is 0:00.00 login [pam] (login)
7333 p3 I 0:00.00 -csh (csh)
23780 p3 S+ 0:00.02 telnet aj

```

show system processes lcc wide (TX Matrix Plus Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

```

-----
PID  TT  STAT    TIME PROVIDER COMMAND
0   ??  WLS    0:00.00 (null)  [swapper]
1   ??  ILs    0:00.19          /packages/mnt/jbase/sbin/init --
2   ??  DL     0:00.02          [g_event]
3   ??  DL     0:00.19          [g_up]
4   ??  DL     0:00.13          [g_down]
5   ??  DL     0:00.00          [thread taskq]
6   ??  DL     0:00.00          [kqueue taskq]
7   ??  DL     0:00.00          [pagedaemon]
8   ??  DL     0:00.00          [vmdaemon]
9   ??  DL     0:01.77          [pagezero]
10  ??  DL     0:00.00          [ktrace]
11  ??  RL    20:33.81          [idle]
12  ??  WL     0:00.38          [swi2: net]
13  ??  WL     0:01.43          [swi7: clock sio]
14  ??  WL     0:00.00          [swi6: vm]
15  ??  DL     0:00.14          [yarrow]
16  ??  WL     0:00.00          [swi9: +]
17  ??  WL     0:00.00          [swi8: +]
18  ??  WL     0:00.00          [swi5: cambio]
19  ??  WL     0:00.00          [swi9: task queue]
20  ??  WL     0:03.18          [irq10: bcm0 uhci1*]
21  ??  WL     0:00.03          [irq11: cb0 uhci0+*]
22  ??  DL     0:00.00          [usb0]
23  ??  DL     0:00.00          [usbtask]
24  ??  DL     0:00.00          [usb1]
25  ??  DL     0:00.06          [usb2]
26  ??  DL     0:00.00          [usb3]
27  ??  DL     0:00.00          [usb4]
28  ??  DL     0:00.00          [usb5]
29  ??  DL     0:00.05          [usb6]
30  ??  DL     0:00.00          [usb7]
31  ??  WL     0:00.00          [irq14: ata0]
32  ??  WL     0:00.00          [irq15: ata1]
33  ??  WL     0:00.00          [irq1: atkbd0]
34  ??  WL     0:00.00          [swi0: sio]
35  ??  WL     0:00.00          [swi3: ip6opt ipopt]
36  ??  WL     0:00.00          [swi4: ip6mismatch+]
37  ??  WL     0:00.00          [swi1: ipfwd]
38  ??  DL     0:00.00          [bufdaemon]
39  ??  DL     0:00.00          [vn1ru]
40  ??  DL     0:00.02          [syncer]
41  ??  DL     0:00.01          [softdepflush]
42  ??  DL     0:00.00          [netdaemon]
43  ??  DL     0:00.00          [vmuncachedaemon]
44  ??  DL     0:00.00          [if_pic_listen]

```

45	??	DL	0:00.03	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.84	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.72	[md4]
158	??	DL	0:00.15	[md5]
179	??	DL	0:00.00	[md6]
215	??	DL	0:00.03	[md7]
225	??	DL	0:00.03	[md8]
1052	??	DL	0:00.00	[jsr_kkcm]
1337	??	SL	0:00.11	[bcmTX]
1338	??	SL	0:00.12	[bcmXGS3AsyncTX]
1339	??	SL	0:03.82	[bcmLINK.0]
1344	??	Is	0:00.00	/usr/sbin/cron
1496	??	I	0:00.00	/sbin/watchdog -t-1
1497	??	S	0:00.06	/usr/libexec/bslockd -mp -N
1498	??	I	0:00.01	/usr/sbin/tnetd -N
1500	??	S	0:09.93	/usr/sbin/chassisd -N
1501	??	S	0:00.05	/usr/sbin/alarmd -N
1502	??	I	0:00.39	/usr/sbin/craftd -N
1503	??	S	0:00.09	/usr/sbin/mgd -N
1506	??	I	0:00.05	/usr/sbin/inetd -N
1507	??	I	0:00.00	/usr/sbin/tnp.sntpd -N
1508	??	I	0:00.00	/usr/sbin/tnp.sntpc -N
1510	??	S	0:00.01	/usr/sbin/smartd -N
1514	??	I	0:00.07	/usr/sbin/jcsd -N
1515	??	S	0:00.17	/usr/sbin/idpd -N
1516	??	I	0:00.00	/usr/libexec/getty Pc ttyv0
2591	??	DL	0:00.01	[peer proxy]
2592	??	DL	0:00.01	[peer proxy]
2593	??	DL	0:00.01	[peer proxy]
2597	??	DL	0:00.01	[peer proxy]
3192	??	S	0:00.02	/usr/sbin/irsd -N
3193	??	S	0:00.05	/usr/sbin/snmpd -N
3194	??	S	0:00.04	/sbin/dcd -N
3195	??	I	0:00.01	/usr/sbin/pfed -N
3196	??	S	0:00.02	/usr/sbin/mib2d -N
3197	??	I	0:00.03	/usr/sbin/dfwd -N
3198	??	S	0:00.15	/usr/sbin/ksyncd -N
3559	??	Ss	0:00.00	mgd: (mgd) (root) (mgd)
3560	??	R	0:00.00	/bin/ps -ax -Jpww
1138	d0-	S	0:00.00	/usr/sbin/usbd -N
1156	d0-	S	0:00.50	/usr/sbin/eventd -N -r -s -A
1517	d0	Is+	0:00.00	/usr/libexec/getty std.9600 ttyd0

show system processes (QFX Series)

```

user@switch> show system processes
PID TT STAT TIME COMMAND
0 ?? Wls -2341043:-31.01 [swapper]
1 ?? SLs 0:01.34 /packages/mnt/jbase/sbin/init --

```

```

 2 ?? DL      2:48.31 [g_event]
 3 ?? DL      1:47.44 [g_up]
 4 ?? DL      1:37.82 [g_down]
 5 ?? DL      0:00.00 [kdm_tcp_poller]
 6 ?? DL      0:00.00 [thread taskq]
 7 ?? DL      0:04.86 [kqueue taskq]
 9 ?? DL      0:03.94 [pagedaemon]
10 ?? DL      0:00.00 [ktrace]
11 ?? RL      0:00.00 [idle: cpu31]
12 ?? RL      0:00.00 [idle: cpu30]
13 ?? RL      0:00.00 [idle: cpu29]
14 ?? RL      0:00.00 [idle: cpu28]
15 ?? RL      0:00.00 [idle: cpu27]
16 ?? RL      0:00.00 [idle: cpu26]
17 ?? RL      0:00.00 [idle: cpu25]
18 ?? RL      0:00.00 [idle: cpu24]
19 ?? RL      0:00.00 [idle: cpu23]
20 ?? RL      0:00.00 [idle: cpu22]
21 ?? RL      0:00.00 [idle: cpu21]
22 ?? RL      0:00.00 [idle: cpu20]
23 ?? RL      0:00.00 [idle: cpu19]
24 ?? RL      0:00.00 [idle: cpu18]
25 ?? RL      0:00.00 [idle: cpu17]
26 ?? RL      0:00.00 [idle: cpu16]
27 ?? RL      0:00.00 [idle: cpu15]
28 ?? RL      0:00.00 [idle: cpu14]
29 ?? RL      0:00.00 [idle: cpu13]
30 ?? RL      0:00.00 [idle: cpu12]
31 ?? RL      0:00.00 [idle: cpu11]
32 ?? RL      0:00.00 [idle: cpu10]
33 ?? RL      0:00.00 [idle: cpu9]
34 ?? RL      18184:07.25 [idle: cpu8]
35 ?? RL      0:00.00 [idle: cpu7]
36 ?? RL      17862:11.31 [idle: cpu6]
37 ?? RL      19343:45.16 [idle: cpu5]
38 ?? RL      5192:38.30 [idle: cpu4]
39 ?? RL      0:00.00 [idle: cpu3]
40 ?? RL      19278:02.24 [idle: cpu2]
41 ?? RL      19291:00.72 [idle: cpu1]
42 ?? RL      18910:31.21 [idle: cpu0]
43 ?? WL      19:03.74 [swi2: net]
44 ?? WL      261:43.82 [swi7: clock sio]
45 ?? WL      0:00.00 [swi6: vm]
46 ?? DL      2:18.57 [yarrow]
47 ?? WL      0:00.00 [swi9: +]
48 ?? WL      0:00.00 [swi8: +]
49 ?? WL      0:12.36 [swi5: cambio]
50 ?? WL      0:00.00 [swi9: task queue]
51 ?? WL      0:00.00 [swi0: sio]
52 ?? WL      0:32.40 [irq39: ehci0]
53 ?? DL      0:00.21 [usb0]
54 ?? DL      0:00.00 [usbtask]
55 ?? WL      0:00.00 [irq22: xlr_lbus0]
56 ?? WL      0:00.00 [irq38: xlr_lbus0]
57 ?? WL      0:00.00 [swi3: ip6opt ipopt]
58 ?? WL      0:00.00 [swi4: ip6mismatch+]
59 ?? WL      0:00.00 [swi1: ipfwd]
60 ?? DL      0:18.65 [pagezero]
61 ?? DL      0:18.59 [bufdaemon]
62 ?? DL      1:10.44 [vnlr_u_mem]
63 ?? DL      1:51.66 [syncer]

```

```

64 ?? DL      0:20.22 [vn1ru]
65 ?? DL      0:40.48 [softdepflush]
66 ?? DL      0:00.00 [netdaemon]
67 ?? DL      20:47.67 [vmkmemdaemon]
68 ?? DL      0:00.00 [if_pfe_listen]
69 ?? SL      0:02.80 [kdm_checkkcore]
70 ?? SL      0:03.34 [kdm_savekcore]
71 ?? SL      0:04.31 [kdm_livekcore]
72 ?? SL      0:06.14 [kdm_logger]
73 ?? SL      0:04.31 [kdm_kdb]
74 ?? SL      0:00.02 [devrt_kernel_thread]
75 ?? DL      0:21.54 [vmuncachedaemon]
76 ?? DL      0:00.00 [if_pic_listen0]
77 ?? SL      0:00.00 [nfsiod 0]
78 ?? SL      0:00.00 [nfsiod 1]
79 ?? SL      0:00.00 [nfsiod 2]
80 ?? SL      0:00.00 [nfsiod 3]
81 ?? WL      5:59.98 [irq13: +]
82 ?? RL      105:06.81 [pkt_sender: cpu0]
83 ?? DL      0:03.62 [md0]
95 ?? DL      0:37.04 [md1]
115 ?? DL     0:06.01 [md2]
135 ?? DL     0:00.75 [md3]
155 ?? DL     0:21.17 [md4]
175 ?? DL     0:01.90 [md5]
195 ?? DL     0:06.26 [md6]
231 ?? DL     0:00.01 [md7]
755 ?? Ss     0:04.17 /usr/sbin/cron
847 ?? S      0:00.10 /usr/sbin/tnetd -N
849 ?? S      0:06.82 /usr/sbin/mgd -N
850 ?? S      0:00.32 /usr/sbin/inetd -N
852 ?? S      1:05.34 /usr/sbin/dhcpd -N
853 ?? S      0:00.18 /usr/sbin/inetd -p /var/run/inetd_4.pid -N -JU __juni
855 ?? L      1181:02.21 /usr/sbin/dc-pfe -N (pafxpc)
857 ?? S      17:55.86 /usr/sbin/vccpd -N
896 ?? S      93:43.45 /usr/sbin/chassism -N
953 ?? S      0:02.89 /sbin/watchdog -t-1
954 ?? S      3:34.00 /sbin/dcd -N
955 ?? S      10:30.13 /usr/sbin/chassisd -N
956 ?? DL     0:00.21 [peer proxy]
957 ?? S      4:07.43 /usr/sbin/alarmd -N
958 ?? S      0:31.69 /usr/sbin/craftd -N
959 ?? S      0:55.16 /usr/sbin/mib2d -N
960 ?? S      3:40.64 /usr/sbin/rpd -N
961 ?? S      0:00.03 /usr/sbin/tnp.snmpd -N
962 ?? S      0:51.94 /usr/sbin/pfed -N
963 ?? S      0:47.31 /usr/sbin/rmopd -N
964 ?? S      0:33.65 /usr/sbin/cosd
965 ?? S      1:48.41 /usr/sbin/ppmd -N
966 ?? S      0:07.18 /usr/sbin/dfwd -N
967 ?? S      1:02.56 /usr/sbin/bfdd -N
968 ?? S      0:00.63 /usr/sbin/rdd -N
969 ?? S      0:40.61 /usr/sbin/dfcd -N
971 ?? S      0:07.81 /usr/sbin/bdbrepd -N
972 ?? S      0:00.28 /usr/sbin/sendd -N
973 ?? S      1:37.69 /usr/sbin/xntpd -j -N -g -JU __juniper_private4__ (nt
974 ?? S      5:56.28 /usr/sbin/snmpd -N -JU __juniper_private4__
975 ?? S      16:46.82 /usr/sbin/jdiameterd -N
976 ?? S      2:34.13 /usr/sbin/eswd -N
977 ?? S      1:03.05 /usr/sbin/sflowd -N
978 ?? S      0:22.30 /usr/sbin/fcd -N

```

```

979 ?? S      1:07.01 /usr/sbin/vccpdf -N
982 ?? S      0:25.25 /usr/sbin/mcsnoopd -N
983 ?? S      3:45.68 /usr/sbin/rpdf -N
1043 ?? S      0:37.87 /usr/sbin/lacpd -N
1048 ?? DL     0:01.29 [peer proxy]
1111 ?? WL     0:00.00 [swi2: FMNITHRD+]
1112 ?? DL     0:00.03 [peer proxy]
12816 ?? S     15:35.32 /usr/sbin/sfid -N
30893 ?? Ss    0:00.65 sshd: tlewis@tty0 (sshd)
30897 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty0 (mgd)
30905 ?? Ss    0:00.64 sshd: tlewis@tty1 (sshd)
30909 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty1 (mgd)
30910 ?? Ss    0:01.26 sshd: tcheng@tty2 (sshd)
30914 ?? Ss    0:00.80 mgd: (mgd) (tcheng)/dev/tty2 (mgd)
30937 ?? R      0:00.03 /bin/ps -ax
661 d0- S      0:21.24 /usr/sbin/eventd -N -r -s -A
860 d0 Ss+     0:00.07 /usr/libexec/getty std.9600 ttyd0
30896 p0 Ss+    0:00.55 -cli (cli)
30908 p1 Ss+    0:00.50 -cli (cli)
30913 p2 Ss+    0:00.85 -cli (cli)

```

show system reboot

List of Syntax	Syntax on page 1010 Syntax (EX Series Switches) on page 1010 Syntax (TX Matrix Router) on page 1010 Syntax (TX Matrix Plus Router) on page 1010 Syntax (MX Series Router) on page 1010 Syntax (QFX Series) on page 1010
Syntax	show system reboot <both-routing-engines>
Syntax (EX Series Switches)	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system reboot <all-chassis all-lcc lcc <i>number</i> scc> <both-routing-engines>
Syntax (TX Matrix Plus Router)	show system reboot <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <both-routing-engines>
Syntax (MX Series Router)	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system reboot <both-routing-engines> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display pending system reboots or halts.
Options	none —Display pending reboots or halts on the active Routing Engine. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix router, display halt or reboot request information for all the T1600 or T4000 routers in the chassis that are connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

all-lcc—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display system halt or reboot request information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for all connected T1600 or T4000 LCCs.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

infrastructure *name*—(QFabric systems only) (Optional) Display reboot request information on the fabric manager Routing Engines and fabric control Routing Engines.

interconnect-device *name*—(QFabric systems only) (Optional) Display reboot request information on the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display reboot request information on the Node group.

scc—(TX Matrix router only) (Optional) Display halt or reboot request information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display halt or reboot request information for the TX Matrix Plus router.

Additional Information By default, when you issue the **show system reboot** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level maintenance

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system reboot on page 1012](#)
[show system reboot all-lcc \(TX Matrix Router\) on page 1012](#)
[show system reboot sfc \(TX Matrix Plus Router\) on page 1012](#)
[show system reboot \(QFX3500 Switch\) on page 1012](#)

Sample Output

[show system reboot](#)

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

[show system reboot all-lcc \(TX Matrix Router\)](#)

```
user@host> show system reboot all-lcc
lcc0-re0:
```

```
-----
No shutdown/reboot scheduled.
```

```
lcc2-re0:
```

```
-----
No shutdown/reboot scheduled.
```

[show system reboot sfc \(TX Matrix Plus Router\)](#)

```
user@host> show system sfc 0
No shutdown/reboot scheduled.
```

[show system reboot \(QFX3500 Switch\)](#)

```
user@switch> show system reboot
No shutdown/reboot scheduled.
```


show system resource-cleanup processes

Syntax	show system resource-cleanup processes <detail> <pid <i>number</i> > <process-name <i>name</i> >
Release Information	Command introduced in Junos OS Release 9.3. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the list of processes that have been registered for resource cleanup services.
Options	<p>detail—(Optional) Display the list of processes that have been registered for resource cleanup services, along with the resources that have been requested for cleanup.</p> <p>pid <i>number</i>—(Optional) Display a process that has been registered for resource cleanup services by specifying the Process Identifier number.</p> <p>process-name <i>name</i>—(Optional) Display a process that has been registered for resource cleanup services by name of the process.</p>
Required Privilege Level	view
List of Sample Output	show system resource-cleanup processes on page 1013 show system resource-cleanup processes detail on page 1013
Output Fields	For a description of the output fields, see Table 89 on page 1013 . Output fields are listed in the approximate order in which they appear.

Table 89: show system resource-cleanup processes Output Fields

Field Name	Field Description
PID	Process ID, a number that identifies a process.
Process name	String that identifies the process.
Resources to clean	Resources that have been registered to be cleaned up.

Sample Output

show system resource-cleanup processes


```
user@host> show system resource-cleanup processes
PID      Process name      Resources to clean
420      jnx-exampled      GENCFG, SYSV shared memory
```

show system resource-cleanup processes detail

```
user@host> show system resource-cleanup processes detail
PID      Process name      Resources to clean
420      jnx-exampled      GENCFG blob major ID 0x8000, minor ID 0x0000
```

SYSV shared memory ID 65536, key 1108955839
SYSV shared memory ID 65537, key 1108955837

show system rollback

Syntax	<code>show system rollback <i>number</i></code> <code><compare <i>number</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a previously committed configuration, or the differences between two previously committed configurations.
<div>  NOTE: The <code>show system rollback</code> command is a purely operational mode command and cannot be issued with <code>run</code> from the configuration mode. </div>	
Options	<p><i>number</i>—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.</p> <p><code>compare <i>number</i></code>—(Optional) Number of another previously committed (rollback) configuration to compare to rollback <i>number</i>. The output displays the differences between the two configurations. The range of values is 0 through 49.</p>
Required Privilege Level	view
List of Sample Output	show system rollback compare on page 1015

Sample Output

show system rollback compare

```

user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
+       }
+     }
+   }
+ }

```

```

+      }
+    }
+    ge-1/3/0 {
+      unit 0 {
+        family inet {
+          filter {
+            input mf_plp;
+          }
+          address 12.1.1.1/30;
+        }
+      }
+    }
+  }
+}

```

show system services service-deployment

Syntax	show system services service-deployment
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about a Session and Resource Control (SRC) client.
Options	This command has no options.
Required Privilege Level	system view
List of Sample Output	show system services service-deployment on page 1017
Output Fields	Table 90 on page 1017 lists the output fields for the show system services service-deployment command. Output fields are listed in the approximate order in which they appear.

Table 90: show system services service-deployment Output Fields

Field Name	Field Description
PDT Keepalive settings	Configured PDT keepalive interval, in seconds.
Keepalives sent	Number of keepalives sent.
Notifications sent	Number of notifications sent.
Last update from peer	Time at which the last update from a peer was received.

Sample Output

show system services service-deployment

```

user@host> show system services service-deployment
Connected to 192.4.4.4 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings:
Interval 15 seconds Keepalives sent: 750 Notifications sent: 0 Last update from
peer: 00:00:06 ago

```

show system software

List of Syntax	Syntax on page 1018 Syntax (EX Series Switches) on page 1018 Syntax (TX Matrix Router) on page 1018 Syntax (TX Matrix Plus Router) on page 1018 Syntax (J Series Routers) on page 1018 Syntax (QFX Series) on page 1018
Syntax	show system software <detail>
Syntax (EX Series Switches)	show system software <all-members> <detail> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system software <all-chassis all-lcc lcc <i>number</i> scc> <detail>
Syntax (TX Matrix Plus Router)	show system software <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <detail>
Syntax (J Series Routers)	show system software <backup> <detail>
Syntax (QFX Series)	show system software <detail> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the Junos OS extensions loaded on your router or switch.
Options	none —Display standard information about all loaded Junos OS extensions. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system software information for all the T640 routers (TX Matrix Router) or all the routers (TX Matrix Plus Router) in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for all T640 routers connected to the

TX Matrix router. On a TX Matrix Plus router, display system software information for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.

backup—(J Series routers only) (Optional) Display the status of old system software packages only.

detail—(Optional) Display detailed information about available Junos OS extensions.

infrastructure name—(QFabric systems only) (Optional) Display the system software running on the fabric control Routing Engine and the fabric manager Routing Engine.

interconnect-device name—(QFabric systems only) (Optional) Display the system software running on the Interconnect device.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

member member-id—(EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

node-group name—(QFabric systems only) (Optional) Display the system software running on the Node group.

scc—(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display system software information for the TX Matrix Plus router.

Required Privilege Level maintenance

Related Documentation	<ul style="list-style-type: none">• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	show system software on page 1020 show system software (TX Matrix Plus Router) on page 1020 show system software (QFX Series) on page 1024
Output Fields	When you enter this command, you are provided a list of Junos OS packages installed on the router and their corresponding Junos OS release number.

Sample Output

[show system software](#)

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]

Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]

Information for junos:

Comment:
JUNOS Base OS boot [7.2R1.7]
```

[show system software \(TX Matrix Plus Router\)](#)

```
user@host> show system software
sfc0-re0:
-----
Information for jbase:
```


Comment:
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:
JUNOS Online Documentation [9.6-20090515.0]
Information for jkernel:

Comment:
JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:Comment:
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aac1:

Comment:
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:

JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:

JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:

JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

...

lcc0-re0:

Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [9.6-20090515.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:

Comment:
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:
JUNOS Services AAACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:
JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:
JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

lcc1-re0:

Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

...

show system software (QFX Series)

user@switch> **show system software**

Information for jbase:

Comment:

JUNOS Base OS Software Suite [11.3-20110730.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [11.3-20110730.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [11.3-20110730.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [11.3-20110730.0]

Information for jpfe:

Comment:

JUNOS Packet Forwarding Engine Support (QFX) [11.3-20110730.0]

Information for jroute:

Comment:

JUNOS Routing Software Suite [11.3-20110730.0]

Information for jswitch:

Comment:

JUNOS Enterprise Software Suite [11.3-20110730.0]

Information for junos:

Comment:

JUNOS Base OS boot [11.3-20110730.0]

Information for jweb:

Comment:

JUNOS Web Management [11.3-20110730.0]

show system statistics

List of Syntax	Syntax on page 1026 Syntax (EX Series Switches) on page 1026 Syntax (TX Matrix Router) on page 1026 Syntax (TX Matrix Plus Router) on page 1026 Syntax (MX Series Router) on page 1026 Syntax (QFX Series) on page 1026
Syntax	show system statistics
Syntax (EX Series Switches)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system statistics
Release Information	Command introduced before JUNOS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display system-wide protocol-related statistics.
Options	none —Display system statistics for all the following protocols: <ul style="list-style-type: none">• arp—Address Resolution Protocol• bridge—IEEE 802.1 Bridging• clns—Connectionless Network Service• esis—End System-to-Intermediate System• ethoamcfm—Ethernet OAM protocol for connectivity fault management• ethoamlfm—Ethernet OAM protocol for link fault management• icmp—Internet Control Message Protocol• icmp6—Internet Control Message Protocol version 6• igmp—Internet Group Management Protocol

- **ip**—Internet Protocol version 4
- **ip6**—Internet Protocol version 6
- **mpls**—Multiprotocol Label Switching
- **rdp**—Reliable Datagram Protocol
- **tcp**—Transmission Control Protocol
- **tnp**—Trivial Network Protocol
- **ttp**—TNP Tunneling Protocol
- **tudp**—Trivial User Datagram Protocol
- **udp**—User Datagram Protocol
- **vpls**—Virtual Private LAN Service

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for all the routers in the chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for all routers (line-card chassis) connected to the TX Matrix Plus router

all-members—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for all members of the Virtual Chassis configuration.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the local Virtual Chassis member.

member member-id—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the specified member of the Virtual Chassis

configuration. For EX4200 switches, replace **member-id** with a value from 0 through 9. For an MX Series Virtual Chassis, replace **member-id** with a value of 0 or 1.

scc—(TX Matrix routers only) (Optional) Display system statistics for a protocol for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with 0.

Additional Information By default, when you issue the **show system statistics** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show system statistics on page 1028](#)
[show system statistics \(EX Series Switches\) on page 1035](#)
[show system statistics \(TX Matrix Router\) on page 1044](#)
[show system statistics \(QFX Series\) on page 1051](#)

Sample Output

show system statistics

```
user@host> show system statistics
ip:
    3682087 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped (queue overflow)
    0 fragments dropped after timeout
    0 fragments dropped due to over limit
    0 packets reassembled ok
    3664774 packets for this host
    17316 packets for unknown/unsupported protocol
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    6528 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs
    0 output packets discarded due to no route
```



```

0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
1123 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
1123 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
    echo reply: 75
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 75
    router advertisement: 130
75 message responses generated
tcp:
3844 packets sent
    3618 data packets (1055596 bytes)
    0 data packets (0 bytes) retransmitted
    0 resends initiated by MTU discovery
    205 ack-only packets (148 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    1079 control packets
5815 packets received
    3377 acks (for 1055657 bytes)
    24 duplicate acks
    0 acks for unsent data
    2655 packets (15004 bytes) received in-sequence
    1 completely duplicate packet (0 bytes)
    0 old duplicate packets
    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts

```

```
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
    27 connections updated cached RTT on close
    27 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
344 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    32 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
1058 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors

udp:
3658884 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
3657342 dropped due to no socket
3657342 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
4291311496 delivered
1551 datagrams output

ipsec:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
```

```

0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route

igmp:
17186 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

arp:
44181302 datagrams received
2 ARP requests received
2028 ARP replies received
3156 resolution requests received
0 unrestricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 with bogus interface
787 with incorrect length
712 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
7611 with multicast target address
0 with my own hardware address
14241699 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29929250 which were not for me
0 packets discarded waiting for resolution
6 packets sent after waiting for resolution
17812 ARP requests sent
2 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry

ip6:
0 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable

```

```
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
Mbuf statistics:
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too may headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol

icmp6:
0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Histogram of error messages to be generated:
    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

ipsec6:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
```

```

0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route
c1n1:
0 total packets received
0 packets delivered
0 too small
0 bad header length
0 bad checksum
0 bad version
0 unknown or unsupported protocol
0 bogus sdl size
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 address fields were not reasonable
0 segment information forgotten
0 forwarded packets
0 total packets sent
0 output packets discarded
0 non-forwarded packets
0 packets fragmented
0 fragments sent
0 fragments discarded
0 fragments timed out
0 fragmentation prohibited
0 packets reconstructed
0 packets destined to dead nexthop
0 packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 total pkts received
0 total packets consumed by protocol
0 pdus received with bad checksum
0 pdus received with bad version number
0 pdus received with bad type field
0 short pdus received
0 bogus sdl size
0 bad header length
0 unknown or unsupported protocol
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 ISO family not configured
tnp:
146776365 unicast packets received
0 broadcast packets received
0 fragmented packets received
0 hello packets dropped
0 fragments dropped
0 fragment reassembly queue flushes
0 hello packets received
0 control packets received
49681642 rdp packets received
337175 udp packets received
96757548 tunnel packets received
0 input packets discarded with no protocol
98397591 unicast packets sent

```

```
0 broadcast packets sent
0 fragmented packets sent
0 hello packets dropped
0 fragments dropped
0 hello packets sent
0 control packets sent
49681642 rdp packets sent
337175 udp packets sent
48378774 tunnel packets sent
0 packets sent with unknown protocol

rdp:
49681642 input packets
0 discards for bad checksum
0 discards bad sequence number
0 refused connections
2031964 acks received
0 dropped due to full socket buffers
49692 retransmits
49681642 output packets
24815968 acks sent
28 connects
0 closes
22783990 keepalives received
22783990 keepalives sent

tudp:
337175 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
0 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
337175 delivered
337175 datagrams output

ttp:
398749 packets sent
0 packets sent while unconnected
0 packets sent while interface down
0 packets sent couldn't get buffer
0 packets sent couldn't find neighbor
44696687 L2 packets received
0 unknown L3 packets received
3682087 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 cyclotron cycle L3 packets received
0 cyclotron send L3 packets received
0 packets received while unconnected
0 packets received from unknown ifl
0 input packets couldn't get buffer
0 input packets with bad type
0 input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
```

```

68877 Input packets dropped based on tlv result
0 input packets for which rt lookup is bypassed

mpls:
0 total mpls packets received
0 packets forwarded
0 packets dropped
0 with header too small
0 after tagging, can't fit link MTU
0 with IPv4 explicit NULL tag
0 with IPv4 explicit NULL cksum errors
0 with router alert tag
0 lsp ping packets (ttl-expired/router alert)
0 with ttl expired
0 with tag encoding error
0 packets discarded, no route

vpls:
0 total packets received
0 with size smaller than minimum
0 with incorrect version number
0 packets for this host
0 packets with no logical interface
0 packets with no family
0 packets with no route table
0 packets with no auxiliary table
0 packets with no corefacing entry
0 packets with no CE-facing entry
0 mac route learning requests
0 mac routes learnt
0 requests to learn an existing route
0 learning requests while learning disabled on interface
0 learning requests over capacity
0 mac routes moved
0 requests to move static route
0 mac route aging requests
0 mac routes aged
0 bogus address in aging requests
0 requests to age static route
0 requests to re-ageout aged route
0 requests involving multiple peer FEs
0 aging acks from PFE
0 aging non-acks from PFE
0 aging requests timed out waiting on FEs
0 aging requests over max-rate
0 errors finding peer FEs

```

show system statistics (EX Series Switches)

```

user@host> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks

```

```
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)
3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
    47 connections updated cached RTT on close
    47 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
3028 keepalive timeouts
    3027 keepalive probes sent
    1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
    0 retransmitted
    0 dupsyn
    4 dropped
    78 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
```



```
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error
```

```
0 errors not generated because old message was icmp
Output histogram:
    295 echo reply
    9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    295 echo
295 message responses generated

igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent

raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMPv2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMPv2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket

arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
```

```

0 datagrams with multicast source address
0 datagrams with multicast source address
0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too may headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f

icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums

```

```
0 Messages with bad length
    0 No route
    0 Administratively prohibited
    0 Beyond scope
    0 Address unreachable
    0 Port unreachable
    0 packet too big
    0 Time exceed transit
    0 Time exceed reassembly
    0 Erroneous header field
    0 Unrecognized next header
    0 Unrecognized option
    0 redirect
    0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
    0 reserved
    0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
    0 reserved
    0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
crlnl:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
```

```

0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections

```

```
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent

tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output

ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result
0 Input packets for which rt lookup is bypassed

mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
```

```

0 Packets used first nexthop in ecmp unilist
vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop

```

```
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
  0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```

show system statistics (TX Matrix Router)

```
user@host> show system statistics
sfc0-re0:
```

Tcp:

```

361694 packets sent
    326507 data packets (103237236 bytes)
    2343 data packets retransmitted (2673324 bytes)
    0 resends initiated by MTU discovery
    33857 ack only packets (31613 packets delayed)
    0 URG only packets
    14 window probe packets
    387 window update packets
    1108 control packets
345879 packets received
    298207 acks(for 103141728 bytes)
    438 duplicate acks
    0 acks for unsent data
    204578 packets received in-sequence(13820995 bytes)
    6 completely duplicate packets(18 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    899 window update packets
    166 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
406 connection requests
233 connection accepts
0 bad connection attempts
0 listen queue overflows
616 connections established (including accepts)
911 connections closed (including 41 drops)
    346 connections updated cached RTT on close
    346 connections updated cached RTT variance on close
    200 connections updated cached ssthresh on close
23 embryonic connections dropped
298155 segments updated rtt(of 287216 attempts)
1163 retransmit timeouts
    27 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
5 keepalive timeouts
    5 keepalive probes sent
    0 connections dropped by keepalive
69922 correct ACK header predictions
34993 correct data packet header predictions
233 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    233 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received

```

- 23 SACK recovery episodes
- 68 segment retransmits in SACK recovery episodes
- 71542 byte retransmits in SACK recovery episodes
- 158 SACK options (SACK blocks) received
- 0 SACK options (SACK blocks) sent
- 0 SACK scoreboard overflow
- 0 ACKs sent in response to in-window but not exact RSTs
- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 259 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

1cc0-re0:

Tcp:

- 346 packets sent
 - 222 data packets (22894 bytes)
 - 0 data packets retransmitted (0 bytes)
 - 0 resends initiated by MTU discovery
 - 80 ack only packets (12 packets delayed)
 - 0 URG only packets
 - 0 window probe packets
 - 5 window update packets
 - 42 control packets
- 358 packets received
 - 268 acks(for 22939 bytes)
 - 9 duplicate acks
 - 0 acks for unsent data
 - 203 packets received in-sequence(33820 bytes)
 - 0 completely duplicate packets(0 bytes)
 - 0 old duplicate packets
 - 0 packets with some duplicate data(0 bytes duped)
 - 0 out-of-order packets(0 bytes)
 - 0 packets of data after window(0 bytes)
 - 0 window probes
 - 6 window update packets
 - 0 packets received after close
 - 0 discarded for bad checksums
 - 0 discarded for bad header offset fields
 - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 35 connections closed (including 2 drops)
 - 3 connections updated cached RTT on close
 - 3 connections updated cached RTT variance on close
 - 0 connections updated cached ssthresh on close
- 0 embryonic connections dropped
- 268 segments updated rtt(of 247 attempts)
- 0 retransmit timeouts
 - 0 connections dropped by retransmit timeout
- 0 persist timeouts
 - 0 connections dropped by persist timeout
- 0 keepalive timeouts
 - 0 keepalive probes sent

```

    0 connections dropped by keepalive
0 correct ACK header predictions
42 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

1cc1-re0:

 Tcp:

```

348 packets sent
    223 data packets (22895 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    81 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    42 control packets
360 packets received
    269 acks(for 22940 bytes)
    9 duplicate acks
    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields

```

```
    0 discarded because packet too short
13 connection requests
18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
36 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
269 segments updated rtt(of 248 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
43 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
```

1cc2-re0:

Tcp:

```
405 packets sent
    271 data packets (23926 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    86 ack only packets (13 packets delayed)
    0 URG only packets
```

```

    0 window probe packets
    5 window update packets
    46 control packets
418 packets received
    321 acks(for 23975 bytes)
    9 duplicate acks
    0 acks for unsent data
    234 packets received in-sequence(34403 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
15 connection requests
19 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)
39 connections closed (including 2 drops)
    4 connections updated cached RTT on close
    4 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
321 segments updated rtt(of 299 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
48 correct data packet header predictions
19 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    19 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs

```

- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 5 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

lcc3-re0:

Tcp:

- 346 packets sent
 - 221 data packets (22895 bytes)
 - 0 data packets retransmitted (0 bytes)
 - 0 resends initiated by MTU discovery
 - 81 ack only packets (13 packets delayed)
 - 0 URG only packets
 - 0 window probe packets
 - 5 window update packets
 - 42 control packets
- 360 packets received
 - 267 acks(for 22940 bytes)
 - 9 duplicate acks
 - 0 acks for unsent data
 - 203 packets received in-sequence(33820 bytes)
 - 0 completely duplicate packets(0 bytes)
 - 0 old duplicate packets
 - 0 packets with some duplicate data(0 bytes duped)
 - 0 out-of-order packets(0 bytes)
 - 0 packets of data after window(0 bytes)
 - 0 window probes
 - 6 window update packets
 - 0 packets received after close
 - 0 discarded for bad checksums
 - 0 discarded for bad header offset fields
 - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 35 connections closed (including 2 drops)
 - 3 connections updated cached RTT on close
 - 3 connections updated cached RTT variance on close
 - 0 connections updated cached ssthresh on close
- 0 embryonic connections dropped
- 267 segments updated rtt(of 246 attempts)
- 0 retransmit timeouts
 - 0 connections dropped by retransmit timeout
- 0 persist timeouts
 - 0 connections dropped by persist timeout
- 0 keepalive timeouts
 - 0 keepalive probes sent
 - 0 connections dropped by keepalive
- 0 correct ACK header predictions
- 43 correct data packet header predictions
- 18 syncache entries added
 - 0 retransmitted
 - 0 dupsyn
 - 0 dropped

```

18 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

show system statistics (QFX Series)

```

user@switch> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)
3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)

```

```
546596 connections closed (including 6 drops)
47 connections updated cached RTT on close
47 connections updated cached RTT variance on close
0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
0 connections dropped by retransmit timeout
0 persist timeouts
0 connections dropped by persist timeout
3028 keepalive timeouts
3027 keepalive probes sent
1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
0 retransmitted
0 dupsyn
4 dropped
78 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
```



```

0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
295 echo reply
9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
295 echo
295 message responses generated
igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum

```

```
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent
raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMP2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMP2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket
arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast source address
0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
```

```
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
0 No route
0 Administratively prohibited
0 Beyond scope
0 Address unreachable
0 Port unreachable
0 packet too big
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 redirect
0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
```

```
0 Bytes sent from userland
histogram by message type:
0 reserved
0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
0 reserved
0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
c1n1:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
```

```
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
```

```
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result0 Input packets for which rt lookup
  is bypassed
mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist
vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
582 Copyright © 2010, Juniper Networks, Inc.
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
```

```
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
```

```
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```


show system storage

List of Syntax	Syntax on page 1061 Syntax (EX Series Switches) on page 1061 Syntax (TX Matrix Router) on page 1061 Syntax (TX Matrix Plus Router) on page 1061 Syntax (MX Series Router) on page 1061 Syntax (QFX Series) on page 1061
Syntax	show system storage <detail>
Syntax (EX Series Switches)	show system storage <detail> <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system storage <detail> <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system storage <detail> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system storage <detail> <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system storage <detail> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display statistics about the amount of free disk space in the router's or switch's file systems.
Options	none —Display standard information about the amount of free disk space in the router's or switch's file systems. detail —(Optional) Display detailed output.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system storage statistics for all the routers in the chassis.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for all routers connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Display system storage statistics for the fabric control Routing Engines or fabric manager Routing Engines.

interconnect-device *name*—(QFabric systems only) (Optional) Display system storage statistics for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display system storage statistics for the Node group.

scc—(TX Matrix routers only) (Optional) Display system storage statistics for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system storage statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system storage** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system storage on page 1063](#)
[show system storage \(TX Matrix Plus Router\) on page 1064](#)
[show system storage \(QFX3500 Switch\) on page 1066](#)

Output Fields [Table 91 on page 1063](#) describes the output fields for the **show system storage** command. Output fields are listed in the approximate order in which they appear.

Table 91: show system storage Output Fields

Field Name	Field Description
Filesystem	Name of the filesystem.
Size	Size of the filesystem.
Used	Amount of space used in the filesystem.
Avail	Amount of space available in the filesystem.
Capacity	Percentage of the filesystem space that is being used.
Mounted on	Directory in which the filesystem is mounted.

Sample Output

show system storage

```

user@host> show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      77M       37M       34M     52%      /
devfs           16K       16K        0B    100%    /dev/
/dev/vn0        12M       12M        0B    100%  /packages/mnt/jbase
/dev/vn1        39M       39M        0B    100%
/packages/mnt/jkernel-7.2R1.7
/dev/vn2        12M       12M        0B    100%
/packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3         2.3M     2.3M        0B    100%
/packages/mnt/jdocs-7.2R1.7
/dev/vn4        14M       14M        0B    100%
/packages/mnt/jroute-7.2R1.7
/dev/vn5         4.5M     4.5M        0B    100%

```

```

/packages/mnt/jcrypto-7.2R1.7
mfs:172          1.5G      4.0K      1.3G      0% /tmp
/dev/ad0s1e      12M       20K       11M       0% /config
procfs          4.0K      4.0K       0B      100% /proc
/dev/ad1s1f      9.4G      4.9G      3.7G      57% /var

```

show system storage (TX Matrix Plus Router)

```
user@host> show system storage
```

```
sfc0-re0:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.4G     178M     2.9G      6%      /
devfs           1.0K     1.0K       0B     100%    /dev
devfs           1.0K     1.0K       0B     100%    /dev/
/dev/md0        33M      33M       0B     100%    /packages/mnt/jbase
/dev/md1        216M     216M       0B     100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2         66M      66M       0B     100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3         4.1M     4.1M       0B     100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4         57M      57M       0B     100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5         15M      15M       0B     100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6         34M      34M       0B     100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7         2.0G     10.0K     1.8G      0%    /tmp
/dev/md8         2.0G      1.0M     1.8G      0%    /mfs
/dev/ad0s1e     383M      82K     352M      0%    /config
procfs          4.0K     4.0K       0B     100%    /proc
/dev/ad1s1f     52G      7.5G     40G      16%    /var

```

```
lcc0-re0:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.4G     178M     2.9G      6%      /
devfs           1.0K     1.0K       0B     100%    /dev
devfs           1.0K     1.0K       0B     100%    /dev/
/dev/md0        33M      33M       0B     100%    /packages/mnt/jbase
/dev/md1        216M     216M       0B     100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2         66M      66M       0B     100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3         4.1M     4.1M       0B     100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4         57M      57M       0B     100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5         15M      15M       0B     100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6         34M      34M       0B     100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7         2.0G     10.0K     1.8G      0%    /tmp
/dev/md8         2.0G     540K     1.8G      0%    /mfs
/dev/ad0s1e     383M      88K     352M      0%    /config
procfs          4.0K     4.0K       0B     100%    /proc
/dev/ad1s1f     52G      6.3G     41G      13%    /var

```

```
lcc1-re0:
```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	23G	13G	7.7G	64%	/var

lcc2-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	64K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	23G	3.7G	17G	18%	/var

lcc3-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					

/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	34K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	23G	18G	3.5G	84%	/var

show system storage (QFX3500 Switch)

```
user@switch> show system storage
```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/da0s2a	343M	192M	123M	61%	/
devfs	1.0K	1.0K	0B	100%	/dev
/dev/md0	119M	119M	0B	100%	/packages/mnt/jbase
/dev/md1	513M	513M	0B	100%	
/packages/mnt/jkernel-qfx-11.1R1.5					
/dev/md2	37M	37M	0B	100%	
/packages/mnt/jpfe-qfx-e9xxx-11.1R1.5					
/dev/md3	6.0M	6.0M	0B	100%	
/packages/mnt/jdocs-qfx-11.1R1.5					
/dev/md4	216M	216M	0B	100%	
/packages/mnt/jroute-qfx-11.1R1.5					
/dev/md5	59M	59M	0B	100%	
/packages/mnt/jcrypto-qfx-11.1R1.5					
/dev/md6	85M	85M	0B	100%	
/packages/mnt/jswitch-qfx-11.1R1.5					
/dev/md7	63M	8.0K	58M	0%	/tmp
/dev/da0s2f	228M	14M	196M	7%	/var
/dev/da0s3d	590M	3.0M	540M	1%	/var/tmp
/dev/da0s3e	104M	162K	95M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc

show system uptime

List of Syntax	Syntax on page 1067 Syntax (EX Series Switches) on page 1067 Syntax (QFX Series) on page 1067 Syntax (TX Matrix Router) on page 1067 Syntax (TX Matrix Plus Router) on page 1067 Syntax (MX Series Router) on page 1067
Syntax	show system uptime
Syntax (EX Series Switches)	show system uptime <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system uptime <director-group <i>name</i> > <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Syntax (TX Matrix Router)	show system uptime <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system uptime <detail> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system uptime <all-members> <invoke-on> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.
Options	none —Show time since the system rebooted and processes started. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus

router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

director-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

infrastructure *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

invoke-on—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

scc—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system uptime** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation

- [Monitoring System Process Information on page 315](#)
- [Monitoring System Properties on page 316](#)
- [10-Gigabit Ethernet LAN/WAN PIC with XFP \(T640 Router\)](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

- [show system uptime on page 1070](#)
- [show system uptime all-lcc \(TX Matrix Router\) on page 1070](#)
- [show system uptime all-lcc \(TX Matrix Plus Router\) on page 1070](#)
- [show system uptime \(QFX Series\) on page 1071](#)

Output Fields Table 92 on page 1069 describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

Table 92: show system uptime Output Fields

Field Name	Field Description
Current time	Current system time in UTC.
System booted	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
Protocols started	Date and time when the routing protocols were last started and how long they have been running.
Last configured	Date and time when a configuration was last committed. Also shows the name of the user who issued the last commit command.
time and up	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

Sample Output

show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00
```

show system uptime all-lcc (TX Matrix Plus Router)

```
user@host> show system uptime all-lcc
sfc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

lcc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc1-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc2-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00
```

lcc3-re0:

Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

show system uptime (QFX Series)

```
user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00
```

show system users

List of Syntax	Syntax on page 1072 Syntax (TX Matrix Router) on page 1072 Syntax (TX Matrix Plus Router) on page 1072 Syntax (MX Series Router) on page 1072
Syntax	show system users <no-resolve>
Syntax (TX Matrix Router)	show system users <all-chassis all-lcc lccnumber scc> <no-resolve>
Syntax (TX Matrix Plus Router)	show system users <detail> <all-chassis all-lcc lcc number sfc number> <no-resolve>
Syntax (MX Series Router)	show system users <all-members> <local> <member member-id> <no-resolve>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in JUNOS OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	List information about the users who are currently logged in to the router or switch.



NOTE: The **show system users** command lists the information about administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client. The output does not list information about web users or automated users that are logged in from a remote client application using Junos XML APIs, such as NETCONF.

- Options** **none**—List information about the users who are currently logged in to the router or switch.
- all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show users currently logged in to all the routers in the chassis.
- all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to all connected T1600 or T4000 LCCs.
- all-members**—(MX Series routers only) (Optional) Display users currently logged in to all members of the Virtual Chassis configuration.

lcc number—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display users currently logged in to the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display users currently logged in to the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

no-resolve—(Optional) Do not attempt to resolve IP addresses to hostnames.

scc—(TX Matrix routers only) (Optional) Show users currently logged in to the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Show users currently logged in to the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system users** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation [• Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system users on page 1074](#)
[show system users lcc no-resolve \(TX Matrix, TX Matrix Plus Router\) on page 1074](#)
[show system users \(TX Matrix Plus Router\) on page 1074](#)
[show system users \(QFX Series\) on page 1075](#)
[show system users no-resolve \(QFX Series\) on page 1075](#)

Output Fields Table 93 on page 1074 describes the output fields for the **show system users** command. Output fields are listed in the approximate order in which they appear.

Table 93: show system users Output Fields

Field Name	Field Description
time and up	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
USER	Username.
TTY	Terminal through which the user is logged in.
FROM	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
LOGIN@	Time when the user logged in.
IDLE	How long the user has been idle.
WHAT	Processes that the user is running.

Sample Output

show system users

```
user@host> show system users
 7:30PM up 4 days, 2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER   TTY FROM          LOGIN@  IDLE WHAT
root   d0  -             Fri05PM 4days -csh (csh)
blue   p0  leve15.compan 7:30PM  - cli
```

show system users lcc no-resolve (TX Matrix, TX Matrix Plus Router)

```
user@host> show system users lcc 2 no-resolve

lcc2-re0:
-----
10:34AM PDT up 1 day, 7:11, 5 users, load averages: 0.03, 0.01, 0.00
USER   TTY   FROM          LOGIN@  IDLE WHAT
root   d0    -             3:21AM  7:12 /bin/csh
regress p0    scc-re0       10:15AM - telnet hostA
regress p1    scc-re0       10:16AM - telnet hostA
regress p2    scc-re0       10:19AM - telnet hostA
regress p3    scc-re0       10:24AM - telnet hostA
```

show system users (TX Matrix Plus Router)

```
user@host> show system users
sfc0-re0:
-----
1:41AM up 26 mins, 3 users, load averages: 0.08, 0.04, 0.03
```

```

USER      TTY      FROM                                LOGIN@  IDLE WHAT
regress   p0       10.209.208.123                     1:18AM  21 cli
regress   p1       172.17.29.207                      1:37AM   2 cli
regress   p2       172.17.28.19                       1:40AM   - cli

lcc0-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.00, 0.03

lcc1-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.02, 0.03

lcc2-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.16, 0.06, 0.02

lcc3-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.12, 0.04, 0.04

regress@aj> show system users
sfc0-re0:
-----
1:42AM up 28 mins, 4 users, load averages: 0.02, 0.03, 0.02
USER      TTY      FROM                                LOGIN@  IDLE WHAT
regress   p0       pssraj-t61.jnpr.net                1:18AM  22 cli
regress   p1       eng-shell14.juniper.net             1:37AM   - cli
regress   p2       bigpink.juniper.net                 1:40AM   - cli
regress   p3       sv-cutty-01.englab.juniper.net       1:42AM   - csh (csh)

lcc0-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.02, 0.01, 0.03

lcc1-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.07, 0.04, 0.03

lcc2-re0:
-----
1:42AM up 27 mins, 0 users, load averages: 0.07, 0.06, 0.02

lcc3-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.05, 0.04, 0.04

```

show system users (QFX Series)

```

user@switch> show system users
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                     2:54AM  39 -cli (cli)
tlewis    p1       172.22.18.117                     3:01AM   - -cli (cli)
tcheng    p2       172.22.17.197                     3:08AM  11 -cli (cli)

```

show system users no-resolve (QFX Series)

```

user@switch> show system users no-resolve
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                     2:54AM  39 -cli (cli)

```

tLewis	p1	172.22.18.117	3:01AM	- -cli (cli)
tcheng	p2	172.22.17.197	3:08AM	11 -cli (cli)

show system virtual-memory

List of Syntax	Syntax on page 1077 Syntax (EX Series) on page 1077 Syntax (TX Matrix Router) on page 1077 Syntax (TX Matrix Plus Router) on page 1077 Syntax (MX Series Router) on page 1077 Syntax (QFX Series) on page 1077
Syntax	show system virtual-memory
Syntax (EX Series)	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system virtual-memory <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system virtual-memory <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system virtual-memory <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use the show system virtual-memory command for troubleshooting with Juniper Networks Customer Support.
Options	none —Display kernel dynamic memory usage information. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for all chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.

infrastructure *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Node group.

scc—(TX Matrix routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system virtual-memory** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix

or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.



NOTE: The `show system virtual-memory` command with the `| display XML` pipe option now displays XML output for the command in the parent tags: `<vmstat-memstat-malloc>`, `<vmstat-memstat-zone>`, `<vmstat-sumstat>`, `<vmstat-intr>`, and `<vmstat-kernel-state>` with each child element as a separate XML tag. In Junos OS Releases 10.1 and earlier, the `| display XML` option for this command does not have an XML API element and the entire output is displayed in a single `<output>` tag element.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	show system virtual-memory on page 1081 show system virtual-memory scc (TX Matrix Router) on page 1085 show system virtual-memory sfc (TX Matrix Plus Router) on page 1086 show system virtual-memory display xml on page 1089 show system virtual-memory (QFX Series) on page 1112
Output Fields	Table 94 on page 1080 lists the output fields for the <code>show system virtual-memory</code> command. Output fields are listed in the approximate order in which they appear.

Table 94: show system virtual-memory Output Fields

Field Name	Field Description
Memory statistics by bucket size	
Size	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
In Use	Number of memory blocks of this size that are in use (bytes).
Free	Number of memory blocks of this size that are free (bytes).
Requests	Number of memory allocation requests made.
HighWater	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
Couldfree	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
Memory usage type by bucket size	
Size	Memory block size (bytes).
Type(s)	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
Memory statistics by type	
Type	Kernel module that is using dynamic memory.
InUse	Number of memory blocks used by this type. The number is rounded up.
MemUse	Amount of memory in use, in kilobytes (KB).
HighUse	Maximum memory ever used by this type.
Limit	Maximum memory that can be allocated to this type.
Requests	Total number of dynamic memory allocation requests this type has made.
Type Limit	Number of times requests were blocked for reaching the maximum limit.
Kern Limit	Number of times requests were blocked for the kernel map.
Size(s)	Memory block sizes this type is using.
Memory Totals	
In Use	Total kernel dynamic memory in use (bytes, rounded up).
Free	Total kernel dynamic memory free (bytes, rounded up).

Table 94: show system virtual-memory Output Fields (*continued*)

Field Name	Field Description
Requests	Total number of memory allocation requests.
ITEM	Kernel module that is using memory.
Size	Memory block size (bytes).
Limit	Maximum memory that can be allocated to this type.
Used	Number of memory blocks used by this type. The number is rounded up.
Free	Number of memory blocks available to this type.
Requests	Total number of memory allocation requests this type has made.
interrupt	Timer events and scheduling interruptions.
total	Total number of interruptions for each type.
rate	Interruption rate.
Total	Total for all interruptions.

Sample Output

show system virtual-memory

```

user@host> show system virtual-memory
Memory statistics by bucket size
Size    In Use    Free    Requests  HighWater  Couldfree
16      906      118     154876    1280       0
32      455      313     209956    640        0
64      4412     260     75380     320        20
128     3200     32      19361     160        81
256     1510     10      8844      80         4
512     446      2       5085      40         0
1K      18       2       5901      20         0
2K      1128     2       4445      10        1368
4K      185      1       456       5          0
8K      5        1       2653      5          0
16K     181      0       233       5          0
32K     2        0       1848      5          0
64K     20       0       22        5          0
128K    5        0       5         5          0
256K    2        0       2         5          0
512K    1        0       1         5          0

Memory usage type by bucket size
Size    Type(s)
16      uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
        DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
        MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,

```

```

pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
32 atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
64 isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
cluster_save buffer, vfscache, pcb, soname, proc-args, file,
AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
128 ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
rtnexthop, iffamily
256 iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
512 UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
Gzip trees, ISOFS mount, msg, ioctlops, ATA generic, bus, proc,
pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
1K iftable, temp, devbuf, NQ NFS Lease, kqueue, kld, AD driver,
Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
2K uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
Gzip trees, ioctlops, bus, ipfw, ifstat, rcache
4K memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
8K temp, devbuf, syncache, Gzip trees
16K indirdep, temp, devbuf, shm, msg
32K pagedep, kld, Gzip trees
64K VM pgdata, devbuf, MSDOSFS mount
128K UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K mbuf, vfscache
512K SWAP

```

Memory statistics by type					Type	Kern		
Type	InUse	MemUse	HighUse	Limit	Requests	Limit	Limit	Size(s)
isadev	13	1K	1K127753K	13	0	0	0	64
atkbddev	2	1K	1K127753K	2	0	0	0	32
uc_devlist	24	3K	3K127753K	24	0	0	0	16,2K
nexusdev	3	1K	1K127753K	3	0	0	0	16
memdesc	1	4K	4K127753K	1	0	0	0	4K
mbuf	1	152K	152K127753K	1	0	0	0	256K
iflogical	6	2K	2K127753K	6	0	0	0	256
iftable	17	9K	9K127753K	18	0	0	0	16,64,256,1K,4K
ZONE	15	2K	2K127753K	15	0	0	0	128
VM pgdata	1	64K	64K127753K	1	0	0	0	64K
UFS mount	12	26K	26K127753K	12	0	0	0	512,2K,4K
UFS ihash	1	128K	128K127753K	1	0	0	0	128K
MFS node	6	2K	3K127753K	35	0	0	0	64,256
FFS node	906	227K	227K127753K	1352	0	0	0	256
dirrem	0	0K	4K127753K	500	0	0	0	32
mkdir	0	0K	1K127753K	38	0	0	0	32
diradd	0	0K	6K127753K	521	0	0	0	32
freefile	0	0K	4K127753K	374	0	0	0	32
freeblks	0	0K	8K127753K	219	0	0	0	128
freefrag	0	0K	1K127753K	193	0	0	0	32
allocindir	0	0K	25K127753K	1518	0	0	0	64
indirdep	0	0K	17K127753K	76	0	0	0	32,16K
allocdirect	0	0K	10K127753K	760	0	0	0	64
bmsafemap	0	0K	1K127753K	72	0	0	0	32

newblk	1	1K	1K127753K	2279	0	0	32,256
inodedep	1	128K	175K127753K	2367	0	0	128,128K
pagedep	1	32K	33K127753K	47	0	0	64,32K
temp	1239	92K	96K127753K	8364	0	0	16,32,64K
devbuf	1413	5527K	5527K127753K	1535	0	0	16,32,64,128,256
lockf	38	3K	3K127753K	2906	0	0	64
atexit	1	1K	1K127753K	1	0	0	16
zombie	0	0K	2K127753K	3850	0	0	128
NFS hash	1	128K	128K127753K	1	0	0	128K
NQNFS Lease	1	1K	1K127753K	1	0	0	1K
NFS daemon	1	1K	1K127753K	1	0	0	256
syncache	1	8K	8K127753K	1	0	0	8K
COS	353	44K	44K127753K	353	0	0	16,32,64,128
BPF	189	3K	3K127753K	189	0	0	16
MSDOSFS mount	1	64K	64K127753K	1	0	0	64K
NULLFS hash	1	1K	1K127753K	1	0	0	64
DEVFS mount	2	1K	1K127753K	2	0	0	16
DEVFS name	487	31K	31K127753K	487	0	0	64
DEVFS node	471	58K	58K127753K	479	0	0	16,128
vnodes	28	7K	7K127753K	429	0	0	16,32,64,128,256
mount	15	8K	8K127753K	18	0	0	16,128,512
cluster_save buffer	0	0K	1K127753K	55	0	0	32,64
vfscache	1898	376K	376K127753K	3228	0	0	64,128,256K
BIO buffer	49	98K	398K127753K	495	0	0	512,2K
pcb	159	16K	17K127753K	399	0	0	16,32,64,128,2K
soname	82	10K	10K127753K	42847	0	0	16,32,64,128
proc-args	57	2K	3K127753K	2105	0	0	16,32,64,128,256
ptys	32	16K	16K127753K	32	0	0	512
ttys	254	33K	33K127753K	522	0	0	128,512
kqueue	5	3K	4K127753K	23	0	0	256,1K
sigio	1	1K	1K127753K	27	0	0	32
file	383	24K	24K127753K	16060	0	0	64
file desc	76	19K	20K127753K	3968	0	0	256
shm	1	12K	12K127753K	1	0	0	16K
dev_t	286	36K	36K127753K	286	0	0	128
timecounter	10	2K	2K127753K	10	0	0	128
kld	11	117K	122K127753K	34	0	0	16,32,128,1K,4K
AR driver	1	1K	3K127753K	5	0	0	64,512,2K
AD driver	2	2K	3K127753K	2755	0	0	64,1K
Gzip trees	0	0K	46K127753K	133848	0	0	32,64,128,256
ISOFS node	1136	142K	142K127753K	1189	0	0	128
ISOFS mount	9	132K	132K127753K	10	0	0	512,128K
sem	3	6K	6K127753K	3	0	0	1K,4K
MD disk	2	2K	2K127753K	2	0	0	16,1K
msg	4	25K	25K127753K	4	0	0	512,4K,16K
rman	59	4K	4K127753K	461	0	0	16,64
ioctlops	0	0K	2K127753K	992	0	0	512,2K
taskqueue	2	1K	1K127753K	2	0	0	32
SWAP	2	413K	413K127753K	2	0	0	32,512K
ATA generic	6	3K	3K127753K	6	0	0	16,512
eventhandler	17	1K	1K127753K	17	0	0	32,64
bus	340	30K	31K127753K	794	0	0	16,32,64,128,256
sysctl	0	0K	1K127753K	130262	0	0	16,32,64
uidinfo	4	1K	1K127753K	10	0	0	32,128
cred	22	3K	3K127753K	3450	0	0	128
subproc	156	10K	10K127753K	7882	0	0	32,64,256
proc	2	1K	1K127753K	2	0	0	512
session	12	2K	2K127753K	34	0	0	128
pgrp	16	1K	1K127753K	45	0	0	32
ippool	1	1K	1K127753K	1	0	0	16
pfestat	0	0K	1K127753K	47349	0	0	16,32,64,512

pic	5	1K	1K127753K	5	0	0	64,128
lr	1	1K	1K127753K	1	0	0	512
itable32	110	4K	4K127753K	110	0	0	32
itable16	161	26K	26K127753K	161	0	0	128,256
ifstate	694	159K	160K127753K	1735	0	0	16,32,64,128,1K
pfe_ipc	0	0K	1K127753K	56218	0	0	16,32,64,128,1K
mkey	250	4K	4K127753K	824	0	0	16,32,64
ifaddr	9	1K	1K127753K	9	0	0	64
sysctl	0	0K	1K127753K	30	0	0	256
rtable	49	6K	6K127753K	307	0	0	16,32,64,128,512
ifmaddr	22	1K	1K127753K	22	0	0	16,32
ipfw	23	10K	10K127753K	48	0	0	16,32,64,512,2K
ifstat	698	805K	805K127753K	698	0	0	128,512,2K
rcache	4	8K	8K127753K	4	0	0	2K
rnode	27	1K	1K127753K	285	0	0	16,32
metrics	1	1K	1K127753K	3	0	0	128
rtnexthop	57	9K	9K127753K	312	0	0	32,128,256,512
iffamily	12	2K	2K127753K	12	0	0	128

Memory Totals:	In Use	Free	Requests
	9311K	54K	489068

ITEM	SIZE	LIMIT	USED	FREE	REQUESTS
PIPE:	192,	0,	4,	81,	4422
SWAPMETA:	160,	95814,	0,	0,	0
unpcb:	160,	0,	114,	36,	279
ripcb:	192,	25330,	5,	37,	5
syncache:	128,	15359,	0,	64,	5
tcpcb:	576,	25330,	23,	12,	32
udpcb:	192,	25330,	14,	28,	255
socket:	256,	25330,	246,	26,	819
KNOTE:	96,	0,	27,	57,	71
NFSNODE:	352,	0,	0,	0,	0
NFSMOUNT:	544,	0,	0,	0,	0
VNODE:	224,	0,	2778,	43,	2778
NAMEI:	1024,	0,	0,	8,	40725
VMSPACE:	192,	0,	57,	71,	3906
PROC:	448,	0,	73,	17,	3923
DP fakepg:	64,	0,	0,	0,	0
PV ENTRY:	28,	499566,	44530,	152053,	1525141
MAP ENTRY:	48,	0,	1439,	134,	351075
KMAP ENTRY:	48,	35645,	179,	119,	10904
MAP:	108,	0,	7,	3,	7
VM OBJECT:	92,	0,	2575,	109,	66912

```

792644 cpu context switches
9863474 device interrupts
286510 software interrupts
390851 traps
3596829 system calls
  16 kernel threads created
 3880 fork() calls
   27 vfork() calls
    0 rfork() calls
    0 swap pager pageins
    0 swap pager pages paged in
    0 swap pager pageouts
    0 swap pager pages paged out
  380 vnode pager pageins
  395 vnode pager pages paged in
  122 vnode pager pageouts

```



```

1476 vnode pager pages paged out
    0 page daemon wakeups
    0 pages examined by the page daemon
101 pages reactivated
161722 copy-on-write faults
    0 copy-on-write optimized faults
84623 zero fill pages zeroed
83063 zero fill pages prezeroed
    7 intransit blocking page faults
535606 total VM faults taken
    0 pages affected by kernel thread creation
238254 pages affected by fork()
    2535 pages affected by vfork()
    0 pages affected by rfork()
283379 pages freed
    0 pages freed by daemon
190091 pages freed by exiting processes
17458 pages active
29166 pages inactive
    0 pages in VM cache
10395 pages wired down
134610 pages free
    4096 bytes per page
183419 total name lookups
    cache hits (90% pos + 7% neg) system 0% per-directory
    deletions 0%, falsehits 0%, toolong 0%

interrupt          total          rate
ata0 irq14         113338           3
mux irq7           727643          21
fxp1 irq10         1178671          34
sio0 irq4           833              0
clk irq0           3439769          99
rtc irq8           4403221          127
Total              9863475          286

Kernel direct memory map:
    4423 pages used
    4057340 pages maximum

```

Note: Kernel direct memory map only displays for 64 bit platform.

show system virtual-memory scc (TX Matrix Router)

```
user@host> show system virtual-memory scc
```

```

Memory statistics by bucket size
Size  In Use  Free  Requests  HighWater  Couldfree
16    898    126   749493    1280       0
32    2018   1310  980643    640       632
64    3490   13342 935420    320       5365
...

Memory usage type by bucket size
Size  Type(s)
16    uc_devlist, COS, BPF, DEVFS mount, DEVFS node, vnodes, mount, pcb,
      soname, rman, bus, sysctl, ifstate, pfe_ipc, mkey, socket, rtable,
      ifmaddr, ipfw, rnode, iftable, temp, devbuf, atexit, proc-args, kld,
      MD disk
32    atkbddev, Gzip trees, dirrem, mkdir, diradd, freefile, freefrag,
      indirdep, bmsafemap, newblk, tseg_qent, COS, vnodes,

```

...

```

Memory statistics by type
      Type InUse MemUse HighUse Limit Requests Limit Limit Size(s)
      isadev 12 1K 1K166400K 12 0 0 64
      atkbdddev 2 1K 1K166400K 2 0 0 32
      uc_devlist 24 3K 3K166400K 24 0 0 16,2K
      ....

Memory Totals: In Use Free Requests
                6091K 1554K 2897122

```

show system virtual-memory sfc (TX Matrix Plus Router)

```

user@host> show system virtual-memory sfc 0
sfc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue 1 1K - 1 64
  entropy 1024 64K - 1024 64
  linker 487 6272K - 1163 16,32,64,4096,32768,131072
  USB 127 10K - 127 16,32,64,128,256,1024,2048
  lockf 46 3K - 98418 64
  USBdev 10 2K - 34 16,128,2048,16384
ifstateSLLNode 0 0K - 1096 16
  devbuf 21243 15683K - 21810
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  temp 1283 151K - 2483472
16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072
  ip6ndp 0 0K - 4 64
  in6ifmulti 1 1K - 1 64
  in6grentry 1 1K - 1 64
  iflogical 20 5K - 29 2048
  iffamilly 45 6K - 69 32,1024,2048
  rtnexthop 266 46K - 608013 32,256,512,1024,2048,4096
  metrics 31 4K - 54 256
  rnode 212 4K - 607848 16,32
  rcache 4 8K - 4 65536
  iflist 0 0K - 6 16,64
  ifdevice 11 8K - 17 16,32768
  ifstat 424 472K - 427 512,16384,65536
  ipfw 42 23K - 145
16,32,64,128,256,512,1024,16384,32768,65536,131072
  ifmaddr 415 11K - 415 16,32
  rtable 329 28K - 608066 16,32,64,128,1024,16384
  sysctl 0 0K - 887976 16,32,64,4096,16384,32768
  ifaddr 64 5K - 70 32,64,128
  mkey 331 6K - 12528 16,128
  pfe_ipc 0 0K - 7299115
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  ifstate 1245054 70088K - 3040437
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768
  idxbucket 1 1K - 1 16
  itable16 5069 1250K - 5103 1024,4096
  itable32 157 10K - 157 64
  itable64 2 1K - 2 128
  lr 1 1K - 4 16384
  pic 37 6K - 37 64,16384
  pfestat 0 0K - 6220 32,64,128,256,131072
  gencfg 1486 424K - 2614 16,32,64,256,512,16384,32768,65536

```

```

        jsr      2      1K      -      22  16
        idl      1      4K      -      165
32, 64, 128, 256, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072
        rtmsg    0      0K      -      16  131072
        module   250    16K      -      250  64, 128
        mtx_pool  1      8K      -      1   64, 128
        DEVFS3   113    13K      -      114  256
        DEVFS1   106    24K      -      106  2048
        pgrp     15     1K      -      8600 64
        session  11     2K      -      2829 512
        proc      2     1K      -      2   16384
        subproc   296    572K     -      24689 2048, 131072
        cred      38     5K      -      619244 256
        plimit    18     4K      -      21311 2048
        uidinfo   3      1K      -      10   32, 512
        sysctluid 2701    82K     -      2701 16, 32, 64
        sysctltmp 0      0K      -      15572 16, 32, 64, 1024
        umtx     171    11K      -      171   64
        SWAP      2    277K     -      2     64
        bus      779    125K     -      3072 16, 32, 64, 128, 32768
        bus-sc    67     62K     -      1477
16, 32, 64, 512, 1024, 2048, 8192, 16384, 65536, 131072
        devstat   8     17K     -      8   16, 131072
        eventhandler 46    2K     -      47   32, 128
        kobj      93    186K     -      111  65536
        DEVFS      8     1K     -      9   16, 64
        rman     106     7K     -      490 16, 32, 64
        sbuf       0     0K     -      28234 16, 32, 32768, 131072
...
lcc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue    1     1K      -      1     64
      entropy  1024    64K      -     1024    64
      linker   487   6272K     -     1163 16, 32, 64, 4096, 32768, 131072
      USB     127    10K      -      127 16, 32, 64, 128, 256, 1024, 2048
      lockf    23     2K      -    169585    64
      USBdev   10     2K      -       34 16, 128, 2048, 16384
      devbuf   5128  10760K     -     5310
16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      temp    1285    151K     -     10770
16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      ip6ndp    0     0K      -        4     64
      iflogical 20     5K      -       29  2048
      iffamilly 45     6K      -       69 32, 1024, 2048
      rtnexthop 189    29K     -    1211988 32, 256, 512, 1024, 2048, 4096
      metrics   11     2K      -       16   256
      rnode    135     3K      -    606391 16, 32
      rcache     4     8K      -        4  65536
      iflist     0     0K      -        6  16, 64
      ifdevice  11     8K      -       17 16, 32768
      ifstat    412   471K     -      415 512, 16384, 65536
      ipfw      42    23K      -        91
16, 32, 64, 128, 256, 512, 1024, 16384, 32768, 65536, 131072
      ifmaddr   415    11K     -       415 16, 32
      rtable    225    20K     -    606584 16, 32, 64, 128, 1024, 16384
      sysctl     0     0K      -    2302479 16, 32, 64
      ifaddr    53     4K      -        69 32, 64, 128
      mkey     133     3K      -     8974 16, 128
      pfe_ipc    0     0K      -    19035108
16, 32, 64, 128, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072

```

```

ifstate 710270 42176K - 9583703
16,32,64,128,256,512,1024,2048,8192,16384,32768
idxbucket 1 1K - 1 16
itable16 5045 1245K - 1825178 1024,4096
itable32 157 10K - 157 64
itable64 2 1K - 2 128
lr 1 1K - 4 16384
pic 37 6K - 37 64,16384
pfestat 0 0K - 1682 32,64,128,256,131072
gencfg 1486 424K - 2812 16,32,64,256,512,16384,32768,65536
jsr 0 0K - 22 16
idl 0 0K - 4 32768,131072
rtsmsg 0 0K - 3 131072
module 250 16K - 250 64,128
mtx_pool 1 8K - 1 64,128
DEVFS3 108 12K - 109 256
DEVFS1 101 23K - 101 2048
pgrp 5 1K - 917 64
session 5 1K - 917 512
proc 2 1K - 2 16384
subproc 217 441K - 4867 2048,131072
cred 21 3K - 48719 256
plimit 9 2K - 5255 2048
uidinfo 2 1K - 2 32,512
sysctluid 2786 85K - 2786 16,32,64
sysctltmp 0 0K - 1833 16,32,64,1024
umtx 126 8K - 126 64
SWAP 2 277K - 2 64
bus 780 125K - 2734 16,32,64,128,32768
bus-sc 69 69K - 1194
16,32,64,512,1024,2048,8192,16384,65536,131072
devstat 8 17K - 8 16,131072
eventhandler 45 2K - 46 32,128
kobj 93 186K - 111 65536
DEVFS 8 1K - 9 16,64
rman 94 6K - 477 16,32,64
sbuf 0 0K - 532 16,32,32768,131072
NULLFS hash 1 1K - 1 64
taskqueue 5 1K - 5 64
turnstiles 127 8K - 127 64
Unitno 6 1K - 44 16,64
ioctlops 0 0K - 1771718 16,32,64,128,8192,16384,65536,131072

iov 0 0K - 79425 16,64,128,256,512,1024,2048,131072
msg 4 25K - 4 32768,131072
sem 4 7K - 4 16384,32768,131072
shm 2 13K - 4 32768
ttys 93 16K - 195 512,32768
soname 31 3K - 389284 16,32,64,256
pcb 101 16K - 4374
16,32,64,128,1024,2048,4096,16384,65536
BIO buffer 40 80K - 750 65536
vfscache 1 512K - 1 65536
cluster_save buffer 0 OK - 55 32,64
VFS hash 1 256K - 1 32,64
vnodes 1 1K - 1 512
mount 266 21K - 481 16,32,64,128,256,4096,32768
vnodemarker 0 0K - 2497 16384
pfs_nodes 25 3K - 25 128
pfs_vncache 144 5K - 386 32
STP 1 1K - 1 64

```

```

      GEOM      173      15K      -      1068
16,32,64,128,256,512,2048,16384,32768,131072
      syncache    1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      tlv_stat     0      0K      -      223
16,32,64,128,256,512,2048,16384,32768,131072
      NFS daemon   1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      p1003.1b     1      1K      -      1  16
      MD disk      9     18K      -      9 65536
      ata_generic   2      2K      -      25 16,16384,32768
      ISOFS mount    7      1K      -      13  512
      ISOFS node 1439    135K      -    1453 128
      CAM SIM       1      1K      -      1  64
      CAM XPT       6      1K      -      9 16,64,16384
      CAM periph    1      1K      -      1 128
      ad_driver     2      1K      -      2  256
      pagedep       1     64K      -     105  64
      inodedep      1    256K      -     552 256
      newblk        1      1K      -     327 64,4096
      bmsafemap     0      0K      -      19  64
      allocdirect   0      0K      -     326 128
      freefrag      0      0K      -      31  32
      freeblks      0      0K      -     103 2048
      freefile      0      0K      -     175  32
      diradd        0      0K      -     590  64
      mkdir         0      0K      -     166  32
      dirrem        0      0K      -     382  32
      savedino      0      0K      -     283 512
      UFS mount     15     36K      -      15 2048,65536,131072
      ata_dma       6      1K      -      6  256
      UMAHash       1      4K      -      5 4096,16384,32768,65536,131072
      cdev          26     3K      -      26  256
      file desc    111    25K      -    5199 16,1024,2048,16384
      VM pgdata     2     65K      -      2  64
      sigio         1      1K      -      27  32
      kenv          30     5K      -      33 16,32,64,131072
      atkbddev      2      1K      -      2  32
      kqueue        0      0K      -      88 1024,4096,32768
      proc-args     28     2K      -    3970 32,64,128,256,512,1024
      isadev        23     2K      -      23  64
      zombie        1      1K      -    4651 128
      ithread       92     7K      -      92 16,64,256
      legacydrv     3      1K      -      3  16
      memdesc       1      4K      -      1 131072
      nexusdev      2      1K      -      2  16
      CAM queue     3      1K      -      3  16
      KTRACE        100    10K      -     100 128
      kbdmux        5      9K      -      5 128,2048,65536,131072
ITEM      SIZE      LIMIT      USED      FREE  REQUESTS
UMA Kegs: 136,      0,      71,      1,      71
...
```

show system virtual-memory | display xml

```

user@host> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.2R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
    </vmstat-memstat-malloc>
  </system-virtual-memory-information>
</rpc-reply>
```

```
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>entropy</memstat-name>
<inuse>1024</inuse>
<memuse>64</memuse>
<high-use>--</high-use>
<memstat-req>1024</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>linker</memstat-name>
<inuse>481</inuse>
<memuse>1871</memuse>
<high-use>--</high-use>
<memstat-req>1145</memstat-req>
<memstat-size>16,32,64,4096,32768,131072</memstat-size>
<memstat-name>lockf</memstat-name>
<inuse>56</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>5998</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>devbuf</memstat-name>
<inuse>2094</inuse>
<memuse>3877</memuse>
<high-use>--</high-use>
<memstat-req>2099</memstat-req>

<memstat-size>16,32,64,128,512,1024,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>temp</memstat-name>
<inuse>21</inuse>
<memuse>66</memuse>
<high-use>--</high-use>
<memstat-req>3127</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ip6ndp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6ifmulti</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6grentry</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
```

```

<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

```

```
<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifmaddr</memstat-name>
  <inuse>103</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>103</memstat-req>
  <memstat-size>16,32</memstat-size>
  <memstat-name>rtable</memstat-name>
  <inuse>129</inuse>
  <memuse>14</memuse>
  <high-use>--</high-use>
  <memstat-req>139</memstat-req>
  <memstat-size>16,32,64,128,1024,16384</memstat-size>
  <memstat-name>sysctl</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>14847</memstat-req>
  <memstat-size>16,32,64,4096,16384,32768</memstat-size>
  <memstat-name>ifaddr</memstat-name>
  <inuse>29</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>29</memstat-req>
  <memstat-size>64,128</memstat-size>
  <memstat-name>mkey</memstat-name>
  <inuse>345</inuse>
  <memuse>6</memuse>
  <high-use>--</high-use>
  <memstat-req>2527</memstat-req>
  <memstat-size>16,128</memstat-size>
  <memstat-name>pfe_ipc</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifstate</memstat-name>
  <inuse>594</inuse>
  <memuse>51</memuse>
  <high-use>--</high-use>
  <memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
  <memstat-name>itable16</memstat-name>
  <inuse>276</inuse>
  <memuse>52</memuse>
  <high-use>--</high-use>
  <memstat-req>294</memstat-req>
  <memstat-size>1024,4096</memstat-size>
  <memstat-name>itable32</memstat-name>
  <inuse>160</inuse>
  <memuse>10</memuse>
  <high-use>--</high-use>
  <memstat-req>160</memstat-req>
  <memstat-size>64</memstat-size>
  <memstat-name>itable64</memstat-name>
  <inuse>2</inuse>
  <memuse>1</memuse>
```



```

<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>

```

```
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctlpid</memstat-name>
<inuse>1117</inuse>
<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```

```

<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>

```

```
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>ioctlops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
```

```

<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>

```

```
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD_disk</memstat-name>
<inuse>10</inuse>
```

```

<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISOFs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISOFs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```



```

<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>

```

```
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
```

```

<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>
  <used>30</used>
  <free>20</free>
  <zone-req>30</zone-req>
  <zone-name>32 Bucket:</zone-name>
  <zone-size>140</zone-size>
  <count-limit>0</count-limit>
  <used>33</used>
  <free>23</free>

```

```
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
```

```
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
```

```
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
```

```

<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL UMA zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>

```

```
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
```



```
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
```

```

    <zone-req>0</zone-req>
    <zone-name>ripcb:</zone-name>
    <zone-size>232</zone-size>
    <count-limit>25194</count-limit>
    <used>5</used>
    <free>29</free>
    <zone-req>5</zone-req>
    <zone-name>SWAPMETA:</zone-name>
    <zone-size>276</zone-size>
    <count-limit>94948</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
    <zone-name>FFS inode:</zone-name>
    <zone-size>132</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>72</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS1 dinode:</zone-name>
    <zone-size>128</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>24</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS2 dinode:</zone-name>
    <zone-size>256</zone-size>
    <count-limit>0</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
  </vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>
  <copy-on-write-faults>48364</copy-on-write-faults>
  <copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
  <zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
  <zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
  <transit-blocking-page-faults>85</transit-blocking-page-faults>
  <total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>

```

```

    <pages-affected-by-fork>95343</pages-affected-by-fork>
    <pages-affected-by-vfork>3526</pages-affected-by-vfork>
    <pages-affected-by-rfork>0</pages-affected-by-rfork>
    <pages-freed>221502</pages-freed>
    <pages-freed-by-daemon>0</pages-freed-by-daemon>
    <pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
    <pages-active>45826</pages-active>
    <pages-inactive>13227</pages-inactive>
    <pages-in-vm-cache>49278</pages-in-vm-cache>
    <pages-wired-down>10640</pages-wired-down>
    <pages-free>70706</pages-free>
    <bytes-per-page>4096</bytes-per-page>
    <swap-pages-used>0</swap-pages-used>
    <peak-swap-pages-used>0</peak-swap-pages-used>
    <total-name-lookups>214496</total-name-lookups>
    <positive-cache-hits>92</positive-cache-hits>
    <negative-cache-hits>5</negative-cache-hits>
    <pass2>0</pass2>
    <cache-deletions>0</cache-deletions>
    <cache-falsehits>0</cache-falsehits>
    <toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
    <intr-name>irq0: clk          </intr-name>
    <intr-cnt>1243455</intr-cnt>
    <intr-rate>999</intr-rate>
    <intr-name>irq4: sio0        </intr-name>
    <intr-cnt>1140</intr-cnt>
    <intr-rate>0</intr-rate>
    <intr-name>irq8: rtc         </intr-name>
    <intr-cnt>159164</intr-cnt>
    <intr-rate>127</intr-rate>
    <intr-name>irq9: cbb1 fxp0   </intr-name>
    <intr-cnt>28490</intr-cnt>
    <intr-rate>22</intr-rate>
    <intr-name>irq10: fxp1       </intr-name>
    <intr-cnt>20593</intr-cnt>
    <intr-rate>16</intr-rate>
    <intr-name>irq14: ata0       </intr-name>
    <intr-cnt>5031</intr-cnt>
    <intr-rate>4</intr-rate>
    <intr-name>Total</intr-name>
    <intr-cnt>1457873</intr-cnt>
    <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
    <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
<kernel-direct-mm-size-information>
    <vm-directmm-size-used>4644</vm-directmm-size-used>
    <vm-directmm-size-max>4057334</vm-directmm-size-max>
</kernel-direct-mm-size-information>
</system-virtual-memory-information>
<cli>
    <banner></banner>
</cli>
</rpc-reply>

```

Note: <kernel-direct-mm-size-information> only displays for 64 bit platform.

show system virtual-memory (QFX Series)

```

user@switch> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>entropy</memstat-name>
      <inuse>1024</inuse>
      <memuse>64</memuse>
      <high-use>-</high-use>
      <memstat-req>1024</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>linker</memstat-name>
      <inuse>481</inuse>
      <memuse>1871</memuse>
      <high-use>-</high-use>
      <memstat-req>1145</memstat-req>
      <memstat-size>16, 32, 64, 4096, 32768, 131072</memstat-size>
      <memstat-name>lockf</memstat-name>
      <inuse>56</inuse>
      <memuse>4</memuse>
      <high-use>-</high-use>
      <memstat-req>5998</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>devbuf</memstat-name>
      <inuse>2094</inuse>
      <memuse>3877</memuse>
      <high-use>-</high-use>
      <memstat-req>2099</memstat-req>

      <memstat-size>16, 32, 64, 128, 512, 1024, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>temp</memstat-name>
      <inuse>21</inuse>
      <memuse>66</memuse>
      <high-use>-</high-use>
      <memstat-req>3127</memstat-req>

      <memstat-size>16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>ip6ndp</memstat-name>
      <inuse>0</inuse>
      <memuse>0</memuse>
      <high-use>-</high-use>
      <memstat-req>4</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6ifmulti</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6grentry</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>

```

```

<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>

```

```
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
<memstat-name>ifmaddr</memstat-name>
<inuse>103</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>103</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rtable</memstat-name>
<inuse>129</inuse>
<memuse>14</memuse>
<high-use>--</high-use>
<memstat-req>139</memstat-req>
<memstat-size>16,32,64,128,1024,16384</memstat-size>
<memstat-name>sysctl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>14847</memstat-req>
<memstat-size>16,32,64,4096,16384,32768</memstat-size>
<memstat-name>ifaddr</memstat-name>
<inuse>29</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>29</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mkey</memstat-name>
<inuse>345</inuse>
<memuse>6</memuse>
<high-use>--</high-use>
<memstat-req>2527</memstat-req>
<memstat-size>16,128</memstat-size>
<memstat-name>pfe_ipc</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ifstate</memstat-name>
<inuse>594</inuse>
<memuse>51</memuse>
<high-use>--</high-use>
<memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
<memstat-name>itable16</memstat-name>
<inuse>276</inuse>
<memuse>52</memuse>
<high-use>--</high-use>
<memstat-req>294</memstat-req>
<memstat-size>1024,4096</memstat-size>
<memstat-name>itable32</memstat-name>
```

```

<inuse>160</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>160</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>itable64</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>

```

```
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctloid</memstat-name>
<inuse>1117</inuse>
```



```

<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>

```

```
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
```

```

<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>
<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
```

```

<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD disk</memstat-name>
<inuse>10</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISofs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISofs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>

```

```
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
```

```

<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
```



```

<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>

```

```
<used>30</used>
<free>20</free>
<zone-req>30</zone-req>
<zone-name>32 Bucket:</zone-name>
<zone-size>140</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>23</free>
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
```

```
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
```

```
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
```

```

<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL_UMA_zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>

```

```
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
```

```
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
```

```

<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>24</free>
<zone-req>1306</zone-req>
<zone-name>FFS2 dinode:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
</vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>

```



```

<copy-on-write-faults>48364</copy-on-write-faults>
<copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
<zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
<zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
<transit-blocking-page-faults>85</transit-blocking-page-faults>
<total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>
<pages-affected-by-fork>95343</pages-affected-by-fork>
<pages-affected-by-vfork>3526</pages-affected-by-vfork>
<pages-affected-by-rfork>0</pages-affected-by-rfork>
<pages-freed>221502</pages-freed>
<pages-freed-by-daemon>0</pages-freed-by-daemon>
<pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
<pages-active>45826</pages-active>
<pages-inactive>13227</pages-inactive>
<pages-in-vm-cache>49278</pages-in-vm-cache>
<pages-wired-down>10640</pages-wired-down>
<pages-free>70706</pages-free>
<bytes-per-page>4096</bytes-per-page>
<swap-pages-used>0</swap-pages-used>
<peak-swap-pages-used>0</peak-swap-pages-used>
<total-name-lookups>214496</total-name-lookups>
<positive-cache-hits>92</positive-cache-hits>
<negative-cache-hits>5</negative-cache-hits>
<pass2>0</pass2>
<cache-deletions>0</cache-deletions>
<cache-falsehits>0</cache-falsehits>
<toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0    </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```


show version

List of Syntax	Syntax on page 1135 Syntax (EX Series Switches) on page 1135 Syntax (TX Matrix Router) on page 1135 Syntax (TX Matrix Plus Router) on page 1135 Syntax (MX Series Router) on page 1135 Syntax (QFX Series) on page 1135
Syntax	<pre>show version <brief detail></pre>
Syntax (EX Series Switches)	<pre>show version <all-members> <brief detail> <local> <member member-id></pre>
Syntax (TX Matrix Router)	<pre>show version <brief detail> <all-chassis all-lcc lcc number scc></pre>
Syntax (TX Matrix Plus Router)	<pre>show version <all-chassis all-lcc lcc number sfc number> <brief detail></pre>
Syntax (MX Series Router)	<pre>show version <brief detail> <all-members> <local> <member member-id></pre>
Syntax (QFX Series)	<pre>show version <brief detail> <component component-name all></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display the hostname and version information about the software running on the router or switch.
Options	<p>none—Display standard information about the hostname and version of the software running on the router or switch.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>all-members—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.</p>

component all—(QFabric systems only) (Optional) Display the host name and version information about the software running on all the components on the QFabric system.

component *component-name*—(QFabric systems only) (Optional) Display the host name and version information about the software running on a specific QFabric system component. Replace *component-name* with the name of the QFabric system component. The *component-name* can be the name of a diagnostics Routing Engine, Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect device, or Node group.

local—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

scc—(TX Matrix routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix router (or switch-card chassis).

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the host name and version information about the software running on for a specified T640 router (line-card chassis or LCC) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the host name and version information about the software running for a specified T1600 or T4000 router (LCC) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc *number*—(TX Matrix Plus routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information By default, when you issue the **show version** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same

command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show version on page 1138](#)
[show version \(TX Matrix Plus Router\) on page 1138](#)
[show version \(TX Matrix Plus Router with 3D SIBs\) on page 1141](#)
[show version \(MX Series Router\) on page 1144](#)
[show version \(QFX3500 Switch\) on page 1144](#)
[show version \(QFabric System\) on page 1145](#)
[show version component all \(QFabric System\) on page 1145](#)

Sample Output

show version

```
user@host> show version
Hostname: router1
Model: m20
JUNOS Base OS boot [7.2-20050312.0]
JUNOS Base OS Software Suite [7.2-20050312.0]
JUNOS Kernel Software Suite [7.2R1.7]
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]
JUNOS Routing Software Suite [7.2R1.7]
JUNOS Online Documentation [7.2R1.7]
JUNOS Crypto Software Suite [7.2R1.7]

{master}

user@host> show version psd 1
psd1-re0:
-----
Hostname: china
Model: t640
JUNOS Base OS boot [9.1I20080311_1959_builder]
JUNOS Base OS Software Suite [9.1-20080321.0]
JUNOS Kernel Software Suite [9.1-20080321.0]
JUNOS Crypto Software Suite [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (T-series) [9.1-20080321.0]
JUNOS Online Documentation [9.1-20080321.0]
JUNOS Routing Software Suite [9.1-20080321.0]
labpkg [7.0]
```

show version (TX Matrix Plus Router)

```
user@host> show version
sfc0-re0:
-----
Hostname: host
Model: txp
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services ACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
```

```
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]
```

```
lcc0-re0:
```

```
-----
Hostname: host1
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]
```

```
lcc1-re0:
```

```
-----
Hostname: host2
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
```

JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

lcc2-re0:

Hostname: host3
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

lcc3-re0:

Hostname: host4
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]


```

JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

show version (TX Matrix Plus Router with 3D SIBs)

```

user@host>show version
sfc0-re0:
-----
Hostname: sfc0
Model: txp
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

lcc0-re0:
-----
Hostname: lcc0
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]

```

JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package [13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

lcc2-re0:

Hostname: lcc2
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package [13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

lcc4-re0:

```

-----
Hostname: tcc4
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

```

tcc6-re0:

```

-----
Hostname: tcc6
Model: t1600
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]

```

```
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc7-re0:
```

```
-----
Hostname: lcc7
Model: t1600
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

show version (MX Series Router)

```
user@host5> show version
Hostname: host5
Model: mx80
JUNOS Base OS boot [11.3-20110717.0]
JUNOS Base OS Software Suite [11.3-20110717.0]
JUNOS Kernel Software Suite [11.3-20110717.0]
JUNOS Crypto Software Suite [11.3-20110717.0]
JUNOS Packet Forwarding Engine Support (MX80) [11.3-20110717.0]
JUNOS Online Documentation [11.3-20110717.0]
JUNOS Routing Software Suite [11.3-20110717.0]
```

show version (QFX3500 Switch)

```
user@switch> show version
```

```

Hostname: switch
Model: qfx_s3500
JUNOS Base OS boot [11.1R1]
JUNOS Base OS Software Suite [11.1R1]
JUNOS Kernel Software Suite [11.1R1]
JUNOS Crypto Software Suite [11.1R1]
JUNOS Online Documentation [11.1R1]
JUNOS Enterprise Software Suite [11.1R1]
JUNOS Packet Forwarding Engine Support (QFX) [11.1R1]
JUNOS Routing Software Suite [11.1R1]

```

show version (QFabric System)

```

user@qfabric> show version
Hostname: qfabric
Model: qfx3000-g
Serial Number: qfsn-0123456789
QFabric System ID: f158527a-f99e-11e0-9fbd-00e081c57cda
JUNOS Base Version [12.2I20111018_0215_dc-builder]

```

show version component all (QFabric System)

```

user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

```

```
FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
DRE-0:
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```


```
FM-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
nodedevice1:
-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
interconnectdevice1:
-
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
```

```
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]  
JUNOS Routing Software Suite [11.3R1.6]  
warning: from interconnectdevice0: Disconnected
```

start shell

Syntax	<code>start shell (csh sh)</code> <code><user username></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type exit from the shell.
<div> NOTE:<ul style="list-style-type: none">To issue this command, the user must have the required login access privileges configured by including the permissions statement at the [edit system login class class-name] hierarchy level.UNIX wheel group membership or permissions are no longer required to issue this command.</div>	
Options	csh —Create a UNIX C shell. sh —Create a UNIX Bourne shell. user username —(Optional) Start the shell as another user.
Additional Information	When you are in the shell, the shell prompt has the following format: <code>username@hostname%</code> An example of the prompt is: <code>root@host%</code>
Required Privilege Level	shell and maintenance
List of Sample Output	start shell csh on page 1148
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

start shell csh

```
user@host> start shell csh
%
exit
%
```



```
username@hostname% start shell sh
%

exit
user@host>
```

test configuration

Syntax	<code>test configuration filename</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found.
Options	filename —Name of the configuration file. syntax-only —Check the syntax of a partial configuration file, without checking for commit errors. This option introduced in Junos OS Release 12.1.
Required Privilege Level	view
List of Sample Output	test configuration on page 1150
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed
```

traceroute

List of Syntax [Syntax on page 1151](#)
 [Syntax \(QFX Series\) on page 1151](#)

Syntax `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<clns>`
 `<gateway address>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<logical system logical-system-name>`
 `<monitor host>`
 `<mpls (ldp FEC address | rsvp label-switched-path-name)>`
 `<no-resolve>`
 `<propagate-ttl>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

Syntax (QFX Series) `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<gateway address>`
 `<inet>`
 `<interface interface-name>`
 `<monitor host>`
 `<no-resolve>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 mpls option introduced in Junos OS Release 9.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 propagate-ttl option introduced in Junos OS Release 12.1.

Description Display the route that packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

Options **host**—IP address or name of remote host.

as-number-lookup—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached

network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

clns—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

gateway address—(Optional) Address of a router or switch through which the route transits.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interface *interface-name*—(Optional) Name of the interface over which to send packets.

logical-system *logical-system-name*—(Optional) Perform this operation on all logical systems or on a particular logical system.

monitor *host*—(Optional) Display real-time monitoring information for the specified host.

mpls (*ldp FEC address | rsvp label-switched-path name*)—(Optional) See *traceroute mpls ldp* and *traceroute mpls rsvp*.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

propagate-ttl—(Optional) On the PE router, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only. Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



NOTE: Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *source-address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

Related Documentation • [traceroute monitor on page 1155](#)

List of Sample Output [traceroute on page 1153](#)[traceroute as-number-lookup host on page 1153](#)[traceroute no-resolve on page 1153](#)[traceroute propagate-ttl on page 1154](#)[traceroute \(Between CE Routers, Layer 3 VPN\) on page 1154](#)[traceroute \(Through an MPLS LSP\) on page 1154](#)

Output Fields [Table 95 on page 1153](#) describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 95: traceroute Output Fields

Field Name	Field Description
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.
<i>number-of-hops</i>	Number of hops from the source to the named router or switch.
<i>router-name</i>	Name of the router or switch for this hop.
<i>address</i>	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

Sample Output**traceroute**

```

user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254)  2.370 ms  2.853 ms  0.367 ms
 2 red14 (10.168.255.250) 0.778 ms  2.937 ms  0.446 ms
 3 yellow (10.156.169.254) 7.737 ms  89.905 ms  0.834 ms

```

traceroute as-number-lookup host

```

user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms 0.728 ms 0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms 0.611 ms 0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms 0.808 ms 0.774 ms

```

traceroute no-resolve

```

user@host> traceroute santacruz no-resolve
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 10.168.1.254 0.458 ms 0.370 ms 0.365 ms
 2 10.168.255.250 0.474 ms 0.450 ms 0.444 ms
 3 10.156.169.254 0.931 ms 0.876 ms 0.862 ms

```

traceroute propagate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
    MPLS Label=299776 CoS=0 TTL=1 S=0
    MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
byte packets
 1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
 2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
    MPLS Label=100006 CoS=0 TTL=1 S=1
 3  vpn09.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
    MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

traceroute monitor

List of Syntax	Syntax on page 1155 Syntax (QFX Series) on page 1155
Syntax	<pre>traceroute monitor <i>host</i> <count <i>value</i>> <inet inet 6> <interval <i>seconds</i>> <no resolve> <size <i>value</i>> <source <i>source-address</i>> <summary></pre>
Syntax (QFX Series)	<pre>traceroute monitor <i>host</i> <count <i>value</i>> <inet> <interval <i>seconds</i>> <no resolve> <size <i>value</i>> <source <i>source-address</i>> <summary></pre>
Release Information	<p>Command introduced in Junos OS Release 8.0</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.
Options	<p><i>host</i>—IP address or name of remote host.</p> <p><i>count value</i>—Number of ping requests, in packets, to send in summary mode. The default value is 10.</p> <p><i>inet inet6</i>—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.</p> <p><i>interval seconds</i>—(Optional) Number of seconds to wait before sending ping requests. The default value is 1.</p> <p><i>no resolve</i>—(Optional) Do not attempt to display addresses symbolically.</p> <p><i>size value</i>—(Optional) Receive the specified number of bytes for each packet. The range is 0 through 65468 bytes. The default value is 64.</p> <p><i>source source-address</i>—(Optional) Source address of the outgoing ping packets.</p> <p><i>summary</i>—(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.</p>
Required Privilege Level	network
List of Sample Output	traceroute monitor on page 1156

Output Fields Table 96 on page 1156 describes the output fields for the **traceroute monitor** command. Output fields are listed in the approximate order in which they appear.

Table 96: traceroute monitor Output Fields

Field Name	Field Description
Host	Hostname or IP address of the router at each hop.
Loss%	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
Snt	Number of ping requests sent to the router at this hop.
Last	Most recent round-trip time, in milliseconds, to the router at this hop.
Avg	Average round-trip time, in milliseconds, to the router at this hop.
Best	Shortest round-trip time, in milliseconds, to the router at this hop.
Wrst	Longest round-trip time, in milliseconds, to the router at this hop.
StDev	Standard deviation of round-trip times, in milliseconds, to the router at this hop.

Sample Output

traceroute monitor

```
user@host> traceroute monitor 10.16.0.1
```

	Loss%	Snt	Last	Avg	Best	Wrst	StDev
Host							
1. 10.17.41.254	0.0%	17	0.7	1.0	0.6	5.4	1.2
2. secret.net	0.0%	17	0.6	1.0	0.6	6.6	1.4
3. top-secret.net	0.0%	17	0.6	0.6	0.6	0.6	0.0

CHAPTER 8

Troubleshooting

- [Troubleshooting Procedures on page 1157](#)

Troubleshooting Procedures

- [Rebooting and Halting a QFX Series Product on page 1157](#)
- [Recovering from a Failed Software Installation on page 1158](#)
- [Recovering the Root Password on page 1159](#)
- [Troubleshooting Network Interfaces on page 1160](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 1161](#)

Rebooting and Halting a QFX Series Product

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

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>      Execute this command
at             Time at which to perform the operation
in            Number of minutes to delay before operation
media         Boot media for next boot
message       Message to display to all users
|            Pipe through a command

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

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



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

```
user@switch> request system halt ?
Possible completions:
<[Enter]>      Execute this command
at             Time at which to perform the operation
in            Number of minutes to delay before operation
media         Boot media for next boot
message       Message to display to all users
|            Pipe through a command
```



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

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

Related Documentation

- [clear system reboot on page 332](#)
- [request system reboot on page 390](#)
- [request system halt on page 375](#)
- [request system power-off on page 385](#)
- [Connecting a QFX Series Device to a Management Console](#)

Recovering from a Failed Software Installation

Problem **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

Solution If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message `Loading /boot/defaults/loader.conf` appears, you are prompted with the following message:

Hit [Enter] to boot immediately, or space bar for command prompt.

Press the Spacebar to enter the manual loader. The `loader>` prompt appears.

3. Enter the following command:

```
loader> install [--format] [--external] source
where:
```

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).

- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
 - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
 - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

Recovering the Root Password

If you forget the root password for the QFX3500 switch, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the QFX3500 switch power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the switch into the RJ-45-to-DB-9 serial port adapter supplied with the switch.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the switch.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None

- Stop bits: 1
 - Flow control: None
9. Power on the switch by (if necessary) plugging the power cords into the QFX3500 switch power supply, or turning on the power to the device or switch by switching on the AC power outlet the device is plugged into

The terminal emulation screen on your management device displays the switch's boot sequence.
 10. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
 11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

ok **boot -s**
 12. At the following prompt, enter **recovery** to start the root password recovery procedure.

Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: **recovery**
 13. Enter configuration mode in the CLI.
 14. Set the root password. For example:

user@switch# **set system root-authentication plain-text-password**
 15. At the following prompt, enter the new root password. For example:

New password: **juniper1**
Retype new password:
 16. At the second prompt, reenter the new root password.
 17. After you have finished configuring the password, commit the configuration.

root@host# **commit**
commit complete
 18. Exit configuration mode in the CLI.
 19. Exit operational mode in the CLI.
 20. At the prompt, enter **y** to reboot the switch.

Reboot the system? [y/n] **y**

Related Documentation • [Configuring the Root Password on page 1702](#)

Troubleshooting Network Interfaces

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

Problem	Description: The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.
----------------	---

Symptoms: When you check the status with the CLI command **show interfaces *interface-name*** , the disabled port is not listed.

Cause By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

Troubleshooting an Aggregated Ethernet Interface

Problem **Description:** The **show interfaces terse** command shows that the LAG is down.

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

Related Documentation

- [Verifying the Status of a LAG Interface on page 2630](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)

PART 5

QFabric System Deployment

- [Overview on page 1165](#)
- [Configuration on page 1255](#)
- [Administration on page 1405](#)
- [Troubleshooting on page 1573](#)

CHAPTER 9

Overview

- [Before You Begin on page 1165](#)
- [Hardware Architecture Overview on page 1177](#)
- [Software Architecture Overview on page 1192](#)
- [Software Features on page 1201](#)
- [Licenses on page 1250](#)

Before You Begin

- [QFabric System Overview on page 1165](#)
- [Understanding QFabric System Terminology on page 1169](#)
- [Understanding Interfaces on the QFabric System on page 1173](#)

QFabric System Overview

The architecture of legacy data centers contrasts significantly with the revolutionary Juniper Networks data center solution.


This topic covers:

- [Legacy Data Center Architecture on page 1165](#)
- [QFX Series QFabric System Architecture on page 1167](#)

Legacy Data Center Architecture

Service providers and companies that support data centers are familiar with legacy multi-tiered architectures, as seen in [Figure 13 on page 1166](#).

Figure 13: Legacy Data Center Architecture



The *access layer* connects servers and other devices to a Layer 2 switch and provides an entry point into the data center. Several access switches are in turn connected to intermediate Layer 2 switches at the *aggregation layer* (sometimes referred to as the *distribution layer*) to consolidate traffic. A *core layer* interconnects the aggregation layer switches. Finally, the core switches are connected to Layer 3 routers in the *routing layer* to send the aggregated data center traffic to other data centers or a wide area network (WAN), receive external traffic destined for the data center, and interconnect different Layer 2 broadcast domains within the data center.


The problems that exist with the multi-tiered data center architecture include:

- **Limited scalability**—The demands for electrical power, cooling, cabling, rack space, and port density increase exponentially as the traditional data center expands, which prohibits growth after minimal thresholds are met.
- **Inefficient resource usage**—Up to 50 percent of switch ports in a legacy data center are used to interconnect different tiers rather than support server and storage connections. In addition, traffic that ideally should move horizontally between servers within a data center often must also be sent vertically up through the tiers to reach a router and down through the tiers to reach the required destination server.
- **Increased latency**—By requiring the devices at each tier level to perform multiple iterations of packet and frame processing, the data plane traffic takes significantly longer to reach its destination than if the sending and receiving devices were directly connected. This processing overhead results in potentially poor performance for time-sensitive applications, such as voice, video, or financial transactions.

QFX Series QFabric System Architecture

In contrast to legacy multi-tiered data center architectures, the Juniper Networks QFX Series QFabric System architecture provides a simplified networking environment that solves the most challenging issues faced by data center operators. A fabric is a set of devices that act in concert to behave as a single switch. It is a highly scalable, distributed, Layer 2 and Layer 3 networking architecture that provides a high-performance, low-latency, and unified interconnect solution for next-generation data centers as seen in Figure 14 on page 1167.

Figure 14: QFX Series QFabric System Architecture



9041145

A QFabric system collapses the traditional multi-tiered data center model into a single tier where all access layer devices (known in the QFabric system model as *Node devices*) are essentially directly connected to all other access layer devices across a very large scale fabric backplane (known in the QFabric system model as the *Interconnect device*). Such an architecture enables the consolidation of data center endpoints (such as servers, storage devices, memory, appliances, and routers) and provides better scaling and network virtualization capabilities than traditional data centers.

Essentially, a QFabric system can be viewed as a single, nonblocking, low-latency switch that supports thousands of 10-Gigabit Ethernet ports or 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports to interconnect servers, storage, and the Internet across a high-speed, high-performance fabric. The entire QFabric system is managed as a single entity through a *Director group*, containing redundant hardware and software components that can be expanded and scaled as the QFabric system grows in size. In addition, the Director group automatically senses when devices are added or removed from the QFabric system and dynamically adjusts the amount of processing resources required to support the system. Such intelligence helps the QFabric system use the minimum amount of power to run the system efficiently, but not waste energy on unused components.

As a result of the QFabric system architecture, data center operators are now realizing the benefits of this next-generation architecture, including:

- **Low latency**—Because of its inherent advantages in this area, the QFabric system provides an excellent foundation for mission-critical applications such as financial transactions and stock trades, as well as time-sensitive applications such as voice and video.
- **Enhanced scalability**—The QFabric system can be managed as a single entity and provides support for thousands of data center devices. As Internet traffic continues to grow exponentially with the increase in high-quality video transmissions and rise in the number of mobile devices used worldwide, the QFabric system can keep pace with the demands for bandwidth, applications, and services offered by the data center.
- **Virtualization-enabled**—The QFabric system was designed to work seamlessly with virtual servers, virtual appliances, and other virtual devices, allowing for even greater scalability, expandability, and rapid deployment of new services than ever before. Migrating to virtual devices also results in significant costs savings, fueled by reduced space requirements, decreased needs for power and cooling, and increased processing capabilities.
- **Simplicity**—Although the QFabric system can scale to hundreds of devices and thousands of ports, you can still manage the QFabric system as a single system.
- **Flexibility**—You can deploy the QFabric system as an entire system or in stages.
- **Convergence**—Because the congestion-free fabric is lossless, all traffic in a QFabric system can be converged onto a single network. As a result, the QFabric system supports Ethernet, Fibre Channel over Ethernet, and native Fibre Channel packets and frames.

Flat, nonblocking, and lossless, the network fabric offered by the QFabric system has the scale and flexibility to meet the needs of small, medium, and large-sized data centers for years to come.

Related Documentation

- [Understanding QFabric System Terminology on page 1169](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)
- [Understanding the QFabric System Software Architecture on page 1192](#)

Understanding QFabric System Terminology

To understand the QFabric system environment and its components, you should become familiar with the terms defined in [Table 97 on page 1169](#).

Table 97: QFabric System Terms

Term	Definition
Clos network fabric	Three-stage switching network in which switch elements in the middle stages are connected to all switch elements in the ingress and egress stages. In the case of QFabric system components, the three stages are represented by an ingress chipset, a midplane chipset, and an egress chipset in an Interconnect device (such as a QFX3008-I Interconnect device). In Clos networks, which are well known for their nonblocking properties, a connection can be made from any idle input port to any idle output port, regardless of the traffic load in the rest of the system.
Director device	Hardware component that processes fundamental QFabric system applications and services, such as startup, maintenance, and inter-QFabric system device communication. A set of Director devices with hard drives can be joined to form a <i>Director group</i> , which provides redundancy and high availability by way of additional memory and processing power. (See also <i>Director group</i> .)
Director group	<p>Set of Director devices that host and load-balance internal processes for the QFabric system. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration, startup, and DNS, DHCP, and NFS services. Operating a Director group is a minimum requirement to manage a QFabric system.</p> <p>The Director group runs the Director software for management applications and runs dual processes in active/standby mode for maximum redundancy and high availability. (See also <i>Director software</i> and <i>Director device</i>.)</p>
Director software	Software that handles QFabric system administration tasks, such as fabric management and configuration. The Junos OS-based Director software runs on the <i>Director group</i> , provides a single, consolidated view of the QFabric system, and enables the main QFabric system administrator to configure, manage, monitor, and troubleshoot QFabric system components from a centralized location. To access the Director software, log in to the default partition. (See also <i>Director device</i> and <i>Director group</i> .)
fabric control Routing Engine	Virtual Junos OS Routing Engine instance used to control the exchange of routes and flow of data between QFabric system hardware components within a partition. The fabric control Routing Engine runs on the Director group.
fabric manager Routing Engine	Virtual Junos OS Routing Engine instance used to control the initialization and maintenance of QFabric system hardware components belonging to the default partition. The fabric manager Routing Engine runs on the Director group.
infrastructure	QFabric system services processed by the virtual Junos Routing Engines operating within the Director group. These services, such as fabric management and fabric control, support QFabric system functionality and high availability.

Table 97: QFabric System Terms (*continued*)

Term	Definition
Interconnect device	QFabric system component that acts as the primary fabric for data plane traffic traversing the QFabric system between Node devices. An example of an Interconnect device is a QFX3008-I Interconnect device. (See also <i>Node device</i> .)
Junos Space	Carrier-class network management system for provisioning, monitoring, and diagnosing Juniper Networks routing, switching, security, and data center platforms.
network Node group	Set of one to eight Node devices that connects to an external network.
network Node group Routing Engine	Virtual Junos OS Routing Engine instance that handles routing processes for a network Node group. The network Node group Routing Engine runs on the Director group.
Node device	Routing and switching device that connects to endpoints (such as servers or storage devices) or external network peers, and is connected to the QFabric system through an Interconnect device. You can deploy Node devices similarly to the way a top-of-rack switch is implemented. An example of a Node device is the QFX3500 Node device. (See also <i>Interconnect device</i> and <i>network Node group</i> .)
partition	<p>Collection of physical or logical QFabric system hardware components (such as Node devices) that provides fault isolation, separation, and security.</p> <p>In their initial state, all QFabric system components belong to a <i>default partition</i>.</p>
QFabric system	<p>Highly scalable, distributed, Layer 2 and Layer 3 networking architecture that provides a high-performance, low-latency, and unified interconnect solution for next-generation data centers. A QFabric system collapses the traditional multi-tier data center model, enables the consolidation of data center endpoints (such as servers, storage devices, memory, appliances, and routers), and provides better scaling and network virtualization capabilities than traditional data centers.</p> <p>Essentially, a QFabric system can be viewed as a single, nonblocking, low-latency switch that supports thousands of 10-Gigabit Ethernet ports or 2-Gbps, 4-Gbps or 8-Gbps Fibre Channel ports to interconnect servers, storage, and the Internet across a high-speed, high-performance fabric. The QFabric system must have sufficient resources and devices allocated to handle the <i>Director group</i>, <i>Node device</i>, and <i>Interconnect device</i> functions and capabilities.</p>

Table 97: QFabric System Terms (*continued*)

Term	Definition
QFabric system control plane	<p>Internal network connection that carries control traffic between QFabric system components. The QFabric system control plane includes management connections between the following QFabric system hardware and software components:</p> <ul style="list-style-type: none"> • <i>Node devices</i>, such as the QFX3500 Node device. • <i>Interconnect devices</i>, such as the QFX3008-I Interconnect device. • <i>Director group processes</i>, such as management applications, provisioning, and topology discovery. • <i>Control plane Ethernet switches</i> to provide interconnections to all QFabric system devices and processes. For example, you can use EX Series EX4200 switches running in Virtual Chassis mode for this purpose. <p>To maintain high availability, the QFabric system control plane uses a different network than the QFabric system data plane, and uses a fabric provisioning protocol and a fabric management protocol to establish and maintain the QFabric system.</p>
QFabric system data plane	<p>Redundant, high-performance, and scalable data plane that carries QFabric system data traffic. The QFabric system data plane includes the following high-speed data connections:</p> <ul style="list-style-type: none"> • 10-Gigabit Ethernet connections between QFabric system endpoints (such as servers or storage devices) and Node devices. • 40-Gbps quad small form-factor pluggable plus (QSFP+) connections between Node devices and Interconnect devices. • 10-Gigabit Ethernet connections between external networks and a Node device acting as a network Node group. <p>To maintain high availability, the QFabric system data plane is separate from the QFabric system control plane.</p>
QFabric system endpoint	Device connected to a Node device port, such as a server, a storage device, memory, an appliance, a switch, or a router.
QFabric system fabric	Distributed, multistage network that consists of a queuing and scheduling system that is implemented in the Node device, and a distributed cross-connect system that is implemented in Interconnect devices. The QFabric system fabric is part of the QFabric system data plane.
QFX3500 Node device	<p>Node device that connects to either endpoint systems (such as servers and storage devices) or external networks in a QFabric system. It is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure.</p> <p>The QFX3500 Node device provides up to 48 10-Gigabit Ethernet interfaces to connect to the endpoints. Twelve of these 48 interfaces can be configured to support 2-Gbps, 4-Gbps or 8-Gbps Fibre Channel, and 36 of the interfaces can be configured to support Gigabit Ethernet. Also, there are four uplink connections to connect to Interconnect devices in a QFabric system. These uplinks use 40-Gbps quad small form-factor pluggable plus (QSFP+) interfaces. (See also <i>QFX3500 switch</i>.)</p>

Table 97: QFabric System Terms (*continued*)

Term	Definition
QFX3500 switch	<p>Standalone data center switch with 10-Gigabit Ethernet access ports and 40-Gbps quad, small form-factor pluggable plus (QSFP+) uplink interfaces. You can (optionally) configure some of the access ports as 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports or Gigabit Ethernet ports.</p> <p>The QFX3500 switch can be converted to a QFabric system Node device as part of a complete QFabric system. The switch is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure. (See also <i>QFX3500 Node device</i>.)</p>
QFX3600 Node device	<p>Node device that connects to either endpoint systems (such as servers and storage devices) or external networks in a QFabric system. It is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure.</p> <p>The QFX3600 Node device provides 16 40-Gbps QSFP+ ports. By default, 4 ports (labeled Q0 through Q3) are configured for 40-Gbps uplink connections between your Node device and your Interconnect device, and 12 ports (labeled Q4 through Q15) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect device, and ports Q2 through Q15 for 10-Gigabit Ethernet connections to either endpoint systems or external networks. (See also <i>QFX3600 switch</i>.)</p>
QFX3600 switch	<p>Standalone data center switch with 16 40-Gbps quad, small form-factor pluggable plus (QSFP+) interfaces. By default, all the 16 ports operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the 40-Gbps ports to operate as four 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.</p> <p>The QFX3600 switch can be converted to a QFabric system Node device as part of a complete QFabric system. The switch is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure. (See also <i>QFX3600 Node device</i>.)</p>
redundant server Node group	Set of two Node devices that connect to servers or storage devices. Link aggregation group (LAG) interfaces can span the Node devices within a redundant server Node group.
rolling upgrade	Method used in the QFabric system to upgrade the software for components in a systematic, low-impact way. A rolling upgrade begins with the Director group, proceeds to the fabric (Interconnect devices), and finishes with the Node groups.
Routing Engine	<p>Juniper Networks-proprietary processing entity that implements QFabric system control plane functions, routing protocols, system management, and user access. Routing Engines can be either physical or virtual entities.</p> <p>The Routing Engine functions in a QFabric system are sometimes handled by Node devices (when connected to endpoints), but mostly implemented by the Director group (to provide support for QFabric system establishment, maintenance, and other tasks).</p>

Table 97: QFabric System Terms (*continued*)

Term	Definition
routing instance	Private collection of routing tables, interfaces, and routing protocol parameters unique to a specific customer. The set of interfaces is contained in the routing tables, and the routing protocol parameters control the information in the routing tables. (See also <i>virtual private network</i> .)
server Node group	Set of one or more Node devices that connect to servers or storage devices.
virtual LAN (VLAN)	Unique Layer 2 broadcast domain for a set of ports selected from the components available in a partition. VLANs allow manual segmentation of larger Layer 2 networks and help to restrict access to network resources. To interconnect VLANs, Layer 3 routing is required.
virtual private network (VPN)	Layer 3 routing domain within a partition. VPNs maintain privacy with a tunneling protocol, encryption, and security procedures. In a QFabric system, a Layer 3 VPN is configured as a <i>routing instance</i> .

**Related
Documentation**

- [QFabric System Overview on page 1165](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)
- [Understanding the QFabric System Software Architecture on page 1192](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

Understanding Interfaces on the QFabric System

This topic describes:

- [Four-Level Interface Naming Convention on page 1173](#)
- [QSFP+ Interfaces on page 1174](#)
- [Link Aggregation on page 1176](#)

Four-Level Interface Naming Convention

When you configure an interface on the QFabric system, the interface name needs to follow a four-level naming convention that enables you to identify an interface as part of either a Node device or a Node group. Include the name of the network or server Node group at the beginning of the interface name.

The four-level interface naming convention is:

device-name:type-fpc/pic/port

where *device-name* is the name of the Node device or Node group. The remainder of the naming convention elements are the same as those in the QFX3500 switch interface naming convention.

An example of a four-level interface name is:

node2:xe-0/0/2

QSFP+ Interfaces

The QFX3500 Node device provides four 40-Gbps QSFP+ (quad small form-factor pluggable plus) interfaces (labeled **Q0** through **Q3**) for uplink connections between your Node device and your Interconnect devices.

The QFX3600 Node device provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight interfaces (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and interfaces Q2 through Q15 for 10-Gigabit Ethernet connections to either endpoint systems or external networks (see [“Configuring the Port Type on QFX3600 Node Devices” on page 1367](#)).

[Table 98 on page 1174](#) shows the port mappings for QFX3600 Node devices.

Table 98: QFX3600 Node Device Port Mappings

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q0	fte-0/1/0	Not supported on this port
Q1	fte-0/1/1	Not supported on this port
Q2	fte-0/1/2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11
Q3	fte-0/1/3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15
Q4	fte-0/1/4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19

Table 98: QFX3600 Node Device Port Mappings (*continued*)

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q5	fte-0/1/5	xe-0/0/20
		xe-0/0/21
		xe-0/0/22
		xe-0/0/23
Q6	fte-0/1/6	xe-0/0/24
		xe-0/0/25
		xe-0/0/26
		xe-0/0/27
Q7	fte-0/1/7	xe-0/0/28
		xe-0/0/29
		xe-0/0/30
		xe-0/0/31
Q8	Not supported on this port	xe-0/0/32
		xe-0/0/33
		xe-0/0/34
		xe-0/0/35
Q9	Not supported on this port	xe-0/0/36
		xe-0/0/37
		xe-0/0/38
		xe-0/0/39
Q10	Not supported on this port	xe-0/0/40
		xe-0/0/41
		xe-0/0/42
		xe-0/0/43
Q11	Not supported on this port	xe-0/0/44
		xe-0/0/45
		xe-0/0/46
		xe-0/0/47

Table 98: QFX3600 Node Device Port Mappings (*continued*)

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q12	Not supported on this port	xe-0/0/48
		xe-0/0/49
		xe-0/0/50
		xe-0/0/51
Q13	Not supported on this port	xe-0/0/52
		xe-0/0/53
		xe-0/0/54
		xe-0/0/55
Q14	Not supported on this port	xe-0/0/56
		xe-0/0/57
		xe-0/0/58
		xe-0/0/59
Q15	Not supported on this port	xe-0/0/60
		xe-0/0/61
		xe-0/0/62
		xe-0/0/63

Link Aggregation

Link aggregation enables you to create link aggregation groups across Node devices within a network Node group or redundant server Node group. You can include up to 32 Ethernet interfaces in a LAG. You can have up to 48 LAGs within a redundant server Node group, and 128 LAGs in a network Node group. To configure a LAG, include the **aggregated-devices** statement at the **[edit chassis node-group node-group-name]** hierarchy level and the **device-count** statement at the **[edit chassis node-group node-group-name aggregated-devices ethernet]** hierarchy level. Additionally, include any aggregated Ethernet options (**minimum-links** and **link-speed**) at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level and the **802.3ad** statement at the **[edit interfaces interface-name ether-options]** hierarchy level. To configure the Link Aggregation Control Protocol (LACP), include the **lACP** statement at the **[edit interfaces aggregated-ether-options]** hierarchy level.

Related Documentation

- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)

Hardware Architecture Overview

- [Understanding the QFabric System Hardware Architecture on page 1177](#)
- [Understanding the Director Group on page 1180](#)
- [Understanding Routing Engines in the QFabric System on page 1181](#)
- [Understanding Interconnect Devices on page 1183](#)
- [Understanding Node Devices on page 1186](#)
- [Understanding Node Groups on page 1190](#)
- [Understanding Port Oversubscription on Node Devices on page 1191](#)

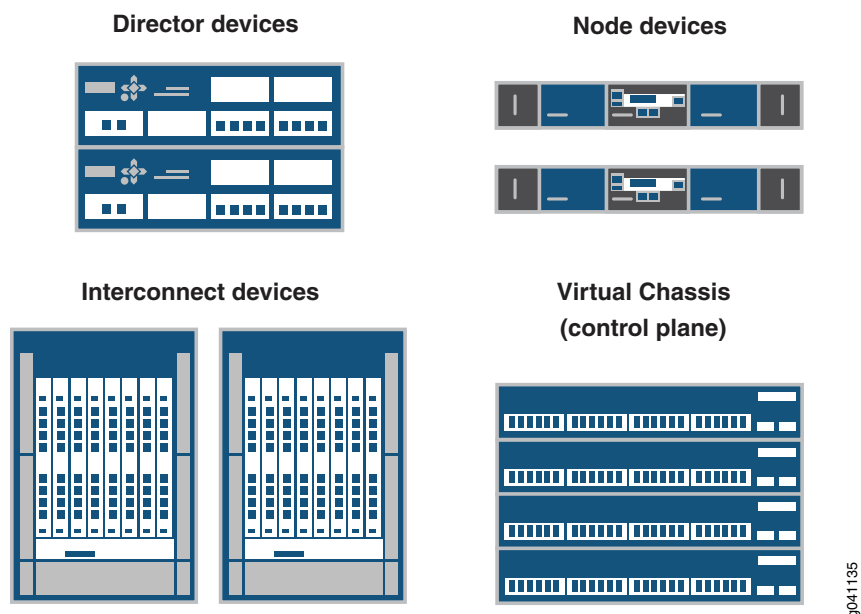
Understanding the QFabric System Hardware Architecture

- [QFabric System Hardware Architecture Overview on page 1177](#)
- [QFX3000-G QFabric System Features on page 1180](#)
- [QFX3000-M QFabric System Features on page 1180](#)

QFabric System Hardware Architecture Overview

The QFabric system is a single-layer networking tier that connects servers and storage devices to one another across a high-speed, unified core fabric. You can view the QFabric system as a single, extremely large, nonblocking, high-performance Layer 2 and Layer 3 switching system. The reason you can consider the QFabric system as a single system is that the Director software running on the Director group allows the main QFabric system administrator to access and configure every device and port in the QFabric system from a single location. Although you configure the system as a single entity, the fabric contains four major hardware components. The hardware components can be chassis-based, group-based, or a hybrid of the two. As a result, it is important to understand the four types of generic QFabric system components and their functions, regardless of which hardware environment you decide to implement. A representation of these components is shown in [Figure 15 on page 1178](#).

Figure 15: QFabric System Hardware Architecture



The four major QFabric system components include the following:

- **Director group**—The *Director group* is a management platform that establishes, monitors, and maintains all components in the QFabric system. It is a set of Director devices that run the Junos operating system (Junos OS) on top of a CentOS foundation. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration and startup, and Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) services. The Director group also runs the software for management applications, hosts and load-balances internal processes for the QFabric system, and starts additional QFabric system processes as requested.
- **Node devices**—A *Node device* is a hardware system located on the ingress of the QFabric system that connects to endpoints (such as servers or storage devices) or external networks, and is connected to the heart of the QFabric system through an Interconnect device. A Node device can be used in a manner similar to how a top-of-rack switch is implemented. By default, Node devices connect to servers or storage devices. However, when you group Node devices together to connect to a network that is external to the QFabric system, the formation is known as a *network Node group*.
- **Interconnect devices**—An *Interconnect device* acts as the primary fabric for data plane traffic traversing the QFabric system between Node devices. To reduce latency to a minimum, the Interconnect device implements multistage Clos switching to provide nonblocking interconnections between any of the Node devices in the system.
- **Control plane network**—The *control plane network* is an out-of-band Gigabit Ethernet management network that connects all QFabric system components. For example, you can use a group of EX4200 Ethernet switches configured as a Virtual Chassis to enable the control plane network. The control plane network connects the Director

group to the management ports of the Node and Interconnect devices. By keeping the control plane network separate from the data plane, the QFabric system can scale to support thousands of servers and storage devices.

The four major QFabric system components can be assembled from a variety of hardware options. Currently supported hardware configurations are shown in [Table 99 on page 1179](#).

Table 99: Supported QFabric System Hardware Configurations

QFabric System Configuration	Director Group	Node Device	Interconnect Device	Control Plane Device
QFX3000-G QFabric system	QFX3100 Director group	QFX3500 and QFX3600 Node device NOTE: There can be a maximum of 128 Node devices in the QFX3000-G QFabric system.	QFX3008-I Interconnect device NOTE: There can be a maximum of four Interconnect devices in the QFX3000-G QFabric system.	Two Virtual Chassis composed of four EX4200 switches each
QFX3000-M QFabric system	QFX3100 Director group NOTE: For a copper-based QFX3000-M QFabric system control plane network, use QFX3100 Director devices with RJ-45 network modules installed. For a fiber-based control plane network, use QFX3100 Director devices with SFP network modules installed.	QFX3500 and QFX3600 Node device NOTE: <ul style="list-style-type: none"> There can be a maximum of 16 Node devices in the QFX3000-M QFabric system. For a copper-based QFX3000-M QFabric system control plane network, use QFX3500 Node devices with a 1000BASE-T management board installed. For a fiber-based control plane network, use QFX3500 Node devices with an SFP management board installed. 	QFX3600-I Interconnect device NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.	Two EX4200 Ethernet switches NOTE: For a copper-based QFX3000-M QFabric system control plane network, use EX4200-24T switches with an SFP+ uplink module installed. For a fiber-based control plane network, use EX4200-24F switches with an SFP+ uplink module installed.

To complete the system, external Routing Engines (such as the fabric manager Routing Engines, network Node group Routing Engines, and fabric control Routing Engines) run on the Director group and implement QFabric system control plane functions. The control plane network provides the control plane connections between the Node devices, the Interconnect devices, and the Routing Engines running on the Director group.

QFX3000-G QFabric System Features

A QFX3000-G QFabric system provides the following key features:

- Support for up to 128 Node devices and 4 Interconnect devices, which provides a maximum of 6144 10-Gigabit Ethernet ports.
- Low port-to-port latencies that scale as the system size grows from 48 to 6144 10-Gigabit Ethernet ports.
- Support for up to 384,000 total ingress queues at each Node device to the QFabric system Interconnect backplane.
- Support for Converged Enhanced Ethernet (CEE) traffic.

QFX3000-M QFabric System Features

A QFX3000-M QFabric system provides the following key features:

- Support for up to 16 Node devices and 4 Interconnect devices, which provides a maximum of 768 10-Gigabit Ethernet ports.
- Low port-to-port latencies that scale as the system size grows from 48 to 768 10-Gigabit Ethernet ports.

Related Documentation

- [Understanding QFabric System Terminology on page 1169](#)
- [Understanding the QFabric System Software Architecture on page 1192](#)
- [Understanding the Director Group on page 1180](#)
- [Understanding Routing Engines in the QFabric System on page 1181](#)
- [Understanding Interconnect Devices on page 1183](#)
- [Understanding Node Devices on page 1186](#)
- [Understanding Node Groups on page 1190](#)
- [Understanding Partitions on page 1194](#)

Understanding the Director Group

Because the Director group provides management services for the QFabric system, it is important to understand the components of the cluster and how the Director group supports the needs of the greater fabric.

- [Director Group Components on page 1180](#)
- [Director Group Services on page 1181](#)

Director Group Components

When you build a Director group, consider the following elements and concepts.

- **Director device**—A single management device for the QFabric system. Director devices with a hard drive provide full processing services and are used to build the Director group.
- **Director group**—A set of Director devices. The Director group is essential to the QFabric system, which cannot operate properly without it. The Director group shares and load-balances processing tasks for the QFabric system, performs topology discovery, assigns identifiers to QFabric system components, and manages interfabric communication. The primary devices in a Director group are Director devices that contain hard drives. The Director devices run dual processes in active or standby mode for maximum redundancy.

When you add additional Director devices to the group, the Director group coordinates their activities and distributes processing loads across all available Director devices. The additional Director devices provide the Director group with additional memory and processing power. Supplementing the Director group with extra Director devices allows the group to scale efficiently and serve the needs of the entire QFabric system as it grows.

Director Group Services

The Director group is a management platform that establishes, monitors, and maintains all components in the QFabric system. It is a set of Director devices that run the Junos operating system (Junos OS) on top of a CentOS foundation. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration and startup, and Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) services. The Director group also runs the software for management applications, hosts and load-balances internal processes for the QFabric system, maintains configuration and topology databases, and starts additional QFabric system processes as requested.

Another critical role provided by the Director group is the hosting of the virtual Junos Routing Engines. These Routing Engines provide services for the QFabric system to keep it operating smoothly.

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Understanding Routing Engines in the QFabric System on page 1181](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding Routing Engines in the QFabric System

Routing Engines perform many important processing tasks in the QFabric system. Knowing where the Routing Engines are located and what services they provide enables you to troubleshoot the QFabric system and ensure that it is running the way it should.

This topic covers:

- [Hardware-Based Routing Engines on page 1182](#)
- [Software-Based External Routing Engines on page 1182](#)

Hardware-Based Routing Engines


A traditional Juniper Networks Routing Engine is a hardware field-replaceable unit that runs routing protocols, builds the routing and switching tables, sends routing information to the Packet Forwarding Engine, and handles several software processes for the device (such as interface control, chassis component monitoring, system management, and user access). Node devices that are part of server Node groups in the QFabric system that connect to servers or storage devices implement Routing Engine functions locally using this traditional hardware method.

Software-Based External Routing Engines

The QFabric system also uses external Routing Engines that run in software on the Director group. In contrast with traditional Routing Engines, the functions and processes provided by software-based Routing Engines are segmented, specialized, and distributed across multiple Routing Engine instances running on the Director group. Such separation provides redundancy for these functions and enables the QFabric system to scale.

[Figure 16 on page 1182](#) shows the external Routing Engine types.

Figure 16: External Routing Engine Types



These special-purpose external Routing Engine instances running on the Director group provide the following major services for the QFabric system:

- **Fabric manager Routing Engine**—Provides services to all devices in the QFabric system, such as system initialization, topology discovery, internal IP address and ID assignment, and interdevice communication. The fabric manager Routing Engine authenticates Interconnect and Node devices, and maintains a database for system components. A single fabric manager Routing Engine instance is generated to manage the entire QFabric system.
- **Fabric control Routing Engine**—Runs the fabric control protocol to share route information between available devices in a partition. A pair of redundant route distribution Routing Engine instances is generated for every partition in the QFabric system, and both instances are active.
- **Diagnostic Routing Engine**—Gathers operational information that allows QFabric system administrators to monitor the health of the QFabric system. A single Routing Engine instance is generated for the entire QFabric system.
- **Network Node group Routing Engine**—Provides Routing Engine functionality for groups of Node devices bundled together as a single Layer 3 routing device, which is used to

connect to external networks. A pair of redundant Routing Engine instances is generated for every network Node group in the QFabric system.

- Related Documentation**
- [Understanding the Director Group on page 1180](#)
 - [Understanding the QFabric System Control Plane on page 1196](#)
 - [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding Interconnect Devices

Interconnect devices in a QFabric system provide a way for the Node devices to connect with one another over a high-speed backplane. By understanding the role of Interconnect devices, you can harness the benefits of low latency, superb scalability, and minimum packet processing offered by a single-tier data center architecture.

This topic covers:

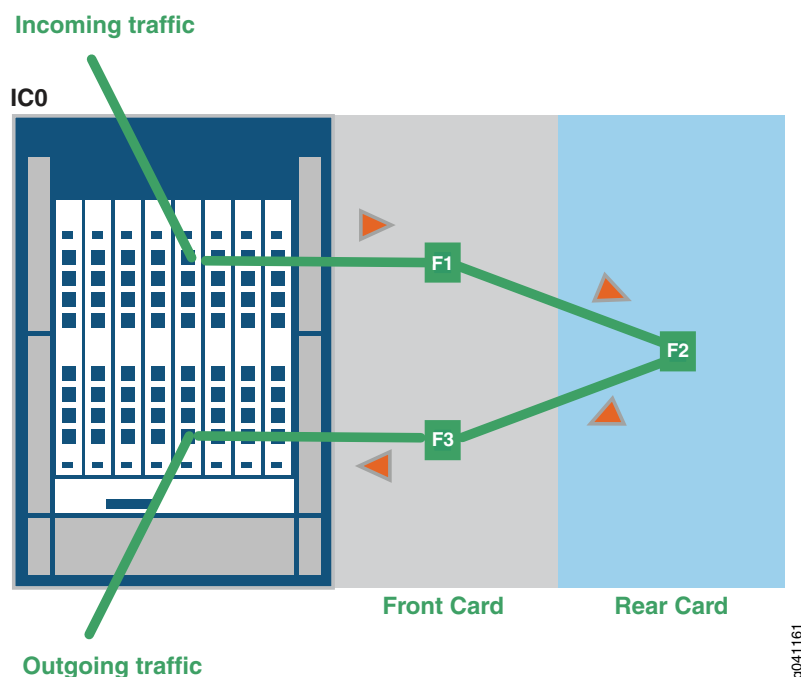
- [Interconnect Device Introduction on page 1183](#)
- [QFX3008-I Interconnect Devices on page 1184](#)
- [QFX3600-I Interconnect Devices on page 1185](#)

Interconnect Device Introduction

Interconnect devices act as the primary fabric for data plane traffic traversing the QFabric system between Node devices. The main task for the Interconnect devices is to transfer traffic between the Node devices as quickly as possible across a high-speed, available path backplane. To reduce latency to a minimum, larger Interconnect devices (such as the QFX3008-I Interconnect device) implement multistage Clos switching to provide nonblocking connections between any of the Node devices in the system.

[Figure 17 on page 1184](#) shows an example of how Clos switching works in the QFX3008-I Interconnect device.

Figure 17: Clos Switching for QFX3008-I Interconnect Devices

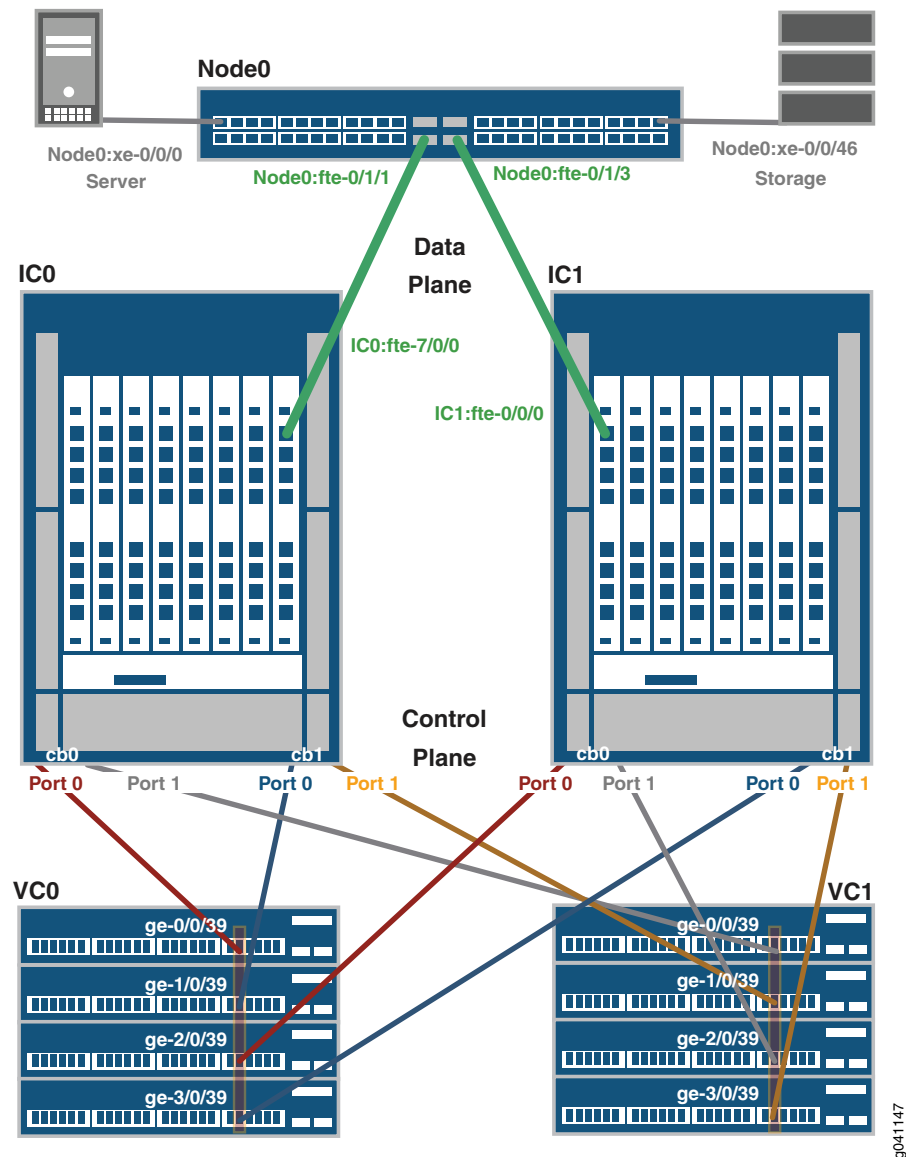


Traffic enters a QSFP+ port from a Node device, and an ingress chipset provides stage F1 processing. For the F2 stage, the frame is sent to a rear card and processed by a midplane chipset. Lastly, an egress chipset on the front card QSFP+ port handles processing tasks for the F3 stage. At each of the three Clos stages, a switching table chooses the best path and determines where to send the frame to reach the next stage. The F1 and F3 stages can be handled by the same front card or different front cards, depending on the best path selected by the fabric. After the frame traverses the Interconnect device backplane, the Interconnect device sends the frame to the egress Node device.

QFX3008-I Interconnect Devices

The QFX3008-I Interconnect device contains eight slots in the front of the chassis. In each slot, you can install a front card containing 16 40-Gbps quad small form-factor pluggable plus (QSFP+) ports. A fully configured system offers a total capacity of 128 QSFP+ connections. These front card ports attach to the high-speed backplane to reach the eight slots in the rear of the chassis, which provide the heavy-duty interconnections for the entire QFX3000-G QFabric system. In addition, four interfaces (two per Control Board) provide Gigabit Ethernet access to the control plane management network. [Figure 18 on page 1185](#) shows an example of the data plane and control plane connections for QFX3008-I Interconnect devices.

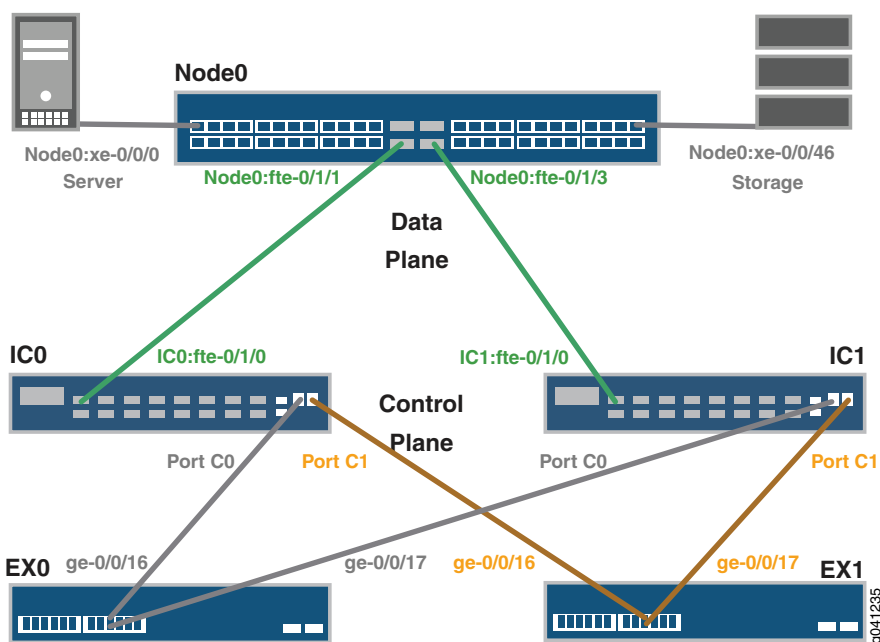
Figure 18: QFX3008-I Data Plane and Control Plane Connections



QFX3600-I Interconnect Devices

The QFX3600-I Interconnect device has 16 40-Gbps quad small form-factor pluggable plus (QSFP+) ports that provide interconnections for the entire QFX3000-M QFabric system. In addition, two management ports provide Gigabit Ethernet access to the control plane management network. [Figure 19 on page 1186](#) shows an example of the data plane and control plane connections for a QFX3600-I Interconnect device.

Figure 19: QFX3600-I Data Plane and Control Plane Connections



Related Documentation

- [Understanding Node Devices on page 1186](#)
- [Understanding the QFabric System Data Plane on page 1199](#)
- [Understanding the QFabric System Control Plane on page 1196](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding Node Devices

Node devices in a QFabric system provide a way for servers, storage devices, and external networks to connect to the QFabric system. By understanding the role of Node devices, you can design your QFabric system topology to take advantage of the unique benefits offered by a single-tier data center architecture.

This topic covers:

- [Node Device Introduction on page 1186](#)
- [QFX3500 Node Devices on page 1187](#)
- [QFX3600 Node Devices on page 1188](#)

Node Device Introduction

A *Node device* in the QFabric system connects either endpoint systems (such as application servers and storage devices) or external networks to Interconnect devices. It can be used similarly to the way a top-of-rack switch is implemented in a data center. Node devices provide an access point to the QFabric system, allowing data to flow into and out of the QFabric system. Because all Node devices in the QFabric system connect through a backplane of Interconnect devices, in essence all Node devices are connected

to one another. This directly connected design model eliminates multiple tiers of aggregation and core devices and provides minimum latency, maximum scalability, and rapid transport of server-to-server traffic and QFabric system-to-external network traffic.

Sets of Node devices can be bundled together into *Node groups*, in which each group operates as a single virtual entity. Node groups that connect to servers and storage devices are known as *server Node groups*, and Node groups that connect to external networks are known as *network Node groups*.

QFX3500 Node Devices

The QFX3500 Node device works as part of a QFabric system. A QFX3500 chassis provides up to 48 10-Gigabit Ethernet interfaces to connect to endpoints or external networks. You can configure 12 of these 48 interfaces to support 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel. You can also configure the remaining 36 interfaces with Gigabit Ethernet.

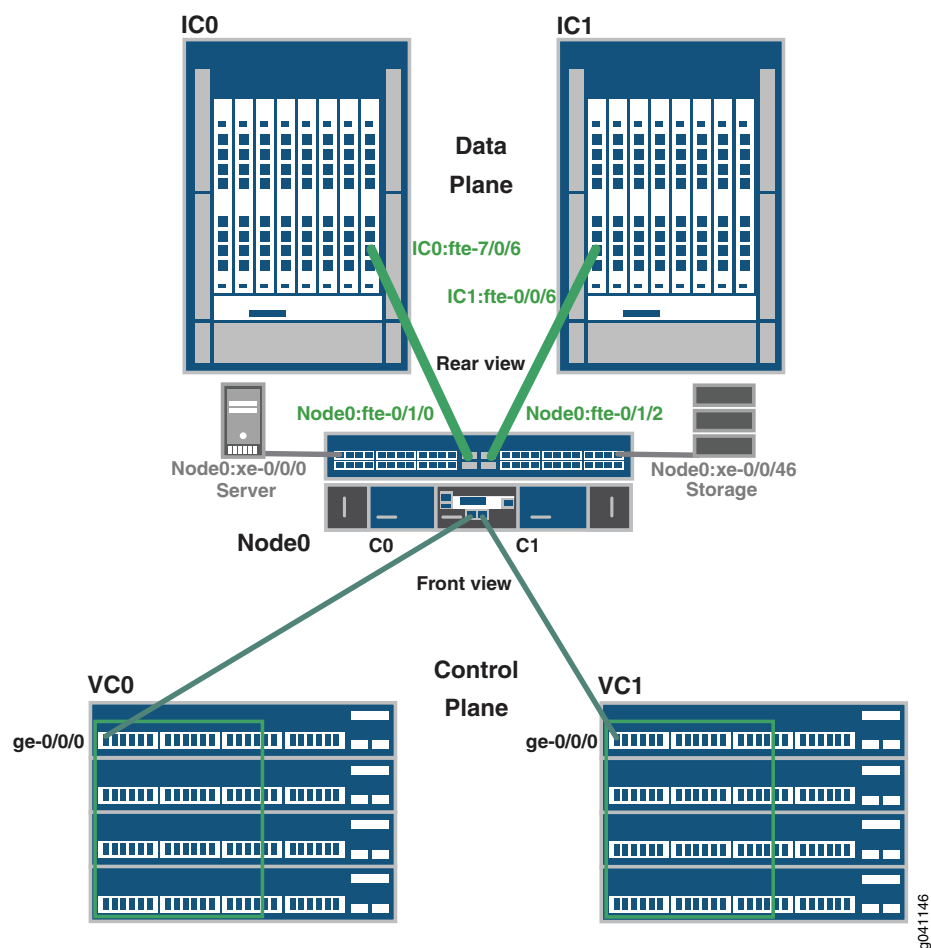


NOTE: You can configure interface ports 0 through 47 as 10-Gigabit Ethernet ports, 0 through 5 and 42 through 47 as Fibre Channel over Ethernet ports, and 6 through 41 as Gigabit Ethernet ports. However, you cannot configure any Fibre Channel over Ethernet ports as Gigabit Ethernet ports or vice versa.

In addition to these server and network interfaces, there are four uplink interfaces to connect the QFX3500 Node device to Interconnect devices in a QFabric system. These uplinks use 40-Gbps quad small form-factor pluggable plus (QSFP+) interfaces.

The control plane requires two management ports on the QFX3500 chassis to connect the Node device to the control plane network. [Figure 20 on page 1188](#) shows an example of the data plane and control plane connections for a QFX3500 Node device.

Figure 20: QFX3500 Data Plane and Control Plane Connections

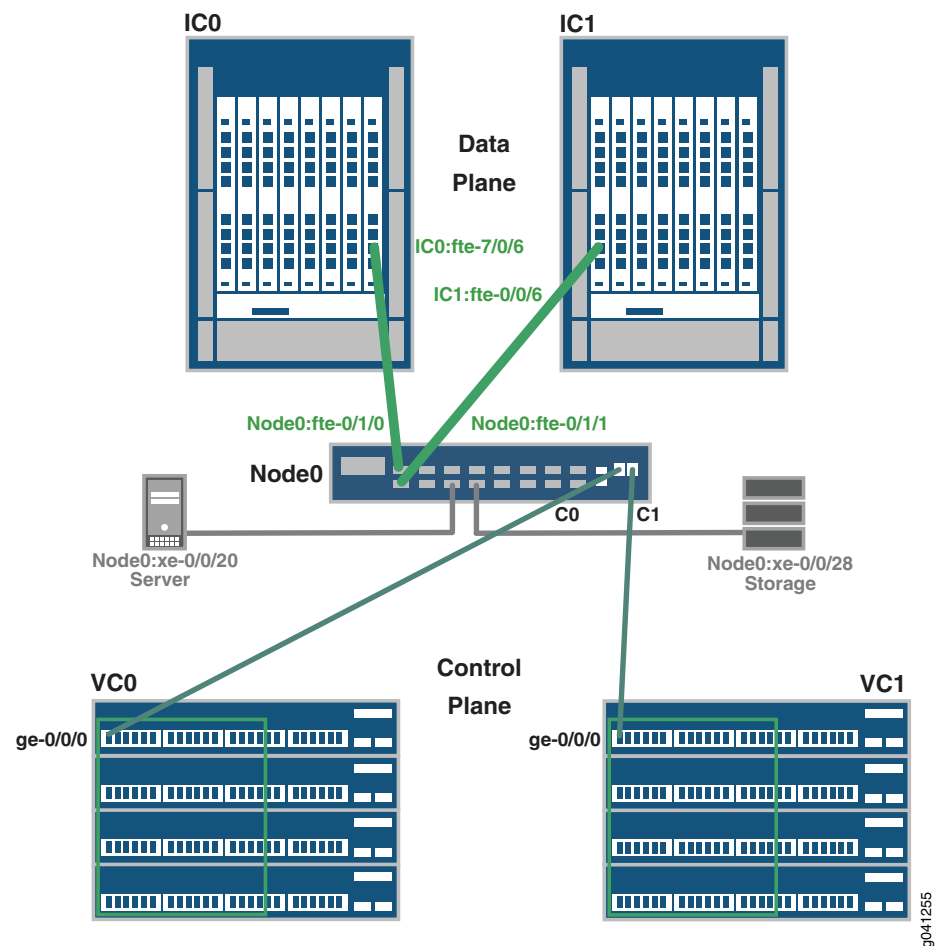


QFX3600 Node Devices

The QFX3600 Node device works as part of a QFabric system. A QFX3600 chassis provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your QFX3600 Node device and your Interconnect device, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure the first eight interfaces (**Q0** through **Q7**) for uplink connections between your Node device and your Interconnect devices, and interfaces **Q2** through **Q15** for 10-Gigabit Ethernet connections to either endpoint systems or external networks.

The control plane requires two management ports on the QFX3600 chassis to connect the Node device to the control plane network. [Figure 21 on page 1189](#) shows an example of the data plane and control plane connections for a QFX3600 Node device.

Figure 21: QFX3600 Data Plane and Control Plane Connections



Related Documentation

- [Converting the Device Mode for a QFabric System Component on page 1258](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)
- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)
- [Understanding Node Groups on page 1190](#)
- [Understanding Interconnect Devices on page 1183](#)
- [Understanding the QFabric System Data Plane on page 1199](#)
- [Understanding the QFabric System Control Plane on page 1196](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding Node Groups

Node groups help you combine multiple Node devices into a single virtual entity within the QFabric system to enable redundancy and scalability at the edge of the data center.

This topic covers:

- [Network Node Groups on page 1190](#)
- [Server Node Groups on page 1190](#)

Network Node Groups

A set of one or more Node devices that connect to an external network is called a *network Node group*. The network Node group also relies on two external Routing Engines running on the Director group. These redundant *network Node group Routing Engines* run the routing protocols required to support the connections from the network Node group to external networks.

When configured, the Node devices within a network Node group and the network Node group Routing Engines work together in tandem as a single entity. By default, network Node group Routing Engines are part of the **NW-NG-0** network Node group but no Node devices are included in the group. As a result, you must configure Node devices to be part of a network Node group.

In a QFabric system deployment that requires connectivity to external networks, you can modify the automatically generated network Node group by including its preset name **NW-NG-0** in the Node group configuration. Within a network Node group, you can include a minimum of one Node device up to a maximum of eight Node devices. By adding more Node devices to the group, you provide enhanced scalability and redundancy for your network Node group.



NOTE: The QFabric system creates a single **NW-NG-0** network Node group for the default partition. You cannot configure a second network Node group inside the default partition. The remaining Node devices within the default partition are reserved to connect to servers, storage, or other endpoints internal to the QFabric system. These Node devices either can be retained in the automatically generated server Node groups or can be configured as part of a redundant server Node group.

Server Node Groups

A *server Node group* is a set of one or more Node devices that connect to servers or storage devices. Unlike Node devices that are part of a network Node group and rely on an external Routing Engine, a Node device within a server Node group connects directly to endpoints and implements the Routing Engine functions locally, using the local CPU built into the Node device itself.

By default, each Node device is placed in its own self-named autogenerated server Node group to connect to servers and storage. You can override the default assignment by

manually configuring a redundant server Node group that contains a maximum of two Node devices. You can use a redundant server Node group to provide multihoming services to servers and storage, as well as configure aggregated LAG connections that span the two Node devices.



NOTE: The Node devices in a redundant server Node group must be of the same type, either a QFX3500 Node device or a QFX3600 Node device. You cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

Related Documentation

- [Configuring Node Groups for the QFabric System on page 1363](#)
- [Understanding Node Devices on page 1186](#)
- [Understanding Routing Engines in the QFabric System on page 1181](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding Port Oversubscription on Node Devices

Each Node device in a QFabric system can have a different port oversubscription configuration. For example, you can have a Node device with 3:1 port oversubscription, another with 6:1 oversubscription, and yet another with 1:1 oversubscription.

The port oversubscription ratio on a Node device is based on the number of uplink connections from the Node device to Interconnect devices. For example, you can configure 1:1 port oversubscription on your QFX3600 Node device by connecting the eight uplink ports (labeled Q0 through Q7) on the Node device to Interconnect devices.

[Table 100 on page 1191](#) shows the oversubscription ratio for ports on Node devices based on the number of Interconnect devices and the number of connections from each Node device to each Interconnect device.

Table 100: Oversubscription Ratio on Node Devices

Number of Interconnect Devices	Number of Connections from Each Node Device to Each Interconnect Device	Oversubscription Ratio on Node Device
2	1	6:1
2	2	3:1
2	4	1:1 (Supported on QFX3600 Node devices only)
4	1	3:1
4	2	1:1 (Supported on QFX3600 Node devices only)

- Related Documentation**
- [*Connecting a QFX3500 Node Device to a QFX3008-I Interconnect Device*](#)
 - [*Connecting a QFX3600 Node Device to a QFX3008-I Interconnect Device*](#)
 - [*Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device*](#)
 - [*Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device*](#)

Software Architecture Overview

- [Understanding the QFabric System Software Architecture on page 1192](#)
- [Understanding the Director Software on page 1193](#)
- [Understanding Partitions on page 1194](#)
- [Understanding the QFabric System Control Plane on page 1196](#)
- [Understanding the QFabric System Data Plane on page 1199](#)

Understanding the QFabric System Software Architecture

The software architecture for the QFabric system environment has been designed to provide a high-speed, low-latency, nonblocking fabric for data center traffic. This topic explores how the software architecture for a QFabric system supports these goals.

Key components of the QFabric system software architecture include:

- A single administrative view of all QFabric system components provides unified management, configuration, monitoring, and troubleshooting of the QFabric system. This view is provided by the QFX Series Director software running on the Director group. A primary administrator can access the unified view through the default partition.
- A fabric control protocol enables rapid transport of data traffic between QFabric system components. This unique feature of the software architecture distributes route information for each device within the QFabric system, and removes the need to run spanning-tree protocols inside the QFabric system network.
- A fabric management protocol provides rapid transport of control traffic between QFabric system components. This protocol helps identify and initialize QFabric system resources, supports device redundancy, and supports management communication throughout the QFabric system.
- A control plane network that is separate from the data plane network provides high availability for the QFabric system.

The software also provides access to relevant features in the Junos operating system (Junos OS) that support QFabric system functionality. Support is available for most switching features available on EX Series Ethernet switches and many routing features available on M Series, MX Series, and T Series routing platforms.

- Related Documentation**
- [Understanding QFabric System Terminology on page 1169](#)
 - [Understanding the QFabric System Hardware Architecture on page 1177](#)
 - [Understanding the Director Software on page 1193](#)

- [Understanding Partitions on page 1194](#)
- [Understanding the QFabric System Control Plane on page 1196](#)
- [Understanding the QFabric System Data Plane on page 1199](#)

Understanding the Director Software

The Director software provides a single view into the QFabric system so that it can be managed as a single entity. This topic explains how the Director software interacts with the components of the QFabric system to maintain operations from a central location.

Because the QFabric system consists of multiple Director, Node, and Interconnect devices, the architects of the QFabric system determined that it would be useful to manage the entire system as a single logical entity. As a result, the Director software handles administration tasks for the entire QFabric system, such as fabric management and configuration. The Director software runs on the *Director group*, provides a single consolidated view of the QFabric system, and enables the main QFabric system administrator to configure, manage, monitor, and troubleshoot QFabric system components from a centralized location. In the Junos operating system (Junos OS) command-line interface (CLI), you can access the Director software by logging in to the default partition.

The Director software handles the following major tasks for the QFabric system:

- Provides command-line interface (CLI) access to all QFabric system components that you have permission to manage or view.
- Evaluates configuration statements and operational mode commands for their scope and sends requests to the applicable Director, Node, and Interconnect devices. (This operation is sometimes referred to as *scattering*.)
- Consolidates responses from Director, Node, and Interconnect devices, and displays output from the devices in a unified, centralized manner. (This operation is sometimes referred to as *gathering*.)
- Coordinates configuration and operational efforts with a database housed in the Director group to store and retrieve configurations, software images, event logs, and system log messages.
- Facilitates control plane communication between the Node devices, the Routing Engine services running on the Director group, and the Interconnect devices.
- Runs parallel processes on the Director group devices to provide high availability for the QFabric system.
- Coordinates interactions with QFabric system components to provide load balancing of processing tasks across the Director group devices.
- Manages user access and privileges.
- Enables you to configure, manage, monitor, and troubleshoot QFabric system components that are assigned to you.
- Gathers QFabric system inventory and topology details.

- Offers a way to manage Director group devices, including the ability to add and delete Director devices in the group, set and switch mastership in the Director group, and monitor Director group status.
- Provides a centralized way to coordinate software upgrades for QFabric system components.

The Director software provides a backbone of functionality that supports the entire QFabric system. It is an essential component of the QFabric system that enables you to implement the system in a logical and efficient way.

**Related
Documentation**

- [Gaining Access to the QFabric System Through the Default Partition on page 1344](#)
- [Understanding the Director Group on page 1180](#)
- [Understanding the QFabric System Software Architecture on page 1192](#)

Understanding Partitions


Partitions provide a way to allocate specified virtual and physical resources within your QFabric system. This topic covers:

- [QFabric System Default Partition on page 1194](#)

QFabric System Default Partition

By default, all equipment and virtual resources in the QFabric system belong to the *default partition*. As a result, the QFabric system in its initial state has a single broadcast domain that is administered by a single main administrator. [Figure 22 on page 1195](#) shows a topology with the default settings—a single collection that contains all the devices in the QFabric system.

Figure 22: QFabric System Topology - Default Partition



g041145



NOTE: The initial release of the QFabric system supports a single default partition. All equipment and resources belong to the default partition.

A partition provides the following functions:

- Fault isolation and separation from other partitions at the control plane level.
- A separate configuration domain for the Node devices within the partition.
- A Layer 2 domain in which MAC learning takes place, and members of the same VLAN can communicate with each other. To provide network connectivity between partitions, you need to enable Layer 3 routing by way of a routed VLAN interface (RVI).

Related Documentation

- [Gaining Access to the QFabric System Through the Default Partition on page 1344](#)
- [Understanding the QFabric System Software Architecture on page 1192](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding the QFabric System Control Plane


The control plane in the QFabric system transports management traffic between QFabric system components to facilitate system operations, configuration, and maintenance. This topic covers:

- [Control Plane Elements on page 1197](#)
- [Control Plane Services on page 1199](#)

Control Plane Elements

Control traffic within a QFabric system is carried across a redundant, scalable, out-of-band, Ethernet switching network called the *control plane* network. To maintain high availability, the QFabric system control plane is separated from the QFabric system data plane. [Figure 23 on page 1197](#) shows a diagram of the QFabric system devices that compose the control plane network.

Figure 23: QFabric System Control Plane Network



The control plane consists of the following elements:

- **Control plane switches**—Provide connectivity to the management interfaces of all QFabric system components in the control plane network, including the Node devices, the Interconnect devices, and the Director group. When you interconnect all QFabric system devices to the control plane switches, the Director group can manage the entire system. Depending on the size and scale of your QFabric system, the control plane switches might be standalone switches or might be groups of switches bundled into a Virtual Chassis (See the Example topics in the Related Documentation section of this topic to learn more about the control plane switch configuration required for your QFabric system.)

For example, the control plane switch for the QFX3000-G QFabric system requires two Virtual Chassis containing four EX4200 switch members each. The two Virtual Chassis connect to each other across a 10-Gigabit Ethernet LAG link to provide maximum resiliency for the QFabric system control plane.

- **Connections between the management interfaces of the Node devices and the control plane switches**—Enable control plane connectivity from the Node devices to the rest of the QFabric system. You must connect two management interfaces from

each Node device to the control plane switches. Connect each interface to a different control plane switch to provide system resiliency.

For the most current guidance on the QFabric control plane configuration and cabling recommendations, see:

- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- **Connections between the management interfaces of the Interconnect devices and the control plane switches**—Enable control plane connectivity from the Interconnect devices to the rest of the QFabric system. You must connect the interfaces in each Interconnect device to the control plane switches. Connect each interface to a different control plane switch to provide system resiliency.

For example, on QFX3008-I Interconnect devices, there are two Control Boards and two interfaces per Control Board, for a total of four connections per Interconnect device. To provide system resiliency, connect one interface from each Control Board to the first Virtual Chassis, and connect the second interface from each Control Board to the second Virtual Chassis.

For the most current guidance on the QFabric control plane configuration and cabling recommendations, see:

- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- **Connections between the network module interfaces of the Director group and the control plane switches**—Enable control plane connectivity from the Director group to the rest of the QFabric system. You must connect some interfaces from the first network module in a Director device to one control plane switch, and connect some interfaces from the second network module in a Director device to the second control plane switch. Also, you must connect the ports from the first network module to the primary control plane switch for each Director device (which may vary depending on the configuration of your Director group).

For the most current guidance on the QFabric control plane configuration and cabling recommendations, see:

- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- **Routing Engines**—Although they are automatically provisioned, specialized Routing Engines implement services such as default QFabric system infrastructure, device

management, route sharing, and diagnostics to support the QFabric system. Routing Engines for control plane functions are virtual entities that run on the Director group.

- **Fabric management protocol**—A link-state protocol runs on the control plane network to identify and initialize QFabric system resources, support device redundancy, and support management communication throughout the QFabric system. The protocol is enabled by default.

Control Plane Services

The QFabric system control plane provides the infrastructure to support the following services for the QFabric system:

- System initialization
- Topology discovery
- Internal IP address and unique ID assignment
- Route information sharing
- Configuration delivery to Node devices
- Interdevice communication between Node devices, Interconnect devices, and the Director group

Many of these services are provided by the external Routing Engines that run in software on the Director group.

Related Documentation

- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- [Understanding the QFabric System Data Plane on page 1199](#)
- [Understanding Routing Engines in the QFabric System on page 1181](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Understanding the QFabric System Data Plane

The data plane in the QFabric system transfers application traffic between QFabric system components rapidly and efficiently. This topic covers:

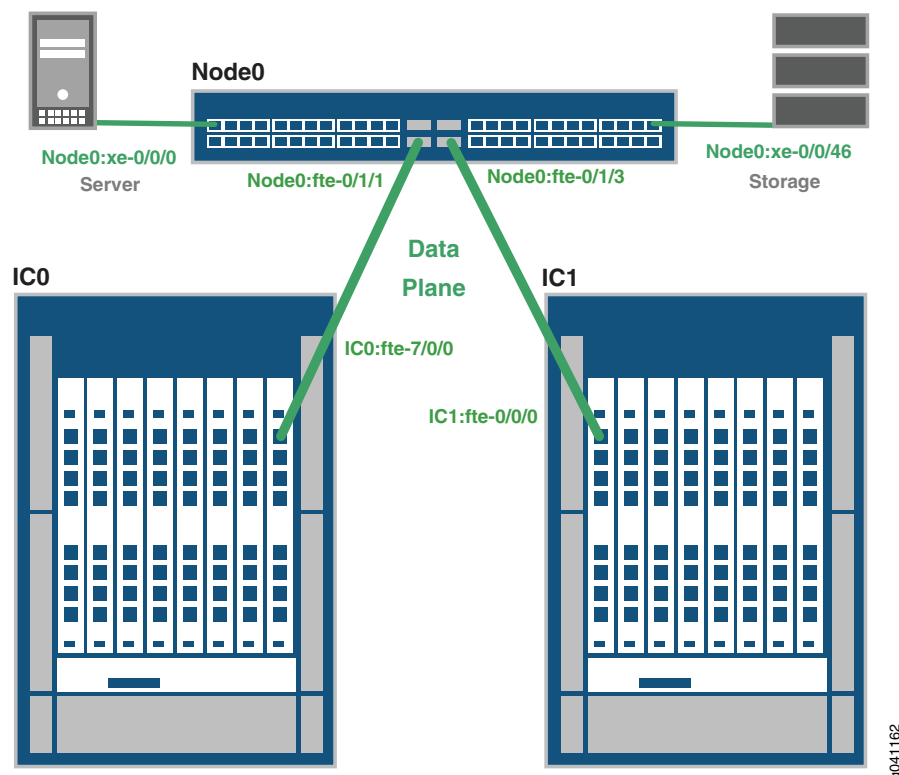
- [Data Plane Components on page 1199](#)
- [QFabric System Fabric on page 1200](#)

Data Plane Components

Data traffic within a QFabric system is carried across a redundant, high-performance, and scalable *data plane*. To maintain high availability, the QFabric system data plane is separated physically from the QFabric system control plane and uses a different network.

Figure 24 on page 1200 shows an example diagram of the QFabric system data plane network.

Figure 24: QFabric System Data Plane Network



The QFabric system data plane includes the following high-speed data connections and elements:


- 10-Gigabit Ethernet or 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel connections between QFabric system endpoints (such as servers or storage devices) and the Node devices.
- 40-Gbps quad, small form-factor pluggable plus (QSFP+) connections between the Node devices and the Interconnect devices.
- 10-Gigabit Ethernet connections between external networks and the Node devices contained in the network Node group.
- A fabric control protocol, used to distribute route information to all devices connected to the QFabric system data plane.

QFabric System Fabric

Unlike traditional data centers that employ a multi-tiered hierarchy of switches, a QFabric system contains a single tier of Node devices connected to one another across a backplane of Interconnect devices. The QFabric system fabric is a distributed, multistage network that consists of a fabric queuing and scheduling system implemented in the Node devices, and a distributed cross-connect system implemented in the Interconnect devices. The

cross-connect system for the QFX3008-I Interconnect device is shown as an example in [Figure 25 on page 1201](#).

Figure 25: QFX3008-I Interconnect Device Cross-Connect System



The design of the cross-connect system provides multistage Clos switching, which results in nonblocking paths for data traffic and any-to-any connectivity for the Node devices. Because all Node devices are connected through the Interconnect device, the QFabric system offers very low port-to-port latencies. In addition, dynamic load balancing and low-latency packet flows provide for scaling the port count and bandwidth capacity of a QFabric system.

Related Documentation

- [Understanding the QFabric System Control Plane on page 1196](#)
- [Understanding the QFabric System Hardware Architecture on page 1177](#)

Software Features

- [QFX Series Software Features Overview on page 1202](#)
- [Understanding Software Upgrade on the QFabric System on page 1217](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 1218](#)
- [Understanding Statements and Commands on the QFabric System on page 1223](#)
- [Understanding NTP on the QFabric System on page 1224](#)
- [Understanding Network Management Implementation on the QFabric System on page 1224](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)

- [Understanding the Implementation of System Log Messages on the QFabric System on page 1227](#)
- [Understanding User and Access Management Features on the QFabric System on page 1229](#)
- [Understanding QFabric System Login Classes on page 1230](#)
- [Understanding Interfaces on the QFabric System on page 1231](#)
- [Understanding Layer 3 Features on the QFabric System on page 1234](#)
- [Understanding Security Features on the QFabric System on page 1235](#)
- [Understanding Port Mirroring on the QFabric System on page 1236](#)
- [Understanding Fibre Channel Fabrics on the QFabric System on page 1236](#)
- [Understanding CoS Fabric Forwarding Class Sets on page 1238](#)

QFX Series Software Features Overview

The following tables list the Juniper Networks QFX Series software features and the Juniper Networks Junos operating system (Junos OS) release in which they were introduced:

- [Table 3 on page 16](#)—Access Control Features
- [Table 4 on page 16](#)—Administration Features
- [Table 5 on page 16](#)—Class-of-Service (CoS) Features
- [Table 6 on page 17](#)—High Availability and Resiliency Features
- [Table 7 on page 19](#)—Interface Features
- [Table 8 on page 19](#)—IP Address Management Features
- [Table 9 on page 20](#)—Layer 2 Network Protocol Features
- [Table 10 on page 22](#)—Layer 3 Protocol Features
- [Table 11 on page 23](#)—Multicast Protocol Features
- [Table 12 on page 24](#)—Network Management and Monitoring Features
- [Table 13 on page 26](#)—Port Security Features
- [Table 14 on page 26](#)—QFabric System-Specific Features
- [Table 15 on page 27](#)—Security
- [Table 16 on page 28](#)—Storage and Fibre Channel Features
- [Table 17 on page 29](#)—System Management Features



NOTE: The command-line interface (CLI) on the QFX Series might display configuration statements that are not supported. However, configuring an unsupported statement on a device has no effect on the operation of the device. See [“QFX Series CLI Hierarchy” on page 192](#) to see which statements are supported on the QFX Series.

Table 101: Access Control Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
MAC RADIUS authentication	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
TACACS+	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
SSH authentication keys	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15

Table 102: Administration Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
System logging (syslog) over IPv4	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Licensing	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10

Table 103: CoS Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Class of service (CoS)—Class-based queuing with prioritization	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS—Multidestination	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
CoS support on link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Enhanced transmission selection (ETS)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Interface-specific CoS rewrite rules	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 103: CoS Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Port shaping and queue shaping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control (PFC)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority-based flow control enhancements	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Re-marking of bridged packets	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Priority remapping	Junos OS 12.3X50-D10	Not supported	Not supported	Not supported
Weighted random early detection (WRED) tail-drop profiles	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
802.3X Ethernet PAUSE autonegotiation enhancements	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Layer 3 ingress packet classification and egress rewrite rule class-of-service features	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Software buffer configurability	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 104: High Availability and Resiliency Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Graceful protocol restart for BGP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Graceful protocol restart for OSPF	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10

Table 104: High Availability and Resiliency Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Link aggregation groups (LAGs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Support for 32 members in a link aggregation group (LAG)	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Multichassis link aggregation	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10	Not supported	Not supported
Virtual Router Redundancy Protocol (VRRP)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Graceful restart for FIP snooping to recover gracefully from Ethernet protocol restarts and crashes	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Nonstop software upgrade (NSSU)	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Layer 3 unicast and multicast support for multichassis link aggregation	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 105: Interface Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Interface ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 105: Interface Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
VLAN-tagged Layer 3 logical interfaces	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 106: IP Address Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Static addresses	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IPv4 support for telnet	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081 .	Junos OS 12.2X50-D10. For information about Telnet support on this system, see “Understanding Telnet on the QFabric System” on page 6081 .
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, virtual routing instances, firewall filters, RADIUS, NTP, and SNMP, on network Node group	Not supported	Not supported	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
IPv6 support for neighbor discovery, SSH, Telnet, ping, traceroute, static routing, firewall filters, and SNMP on management interfaces	Junos OS 12.2X50-D20	Not supported	Not applicable	Not applicable

Table 106: IP Address Management Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IPv6 support for neighbor discovery, router advertisements, stateless autoconfiguration, link aggregation, SSH, Telnet, ping, traceroute, path MTU, static routing, dynamic routing (BGPv6, IS-ISv6, OSPFv3, RIPv6), graceful restart, BFD, virtual routers, firewall filters, SNMP, CoS, VRRPv3, Radius, TACACS+, AAA, NTP, and syslog, on network interfaces	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not applicable	Not applicable

Table 107: Layer 2 Network Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
802.1Q VLAN tagging	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
BPDU protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)
Jumbo frames on routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Link Layer Discovery Protocol (LLDP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Loop protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 107: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Root protection for spanning-tree protocols	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Spanning tree:	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10 (network Node group only)	Junos OS 12.2X50-D10 (network Node group only)
<ul style="list-style-type: none"> Spanning Tree Protocol (STP) Rapid Spanning Tree Protocol (RSTP) Multiple Spanning Tree Protocol (MSTP) VLAN Spanning Tree Protocol (VSTP) RSTP and VSTP concurrent configuration 				
Multiple VLAN Registration Protocol (MVRP)	Not supported	Not supported	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15
VLAN ranges	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Private VLANs	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Secondary VLAN trunk ports and promiscuous access ports for private VLANs	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Q-in-Q tunneling	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D15	Junos OS 13.1X50-D15

Table 107: Layer 2 Network Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
VLAN translation	Junos OS 12.1X49-D1	Junos OS 12.1X49-D1	Not supported	Not supported
Layer 2 Tunneling Protocol	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Reflective relay	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Proxy ARP	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported

Table 108: Layer 3 Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Bidirectional Forwarding Detection (BFD)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Border Gateway Protocol (BGP) A separate software license is required for BGP. See "Software Features That Require Licenses on the QFX Series" on page 87.	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
BGPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Intermediate System-to-Intermediate System (IS-IS) A separate software license is required for IS-IS. See "Software Features That Require Licenses on the QFX Series" on page 87.	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IS-ISv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Open Shortest Path First (OSPF) v2	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
OSPFv3	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 108: Layer 3 Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Per-packet load balancing (ECMP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Routing Information Protocol versions 1 and 2 (RIPv1 and RIPv2)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
RIPv6	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
Routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Routing options	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Static routes	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Virtual router routing instances for unicast protocols	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D20
Virtual router routing instances for multicast protocols	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 109: Multicast Protocol Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
IGMPv1 and v2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
IGMPv3	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
IGMPv1 and v2 snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP snooping with routed VLAN interfaces (RVIs)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IGMP querier	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported
IGMP filtering	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 109: Multicast Protocol Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Protocol Independent Multicast sparse mode (PIM SM)	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Protocol Independent Multicast source-specific multicast (PIM SSM)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Bidirectional Forwarding Detection (BFD) for PIM	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static RP	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Anycast RP	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast Source Discovery Protocol (MSDP)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Multicast VLAN registration	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 110: Network Management and Monitoring Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Junos Space	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Local port mirroring	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Remote port mirroring	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Remote port mirroring to IP address (GRE encapsulation)	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 110: Network Management and Monitoring Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Increased numbers of port-mirroring sessions	Not supported	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
RMON	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
sFlow monitoring technology	Junos OS 11.3R1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Simple Network Management Protocol version 1 (SNMPv1) and SNMPv2	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Simple Network Management Protocol version 3 (SNMPv3)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Not supported	Not supported
Juniper enterprise-specific SNMP Utility MIB	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Juniper enterprise-specific SNMP Fabric Chassis MIB	Not applicable	Not applicable	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)	Junos OS 12.2X50-D20 Junos OS 13.1X50-D10 (Added Director device information)
Juniper Class-of-Service MIB	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Uplink failure detection	Junos OS 11.3R4	Junos OS 12.2X50-D20	Not supported	Not supported
Junos OS automation script support	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system control plane debugging and diagnostics support	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Internal fabric OAM monitoring	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DHCP smart relay	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not supported	Not supported

Table 111: Port Security Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Automatic recovery for port error disable conditions	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
MAC move limiting	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Persistent MAC learning (sticky MAC)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Static ARP support	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Storm control (broadcast, unicast, and multicast)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DHCP snooping	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Dynamic ARP inspection (DAI)	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported

Table 112: QFabric System-Specific Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Aliasing of Node devices	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
High availability of QFabric system components	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Node groups	Not applicable	Not applicable	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 112: QFabric System-Specific Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
QFabric system component-level access and login classes	Not applicable	Not applicable	Junos OS 11.3X30.9	Junos OS 12.2X50-D10
Aliases for Director devices and Interconnect devices	Not applicable	Not applicable	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
QFabric system backup and recovery (software and configuration)	Not applicable	Not applicable	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 113: Security

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Firewall filters and rate limiting	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on LAGs	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on loopback interfaces	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Firewall filters on management interfaces	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Not supported	Not supported
Policing	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increased numbers of supported filters and policers	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10
Filter-based forwarding	Junos OS 12.2X50-D20	Junos OS 12.2X50-D20	Not supported	Not supported

(For a list of supported firewall filter match conditions and actions, see [“Firewall Filter Match Conditions and Actions” on page 4644.](#))

Table 114: Storage and Fibre Channel Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Data center bridging technologies	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Data Center Bridging Capability Exchange protocol (DCBX)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
FCoE Initialization Protocol (FIP)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Fibre Channel fabrics	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel interfaces	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel over Ethernet (FCoE) to Fibre Channel gateway	Junos OS 11.1R1	Not Supported	Junos OS 11.3X30.6 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
FCoE-FC gateway link automated load balancing	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
FCoE OxID hash control	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
FIP snooping	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Increase to 2500 FIP snooping sessions	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported
Increase to 2500 FIP snooping sessions when QFX3500 switch acts as an FCoE-FC gateway.	Junos OS 12.3X50-D10	Not Supported	Not Supported	Not Supported
Transit switch	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
DCBX application protocol TLV exchange	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3R1	Junos OS 12.2X50-D10

Table 114: Storage and Fibre Channel Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
User-Configurable DCBX application protocol TLV exchange	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
DCBX version support for IEEE DCBX and DCBX version 1.01	Junos OS 12.2X50-D10	Junos OS 12.2X50-D20	Junos OS 12.2X50-D10	Junos OS 12.2X50-D10
Graceful restart for FCoE-FC gateway	Junos OS 12.1X49-D1	Not supported	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
Fibre Channel load balancing, maximum session limit, disabling fabric WWN verification check	Junos OS 12.1X49-D1	Not supported	Junos OS 12.1X49-D1 (supported only on QFX3500 Node device)	Junos OS 12.2X50-D10 (supported only on QFX3500 Node device)
VN_Port to VN_Port FIP snooping to enable configuring virtual links between VN_Ports without sending traffic through an FCoE forwarder (FCF)	Junos OS 12.2X50-D10	Not supported	Junos OS 13.1X50-D10	Junos OS 13.1X50-D10

Table 115: System Management Features

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
Configuration rollback	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
IP directed broadcast	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10
Online insertion and removal (OIR)	Junos OS 11.1R1	Junos OS 12.2X50-D20	Junos OS 11.3X30.6	Junos OS 12.2X50-D10

Table 115: System Management Features (*continued*)

Feature	QFX3500 Switch	QFX3600 Switch	QFX3000-G QFabric System	QFX3000-M QFabric System
J-Web interface, for switch configuration and management	Junos OS 12.1X49-D1	Junos OS 12.2X50-D20	Not supported	Not supported
Zero Touch Provisioning	Junos OS 12.3X50-D10	Junos OS 12.3X50-D10	Not Supported	Not Supported

Related Documentation

- [QFX3000-G QFabric System Hardware Documentation](#)
- [QFX3000-M QFabric System Hardware Documentation](#)
- [QFX3500 Device Hardware Documentation](#)
- [QFX3600 Device Hardware Documentation](#)

Understanding Software Upgrade on the QFabric System

The QFabric system software package contains software for the QFabric system infrastructure and for all of the different component devices in the QFabric system: Director group, Interconnect devices, and Node devices.

- [Operational Software Commands on page 1217](#)
- [Operational Reboot Commands on page 1218](#)

Operational Software Commands

The **request system software download** CLI command enables you to download the software package to various locations: for example, USB device, remote server, or FTP site.

The following CLI commands enable you to install the software for the Director group, Interconnect devices, Node devices, and the QFabric system infrastructure. You may need to specify the **reboot** option depending on which devices or QFabric infrastructure you are installing the software. The **reboot** option works differently depending on whether you install the software on the QFabric system infrastructure or on a particular device in the QFabric system.

- **request system software add component all**

This command installs software for the Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect devices, and network and server Node groups.

- **request system software add component director-group**

This command installs software for the Director group and the default partition, which is where you access the QFabric system CLI.

- **request system software add component fabric**

This command installs the software for the fabric control Routing Engines and the Interconnect devices.

- **request system software add component *node-group-name***

This command installs software for a server Node group or a network Node group.

Additionally, you can back up your current QFabric configuration file and installation-specific parameters using the **request system software configuration-backup** command. We recommend that you save this file to an external location, like an FTP site or USB device, but you can save it locally.

Operational Reboot Commands

The following commands enable you to reboot the entire QFabric system, various Node devices, or the QFabric system infrastructure:

- **request system reboot all**

This command reboots the Director group, fabric control Routing Engines, fabric manager Routing Engine, Interconnect devices, and network and server Node groups.

- **request system reboot director-group**

This command reboots the Director group and the default partition, which is where you access the QFabric system CLI.

- **request system reboot fabric**

This command reboots the fabric control Routing Engines and the Interconnect devices.

- **request system reboot node-group**

This command reboots a server Node group or a network node group.

Related Documentation

- [Upgrading Software on a QFabric System on page 134](#)

Understanding Nonstop Software Upgrade for QFabric Systems

The framework that underlies a nonstop software upgrade in a QFabric system enables you to upgrade the system in a step-by-step manner and minimize the impact to the continuous operation of the system. This topic explains how a nonstop software upgrade works in a QFabric system, the steps that are involved, and the procedures that you need to implement to experience the benefits of this style of software upgrade.

Nonstop software upgrade enables some QFabric system components to continue operating while similar components in the system are being upgraded. In general, the QFabric system upgrades redundant components in stages so that some components remain operational and continue forwarding traffic while their equivalent counterparts upgrade to a new version of software.



TIP: Use the following guidelines to decide when to implement a nonstop software upgrade:

- If you need to upgrade all components of the system in the shortest amount of time (approximately one hour) and you do not need to retain the forwarding resiliency of the data plane, issue the `request system software add component all` command to perform a standard software upgrade. All components of the QFabric system upgrade simultaneously and expediently, but this type of upgrade does not provide resiliency or switchover capabilities.
- If you need to minimize service impact, preserve the forwarding operations of the data plane during the upgrade, and are willing to take the extra time required for component switchovers (in many cases, several hours), issue the three nonstop software upgrade commands (`request system software nonstop-upgrade director-group | fabric | node-group`) described in this topic in the correct order.



NOTE:

- Before you begin a nonstop software upgrade, issue the `request system software download` command to copy the software to the QFabric system.
- Each of the 3 nonstop software upgrade steps must be considered parts of the whole process. You must complete all 3 steps of a nonstop software upgrade in the correct order to ensure the proper operation of the QFabric system.
- Open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected. For more information on verification of the upgrade, see [“Verifying Nonstop Software Upgrade for QFabric Systems” on page 1410](#).
- Issue the `show fabric administration inventory` command to verify that all upgraded components are operational at the end of a step before beginning the next step.
- Once you start the nonstop software upgrade process, we strongly recommend that you complete all 3 steps within 12 hours.

The three steps to a successful nonstop software upgrade must be performed in the following order:

- **Director group**—The first step upgrades the Director devices, the fabric manager Routing Engine, and the diagnostic Routing Engine. To perform the first step, issue the `request system software nonstop-upgrade director-group` command. The key actions that occur during a Director group upgrade are:

1. Connecting to the QFabric system by way of an SSH connection. This action establishes a load-balanced CLI session on one of the Director devices in the Director group.
2. The QFabric system downloads and installs the new software in both Director devices.
3. The Director device hosting the CLI session becomes the master for all QFabric system processes running on the Director group, such as the fabric manager and network Node group Routing Engines.
4. The QFabric system installs the new software for the backup fabric manager Routing Engine on the backup Director device.
5. The backup Director device reboots to activate the new software.
6. The master Director device begins a 15 minute sequence that includes a temporary suspension of QFabric services and a QFabric database transfer. You cannot issue operational mode commands in the QFabric CLI during this period.
7. The QFabric system installs the new software for the fabric manager and diagnostic Routing Engines on the Director group master.
8. The QFabric system switches mastership of all QFabric processes from the master Director device to the backup Director device.
9. The master Director device reboots to activate the new software.
10. The CLI session terminates, and logging back in to the QFabric system with a new SSH connection establishes the session on the new master Director device (the original backup).
11. The previous master Director device resumes operation as a backup and the associated processes (such as the fabric manager and network Node group Routing Engines) become backup as well. The fabric control Routing Engine associated with this Director device returns to active status.



NOTE: After the Director group nonstop software upgrade completes, any Interconnect device or Node device that reboots will automatically download the new software, install it, and reboot again. As a result, try not to restart any QFabric system devices before you complete the rest of the nonstop software upgrade steps.



TIP:

- To enable BGP and OSPF to continue operating on the network Node group during a Director group nonstop service upgrade, we recommend that you configure graceful restart for these routing protocols. For more information on graceful restart, see [“Configuring Graceful Restart for QFabric Systems” on page 1375](#).
- Wait 15 minutes after the second Director device returns to service and hosts Routing Engine processes before proceeding to step 2—the fabric

upgrade. You can verify the operational status of both Director devices by issuing the **show fabric administration inventory director-group status** command. Also, issue the **show fabric administration inventory infrastructure** command to verify when the Routing Engine processes become load balanced (typically, there will be three to four Routing Engines running on each Director device).

- Fabric—The second step upgrades the Interconnect devices and the fabric control Routing Engines. To perform the second step, issue the **request system software nonstop-upgrade fabric** command. The key actions that occur during a fabric upgrade are:
 1. The QFabric system downloads, validates, and installs the new software in all Interconnect devices and fabric control Routing Engines (FC-0 and FC-1).
 2. One fabric control Routing Engine reboots and comes back online.
 3. The other fabric control Routing Engine reboots and comes back online.
 4. The first Interconnect device reboots, comes back online, and resumes the forwarding of traffic.
 5. Subsequent Interconnect devices reboot one at a time, come back online, and return to service.



NOTE:

- If the software does not load properly on any one of the fabric components, all components revert back to the original software version.
- If one of the components in a fabric upgrade does not reboot successfully, issue the **request system reboot fabric** command to reattempt the rebooting process for this fabric component and activate the new software.

- Node group—The third and final step upgrades Node groups. You can choose to upgrade a network Node group, a redundant server Node group, or individual server Node groups. You can upgrade the Node groups one at a time or in groups (known as upgrade groups). However, you must upgrade all Node groups in your QFabric system before you can complete the nonstop software upgrade process. To perform the third step, issue the **request system software nonstop-upgrade node-group** command.

The key actions that occur during a network Node group upgrade are:

1. The QFabric system copies the new software to each Node device one at a time.
2. The QFabric system validates and then installs the new software in all Node devices simultaneously.
3. The system copies the software to the network Node group Routing Engines.
4. The QFabric system validates and then installs the software in the network Node group Routing Engines one at a time -- first the backup, then the master.

5. The backup network Node group Routing Engine reboots and comes back online.
6. The supporting Node devices reboot and come back online one at a time.



NOTE: To reduce the total upgrade duration, configure an upgrade group. All Node devices within the upgrade group reboot at the same time.

7. The master network Node group Routing Engine relinquishes mastership to the backup, reboots, and comes back online.

The key actions that occur during a redundant server Node group upgrade are:

1. The QFabric system copies the new software to the backup Node device, then the master Node device.
2. The QFabric system validates and then installs the new software on the backup Node device, then the master Node device.
3. The backup Node device reboots, comes back online, and becomes the master Node device.
4. The previous master Node device reboots and comes back online as a backup Node device.



NOTE: For redundant server Node groups, both Node devices must be online before the upgrade will proceed. If one of the devices is no longer available, remove the Node device from the Node group configuration before you issue the nonstop software upgrade command.

The key actions that occur during a server Node group upgrade for a Node group that contains one member are:

1. The Node device downloads the software package and validates the software.
2. The Node device installs the software and reboots.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [request system software add on page 394](#)
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)

Understanding Statements and Commands on the QFabric System

- [Chassis Statements on page 1223](#)
- [Chassis Commands on page 1223](#)

Chassis Statements

The following chassis statements enable you to configure various options for your Interconnect devices, Node groups (network and server), and Node devices:

- **interconnect-device**
- **node-group**
- **node-device**

Chassis Commands

The Junos OS CLI contains additions to the existing chassis commands. These additions reflect new options as a result of adding the **interconnect-device**, **node-group**, and **node-device** chassis statements at the **[edit chassis]** hierarchy level.

The following chassis commands enable you to monitor and configure the QFabric system hardware and software options at various hierarchy levels:

- **clear chassis display message**
- **request chassis beacon**
- **request chassis cb** (QFX3000-G QFabric systems only)
- **request chassis fabric** (QFX3000-G QFabric systems only)
- **request chassis fpc**
- **request chassis routing-engine master**
- **set chassis aggregated-devices**
- **set chassis alarm**
- **set chassis container-devices**
- **set chassis craft-lockout**
- **set chassis display**
- **set chassis fpc**
- **set chassis routing-engine**
- **show chassis alarms**
- **show chassis beacon**
- **show chassis environment**
- **show chassis fan** (QFX3000-G QFabric systems only)
- **show chassis fabric**

- **show chassis firmware**
- **show chassis fpc**
- **show chassis hardware**
- **show chassis lcd**
- **show chassis led**
- **show chassis location**
- **show chassis mac-addresses**
- **show chassis nonstop-upgrade**
- **show chassis pic**
- **show chassis routing-engine**
- **show chassis temperature-thresholds**
- **show chassis zones**

Understanding NTP on the QFabric System

Network Time Protocol (NTP) enables you to synchronize the time across the network. This is especially helpful for correlating log events and replicating databases and file systems. The QFabric system synchronizes time with servers that are external to the system and operates in client mode only.

To configure NTP, include the **server address** and **authentication-key** statements at the **[edit system ntp]** hierarchy level.

Understanding Network Management Implementation on the QFabric System

This topic describes network management features on the QFabric system that are implemented differently than on other devices running Junos OS.

The following network management features are supported on the QFabric system:

- **System log messages**—The QFabric system monitors events that occur on its component devices, distributes system log messages about those events to all external system log message servers (hosts) that are configured, and archives the messages. Component devices include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. You configure system log messages at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view messages.
- **Simple Network Management Protocol (SNMP) Version 1 (v1) and v2c**—SNMP monitors network devices from a central location. The SNMP implementation on the QFabric system supports the basic SNMP architecture of Junos OS with some limitations, including a reduced set of MIB objects, read-only access for SNMP communities, and limited support for SNMP requests. You configure SNMP at the **[edit snmp]** hierarchy level. Only the **show snmp statistics** operational mode command is supported, but you can issue SNMP requests using external SNMP client applications.

- **Advanced Insight Solutions (AIS)**—AIS provides tools and processes to automate the delivery of support services for the QFabric system. AIS components include Advanced Insight Scripts (AI-Scripts) and Advanced Insight Manager (AIM). You install AI-Scripts using the **request system scripts add** operational mode command. However, the **jais-activate-scripts.slax** file used during installation is preconfigured for the QFabric system and cannot be changed.



NOTE: Do not install Junos Space and AIS on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

Related Documentation

- [Advanced Insight Scripts \(AI-Scripts\) Release Notes](#)
- [Understanding Device and Network Management Features on page 6077](#)
- [Overview of Junos OS System Log Messages on page 6154](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
- [SNMP MIBs Support on page 6124](#)

Understanding the Implementation of SNMP on the QFabric System

SNMP monitors network devices from a central location. The QFabric system supports the basic SNMP architecture of Junos OS, but its implementation of SNMP differs from that of other devices running Junos OS. This topic provides an overview of the SNMP implementation on the QFabric system.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent resides in the QFabric Director software and is responsible for receiving and distributing all traps as well as responding to all the queries of the SNMP manager. For example, traps that are generated by a Node device are sent to the SNMP agent in the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.



NOTE: In its SNMP implementation, the QFabric system acts as an SNMP proxy server, and requires more time to process SNMP requests than a typical Junos OS device does. The default timeout setting on most SNMP client applications is 3 seconds, which is not enough time for the QFabric system to respond to SNMP requests, so the results of your **mibwalk** command may be incomplete. For this reason, we recommend that you change the SNMP timeout setting to 5 seconds or longer for the QFabric system to complete the responses to your requests.

Support for SNMP on the QFabric system includes:

- Support for the SNMP Version 1 (v1) and v2.



NOTE: Only SNMPv2 traps are supported on the QFabric system.

- Support for the following standard MIBs:
 - RFC 1155, *Structure and Identification of Management Information for TCP/IP-based Internets*
 - RFC 1157, *A Simple Network Management Protocol (SNMP)*
 - RFC 1212, *Concise MIB Definitions*
 - RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II* (partial support, including the system group and interfaces group)
 - RFC 1215, *A Convention for Defining Traps for use with the SNMP*
 - RFC 1901, *Introduction to Community-based SNMPv2*
 - RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
 - RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)*
 - RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol Using SMIv2*
 - RFC 2012, *SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2*
 - RFC 2013, *SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2*
 - RFC 2233, *The Interfaces Group MIB Using SMIv2*
 - RFC 2571, *An Architecture for Describing SNMP Management Frameworks* (read-only access) (excluding SNMPv3)
 - RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (read-only access) (excluding SNMPv3)
 - RFC 2576, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
 - RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
 - RFC 2579, *Textual Conventions for SMIv2*
 - RFC 2580, *Conformance Statements for SMIv2*
 - RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*
 - RFC 2863, *The Interfaces Group MIB*
 - RFC 3410, *Introduction and Applicability Statements for Internet Standard Management Framework* (excluding SNMPv3)

- RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Framework* (excluding SNMPv3)
- RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (excluding SNMPv3)
- RFC 3413, *Simple Network Management Protocol (SNMP) Applications* (excluding SNMPv3)
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)*
- RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 4188, *Definitions of Managed Objects for Bridges*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4363b, *Q-Bridge VLAN MIB*
- Support for the following Juniper Networks enterprise-specific MIBs:
 - Chassis MIB (mib-jnx-chassis.txt)
 - Class-of-Service MIB (mib-jnx-cos.txt)
 - Configuration Management MIB (mib-jnx-cfgmgmt.txt)
 - Fabric Chassis MIB (mib-jnx-fabric-chassis.txt)
 - Interface MIB Extensions (mib-jnx-if-extensions.txt)
 - Power Supply Unit MIB (mib-jnx-power-supply-unit.txt)
 - QFabric MIB (mib-jnx-qf-smi.txt)
 - Utility MIB (mib-jnx-util.txt)
- Support for operational mode commands—Limited to the **show snmp statistics** command. You may issue other SNMP requests, including **get**, **get next**, and **walk** requests, by using external SNMP client applications.

- Related Documentation**
- [SNMP MIBs Support on page 6124](#)
 - [SNMP Traps Support on page 6139](#)

Understanding the Implementation of System Log Messages on the QFabric System

This topic provides an overview of system log (syslog) messages as implemented on the QFabric system.

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers

(hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

You configure system log messages by using the **host** and **file** statements at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view the messages.



NOTE: On the QFabric system, a syslog file named **messages** with a size of 100 MB is configured by default. If you do not configure a filename, you can use the default filename **messages** with the **show log filename** command.

All messages with a severity level of **notice** or higher are logged. Messages with a facility level of **interactive-commands** on Node devices are not logged.

The QFabric system supports the following system log message features:

- The **file filename** and **host hostname** statements at the **[edit system syslog]** hierarchy level are supported. Other statements at that hierarchy level are not supported.
- You can specify the maximum amount of data that is displayed when you issue the **show log filename** command by configuring the **file filename archive maximum-file-size** statement.
- You can specify that one or more system log message servers receive messages, which are sent to each server that is configured.
- If you configured an alias for a device or interface, the alias is displayed in the message for the device or interface.
- The level of detail that is included in a message depends on the facility and severity levels that are configured. Messages include the highest level of detail available for the configured facility and severity levels.
- The unit of time is measured and displayed in seconds, and not milliseconds. If you attempt to configure the **time-format** option in milliseconds, the log output displays **000**.

Starting in Junos OS Release 13.1, the QFabric system supports these additional syslog features:

- You can filter the output of the **show log filename** operational mode command by device type and device ID or device alias when you specify the **device-type (device-id | device-alias)** optional parameters. Device types include **director-device**, **infrastructure-device**, **interconnect-device**, and **node-device**.
- You can specify the syslog structured data output format when you configure the **structured-data** statement at the **[edit system syslog file filename]** and **[edit system syslog host hostname]** hierarchy levels.



NOTE: Information displayed in the structured data output for system logs originating from the Director software may not be complete.

- You can filter the types of logs that the Director group collects from a component device when you configure the **filter all facility severity** or **filter all match “regular-expression”** statements at the **[edit system syslog]** hierarchy level.

Unsupported syslog features include:

- File access to syslog messages
- Monitoring of syslog messages

Related Documentation

- [Example: Configuring System Log Messages on page 1372](#)
- [syslog \(QFabric System\) on page 1403](#)

Understanding User and Access Management Features on the QFabric System

The QFabric system supports the following user and access management features:

- User authentication
- RADIUS
- Link Layer Discovery Protocol (LLDP)
- SSH
- TACACS+
- Access privilege management

The specific functionality, features, options, syntax, and hierarchy levels of some of the user and access management commands and configuration statements implemented on the QFabric system may differ somewhat from the same commands and configuration statements on standard Junos OS. See the configuration statement or command topic in the documentation set for additional information, and use the help (?) command-line function to display specific information as needed.

Some user and access management features are not yet fully supported in the full QFabric architecture, although full support is planned for future releases. The user and access management features currently unsupported on the QFabric system include:

- Full RADIUS server support, including RADIUS accounting
- **accounting-options** configuration statement hierarchy
- **tacplus-options** configuration statement

Understanding QFabric System Login Classes

In some cases (such as device-level troubleshooting), it is useful to log in to individual QFabric system components so you can view and manage issues on a per-device basis. This topic explains the login classes that provide individual component access within a QFabric system.



NOTE: Under normal operating conditions, you should manage the QFabric system as a single entity by using the QFabric system default partition command-line interface (CLI). The default partition CLI provides you with the ability to configure and monitor your entire QFabric system from a central location and should be used as the primary way to manage the system.

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system:

- **qfabric-admin**—Provides the ability to log in to individual QFabric system components and manage them. This class is equivalent to setting the following permissions: **access, admin, clear, firewall, interface, maintenance, network, reset, routing, secret, security, snmp, system, trace, and view**. The *qfabric-admin* class also enables you issue all operational mode commands except **configure**. To provide QFabric system component-level login and management privileges, include the **qfabric-admin** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.
- **qfabric-operator**—Provides the privilege to log in to individual QFabric system components and view component operations and configurations. This class is equivalent to setting the following permissions: **trace** and **view**. The *qfabric-operator* class also enables you issue the **monitor** and **show log messages** operational mode commands. To provide limited QFabric system component-level access, include the **qfabric-operator** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.
- **qfabric-user**—Prevents access to individual QFabric system components. This class is the default setting for all QFabric system users and is equivalent to the preset Junos OS class of **unauthorized**. To prevent a user from accessing individual QFabric system components, include the **qfabric-user** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.

When you perform the initial setup for the Director group, you must specify a username and password for QFabric components. Once configured, this information is stored in the QFabric system and mapped to the QFabric system login classes. Such mapping allows users with the proper login class (**qfabric-admin** or **qfabric-operator**) to log in automatically to a component without being prompted for the username and password.

After you assign the **qfabric-admin** or **qfabric-operator** class to a user, the user can log in to an individual QFabric system component by issuing the **request component login *component-name*** command. You can access Node devices, Interconnect devices, and virtual Junos Routing Engines (diagnostics, fabric control, and fabric manager) one at a

time when you issue this command. To leave the CLI prompt of a component and return to the QFabric system default partition CLI, issue the **exit** command from the component's operational mode CLI prompt.

Related Documentation

- [Example: Configuring QFabric System Login Classes on page 1345](#)
- [remote-debug-permission on page 1401](#)
- [request component login on page 1444](#)
- [Junos OS Login Classes Overview on page 1683](#)

Understanding Interfaces on the QFabric System

This topic describes:

- [Four-Level Interface Naming Convention on page 1231](#)
- [QSFP+ Interfaces on page 1231](#)
- [Link Aggregation on page 1234](#)

Four-Level Interface Naming Convention

When you configure an interface on the QFabric system, the interface name needs to follow a four-level naming convention that enables you to identify an interface as part of either a Node device or a Node group. Include the name of the network or server Node group at the beginning of the interface name.

The four-level interface naming convention is:

device-name:type-fpc/pic/port

where *device-name* is the name of the Node device or Node group. The remainder of the naming convention elements are the same as those in the QFX3500 switch interface naming convention.

An example of a four-level interface name is:

node2:xe-0/0/2

QSFP+ Interfaces

The QFX3500 Node device provides four 40-Gbps QSFP+ (quad small form-factor pluggable plus) interfaces (labeled **Q0** through **Q3**) for uplink connections between your Node device and your Interconnect devices.

The QFX3600 Node device provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight interfaces (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and interfaces Q2 through Q15 for 10-Gigabit Ethernet connections to either endpoint systems or external

networks (see [“Configuring the Port Type on QFX3600 Node Devices” on page 1367](#)).
[Table 98 on page 1174](#) shows the port mappings for QFX3600 Node devices.

Table 116: QFX3600 Node Device Port Mappings

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q0	fte-0/1/0	Not supported on this port
Q1	fte-0/1/1	Not supported on this port
Q2	fte-0/1/2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11
Q3	fte-0/1/3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15
Q4	fte-0/1/4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19
Q5	fte-0/1/5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23
Q6	fte-0/1/6	xe-0/0/24 xe-0/0/25 xe-0/0/26 xe-0/0/27
Q7	fte-0/1/7	xe-0/0/28 xe-0/0/29 xe-0/0/30 xe-0/0/31

Table 116: QFX3600 Node Device Port Mappings (*continued*)

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q8	Not supported on this port	xe-0/0/32
		xe-0/0/33
		xe-0/0/34
		xe-0/0/35
Q9	Not supported on this port	xe-0/0/36
		xe-0/0/37
		xe-0/0/38
		xe-0/0/39
Q10	Not supported on this port	xe-0/0/40
		xe-0/0/41
		xe-0/0/42
		xe-0/0/43
Q11	Not supported on this port	xe-0/0/44
		xe-0/0/45
		xe-0/0/46
		xe-0/0/47
Q12	Not supported on this port	xe-0/0/48
		xe-0/0/49
		xe-0/0/50
		xe-0/0/51
Q13	Not supported on this port	xe-0/0/52
		xe-0/0/53
		xe-0/0/54
		xe-0/0/55
Q14	Not supported on this port	xe-0/0/56
		xe-0/0/57
		xe-0/0/58
		xe-0/0/59

Table 116: QFX3600 Node Device Port Mappings (*continued*)

Port Number	40-Gigabit Data Plane Uplink Interfaces	10-Gigabit Ethernet Interfaces
Q15	Not supported on this port	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63

Link Aggregation

Link aggregation enables you to create link aggregation groups across Node devices within a network Node group or redundant server Node group. You can include up to 32 Ethernet interfaces in a LAG. You can have up to 48 LAGs within a redundant server Node group, and 128 LAGs in a network Node group. To configure a LAG, include the **aggregated-devices** statement at the **[edit chassis node-group node-group-name]** hierarchy level and the **device-count** statement at the **[edit chassis node-group node-group-name aggregated-devices ethernet]** hierarchy level. Additionally, include any aggregated Ethernet options (**minimum-links** and **link-speed**) at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level and the **802.3ad** statement at the **[edit interfaces interface-name ether-options]** hierarchy level. To configure the Link Aggregation Control Protocol (LACP), include the **lacp** statement at the **[edit interfaces aggregated-ether-options]** hierarchy level.

Related Documentation

- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)

Understanding Layer 3 Features on the QFabric System

The QFabric system supports the following Layer 3 features:

- Static routes, which enable you to manually configure and enter routes directly into the routing table.
- Routed VLAN Interfaces, which are a special type of Layer 3 virtual interface that enable you to forward packets between VLANs without using a router to connect the VLANs. Using this approach to connect VLANs reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.
- Routing protocols for routing traffic. The following routing protocols are supported on QFabric systems:
 - Border Gateway Protocol (BGP), which is an exterior gateway protocol (EGP) for routing traffic between autonomous systems (ASs).
 - Open Shortest Path First (OSPF) protocol, which is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). QFabric systems support OSPFv1 and OSPFv2.

**NOTE:**

- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.
- You can configure routing protocols by including statements at the [edit protocols] hierarchy level. If you want to isolate customer traffic on your network, you can configure virtual router routing instances at the [edit routing-instances] hierarchy level, and configure routing protocols for each virtual router routing instance by including statements at the [edit routing-instances *routing-instance-name* protocols] hierarchy level.

**Related
Documentation**

- [Understanding Virtual Router Routing Instances on page 2822](#)

Understanding Security Features on the QFabric System

The QFabric system supports the following security features:

- Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received).
- Policing (rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service. You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.
- MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:
 - Maximum number of MAC addresses—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
 - Allowed MAC—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN.

If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

- Storm control causes a switch to monitor traffic levels and take a specified action when a specified traffic level—called the storm control level—is exceeded, thus preventing packets from proliferating and degrading service. You can configure switches to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

Understanding Port Mirroring on the QFabric System

Port mirroring copies unicast packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Understanding Fibre Channel Fabrics on the QFabric System

A Fibre Channel (FC) fabric on a QFabric system is a construct that you configure on a QFX3500 Node device when the Node device is in FCoE-FC gateway mode. The FC fabric on a QFabric Node device is not the same as an FC fabric on a storage area network (SAN). The FC fabric on a QFabric Node device is local to that particular node device. We call the FC fabric on a QFabric Node device a *local FC fabric* to differentiate it from an FC fabric on the SAN.



NOTE: The QFX3600 Node device does not support FC or FCoE features.

A local FC fabric does not span Node devices and does not span the fabric Interconnect device. Local FC fabrics are entirely contained on a single Node device. A local FC fabric creates associations that connect FCoE devices that have converged network adapters (CNAs) on the Ethernet network to an FC switch or FCoE forwarder (FCF) on the FC network. A local FC fabric consists of:

- A unique fabric name.
- A unique fabric ID.
- One or more FCoE VLAN interfaces that include one or more 10-Gigabit Ethernet interfaces connected to FCoE devices. The FCoE VLANs transport traffic between the FCoE servers and the FCoE-FC gateway. Each FCoE VLAN must carry only FCoE traffic. You cannot mix FCoE traffic and standard Ethernet traffic on the same VLAN.

The 10-Gigabit Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic because FIP VLAN discovery and notification frames are exchanged as untagged packets.

Each FCoE VLAN interface can present multiple VF_Port interfaces to the FCoE network.

- One or more native FC interfaces. The native FC interfaces transport traffic between the gateway and the FC switch or FCF.



TIP: If the network does not use a dual-rail architecture for redundancy, configure more than one native FC interface for each local FC fabric to create redundant connections between the FCoE devices and the FC network. If one physical link goes down, any sessions it carried can log in again and connect to the FC network on a different interface.

All of the FC and FCoE traffic that belongs to a local FC fabric on a Node device must enter and exit that Node device. This means that the FC switch or FCF and the FCoE devices in the Ethernet network must be connected to the same Node device. The interfaces that connect to the FC switch and the interfaces that connect to the FCoE devices must be included in the local FC fabric. You cannot configure a local FC fabric that spans more than one Node device.

Traffic flows from FC and FCoE devices that are not in the same local FC fabric remain separate and cannot communicate with each other through the FCoE-FC gateway.



NOTE: The QFabric system enforces commit checks to ensure that local FC fabrics and FCoE VLANs on FCoE-FC gateways do not span more than one Node device.

Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Understanding CoS Fabric Forwarding Class Sets

Fabric forwarding class sets (fabric fc-sets) are similar to the fc-sets (priority groups) you configure on Node devices. The major differences are:

1. Fabric fc-sets group traffic for transport across the QFX3008-I or QFX3600-I Interconnect device (the fabric). Node device fc-sets group traffic on a Node device for transport across that Node device.
2. Fabric fc-sets are global. They apply to the entire fabric. Node device fc-sets apply only to the Node device on which they are configured.
3. You can configure class of service (CoS) scheduling for Node device fc-sets, but you cannot configure CoS for fabric fc-sets.
4. Fabric fc-sets map to Interconnect device fabric output queues statically—you cannot configure the mapping of fabric fc-sets to fabric output queues. All traffic in a fabric fc-set maps to the same output queue.

Node device fc-sets include forwarding classes that map to Node device output queues, and you can configure the mapping of forwarding classes to output queues (or you can use the default mapping). Because output queues are mapped to forwarding classes, different classes of traffic in a Node device fc-set can be mapped to different output queues.

Node device fc-sets consist of forwarding classes containing traffic that requires similar CoS treatment. (Forwarding classes are default forwarding classes or user-defined forwarding classes.) You can configure CoS for each fc-set to determine how the traffic of its forwarding classes is scheduled on a Node device.

When traffic exits a Node device interface and enters an Interconnect device fabric interface, the Interconnect device uses the same forwarding classes to group traffic. The forwarding classes are mapped to global fabric fc-sets for transport across the fabric. Like fc-sets on a Node device, fabric fc-sets also contain traffic that requires similar CoS treatment. Unlike fc-sets on a Node device, you cannot configure CoS on fabric fc-sets.

Fabric fc-sets reside on the Interconnect device and are global to the QFabric system. Fabric fc-sets apply to all traffic that traverses the fabric. The mapping of forwarding classes to fabric fc-sets is global and applies to all forwarding classes with traffic that traverses the fabric from all connected Node devices. You can change the mapping of forwarding classes to fabric fc-sets. All mapping changes you make are global. For example, if you change the fabric fc-set to forwarding class mapping of the default best-effort forwarding class, then every Node device's best-effort forwarding class traffic that traverses the fabric is mapped to that fabric fc-set.

This topic describes:

- [Default Fabric Forwarding Class Sets on page 1239](#)
- [Fabric Forwarding Class Set Configuration and Implementation on page 1242](#)
- [Fabric Forwarding Class Set Scheduling \(CoS\) on page 1244](#)
- [Support for Flow Control and Lossless Transport Across the Fabric on page 1246](#)

- [Viewing Fabric Forwarding Class Set Information on page 1248](#)
- [Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences on page 1250](#)

Default Fabric Forwarding Class Sets

Interconnect devices have 12 default fabric fc-sets, including five visible default fabric fc-sets, four for unicast traffic and one for multdestination (multicast, broadcast, and destination lookup failure) traffic.

There are also seven hidden default fabric fc-sets. There are three hidden default fabric fc-sets for multdestination traffic that you can use if you want to map different multdestination forwarding classes to different multdestination fabric fc-sets. There are four hidden default fabric fc-sets for lossless traffic that you can use to map different lossless forwarding classes (priorities) to different lossless fabric fc-sets.

[Table 117 on page 1239](#) shows the default fabric fc-sets:

Table 117: Default Fabric Forwarding Class Sets

Fabric Forwarding Class Set Name	Characteristics
<code>fabric_fcset_be</code>	Transports best-effort unicast traffic across the fabric.
<code>fabric_fcset_strict_high</code>	Transports unicast traffic that has been configured with strict-high priority and in the network-control forwarding class across the fabric. This fabric fc-set receives as much bandwidth across the fabric as it needs to service the traffic in the group up to the entire fabric interface bandwidth. For this reason, exercise caution when mapping traffic to this fabric fc-set to avoid starving other traffic.
<code>fabric_fcset_noloss1</code>	Transports unicast traffic in the default fcoe forwarding class across the fabric.
<code>fabric_fcset_noloss2</code>	Transports unicast traffic in the default no-loss forwarding class across the fabric.
<code>fabric_fcset_noloss3</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
<code>fabric_fcset_noloss4</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
<code>fabric_fcset_noloss5</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.

Table 117: Default Fabric Forwarding Class Sets (*continued*)

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_noloss6	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_multicast1	Transports multdestination traffic in the mcast forwarding class across the fabric. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast2	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.

The five default forwarding classes (**best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast**) are mapped to the fabric fc-sets by default as shown in [Table 118 on page 1240](#).

Table 118: Default Forwarding Class to Fabric Forwarding Class Set Mapping

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
best-effort	fabric_fcset_be	0	NA
network-control	fabric_fcset_strict_high	7	NA
fcoe	fabric_fcset_noloss1	1	9K
no-loss	fabric_fcset_noloss2	2	9K
mcast	fabric_fcset_multicast1	8	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss3	3	9k

Table 118: Default Forwarding Class to Fabric Forwarding Class Set Mapping (*continued*)

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss4	4	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss5	5	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss6	6	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast2	9	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast3	10	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast4	11	NA

The maximum fiber cable length between the QFabric system Node device and the QFabric system Interconnect device is 150 meters.



TIP: If you explicitly configure lossless forwarding classes, we recommend that you map each user-configured lossless forwarding class to an unused fabric fc-set (fabric_fcset_noloss3 through fabric_fcset_noloss6) on a one-to-one basis: one lossless forwarding class mapped to one lossless fabric fc-set.

The reason for one-to-one mapping is to avoid fate sharing of lossless flows. Because each fabric fc-set is mapped statically to an output queue, when you map more than one forwarding class to a fabric fc-set, all of the traffic in all of the forwarding classes that belong to the fabric fc-set uses the same output queue. If that output queue becomes congested due to congestion caused by one of the flows, the other flows are also affected. (They share fate because the flow that congests the output queue affects flows that are not experiencing congestion.)

However, it is important to understand that fabric_fcset_noloss1 and fabric_fcset_noloss2 have a scheduling weight of 35, while the other fabric fc-sets have a scheduling weight of 1. The scheduling weights mean that fabric_fcset_noloss1 and fabric_fcset_noloss2 receive most of the bandwidth available to lossless fabric fc-sets if the amount of traffic on fabric_fcset_noloss1 and fabric_fcset_noloss2 requires the bandwidth.

If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will consume most of that bandwidth, then you should place all lossless traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2`. If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will **<emphasis>not</emphasis>** consume most of that bandwidth, then you can map lossless forwarding classes in a one-to-one manner to lossless fabric fc-sets to avoid fate sharing.

If you want to map different multidestination forwarding classes to different multidestination fabric fc-sets, use one or more of the hidden multidestination fabric fc-sets.



NOTE: The global mapping of forwarding classes to fabric fc-sets is independent of the mapping of forwarding classes to Node device fc-sets. Global mapping of forwarding classes to fabric fc-sets occurs only on the Interconnect device. The Node device mapping of forwarding classes to fc-sets does not affect the global mapping of forwarding classes to fabric fc-sets on the Interconnect device, and vice versa.

When you define new forwarding classes on a Node device, you explicitly map those forwarding classes to Node device fc-sets. However, new (user-created) forwarding classes are mapped by default to fabric fc-sets. (You can override the default mapping if you want to configure the forwarding class to fabric fc-set mapping explicitly, as described in the next section.)

By default:

- All best-effort traffic forwarding classes that you create are mapped to the **`fabric_fcset_be`** fabric fc-set.
- All lossless traffic forwarding classes that you create are mapped to the **`fabric_fcset_noloss1`** or **`fabric_fcset_noloss2`** fabric fc-set.
- All multidestination traffic forwarding classes that you create are mapped to the **`fabric_fcset_multicast1`** fabric fc-set.
- All **strict-high** priority traffic and **network-control** forwarding classes that you create are mapped to the **`fabric_fcset_strict_high`** fabric fc-set.

Fabric Forwarding Class Set Configuration and Implementation

You can map forwarding classes to fabric fc-sets, but no other attributes of fabric fc-sets are user-configurable, including CoS. This section describes:

- [Mapping Forwarding Classes to Fabric Forwarding Class Sets on page 1243](#)
- [Fabric Forwarding Class Set Implementation on page 1243](#)

Mapping Forwarding Classes to Fabric Forwarding Class Sets

If you do not want to use the default mapping of forwarding classes to fabric fc-sets, you can map forwarding classes to fabric fc-sets the same way as you map forwarding classes to Node device fc-sets. To do this, use exactly the same statement that you use to map forwarding classes to fc-sets, but instead of specifying a Node device fc-set name, specify a fabric fc-set name.



NOTE: The global mapping of forwarding classes to fabric fc-sets does not affect the mapping of forwarding classes to Node device fc-sets. The global forwarding class mapping to fabric fc-sets pertains to the traffic only when it enters, traverses, and exits the fabric. The forwarding class mapping to fc-sets on a Node device is valid within that Node device.

Mapping forwarding classes to fabric fc-sets does not affect the scheduling configuration of the forwarding classes or fc-sets on Node devices. Fabric fc-set scheduling (which is not user-configurable) pertains to traffic only when it enters, traverses, and exits the Interconnect device fabric.

If you change the mapping of a forwarding class to a fabric fc-set, the new mapping is global and applies to all traffic in that forwarding class, regardless of which Node device forwards the traffic to the Interconnect device.

- To assign one or more forwarding classes to a fabric fc-set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric-forwarding-class-set-name class
forwarding-class-name
```

For example, to map a user-defined forwarding class named **best-effort-2** to the fabric fc-set **fabric_fcset_be**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric_fcset_be class best-effort-2
```



NOTE: Because fabric fc-set configuration is global, in this example all forwarding classes with the name **best-effort-2** on all of the Node devices attached to the fabric use the **fabric_fcset_be** fabric fc-set to transport traffic across the fabric.

Fabric Forwarding Class Set Implementation

The following rules apply to fabric fc-sets:

- You cannot create new fabric fc-sets. Only the twelve default fabric fc-sets are available.
- You cannot delete a default fabric fc-set.
- You cannot attach a fabric fc-set to a Node device interface. Fabric fc-sets are used only on the Interconnect device fabric, not on Node devices.

- You can map only multdestination forwarding classes to multdestination fabric fc-sets.
- You cannot map multdestination forwarding classes to unicast fabric fc-sets.
- You cannot map unicast forwarding classes to multdestination fabric fc-sets.
- You cannot configure CoS for fabric fc-sets. (However, default CoS scheduling properties are applied to traffic on the fabric, and the fabric interfaces use link layer flow control (LLFC) for flow control.)

Fabric Forwarding Class Set Scheduling (CoS)

Although fabric fc-set CoS is not user-configurable, CoS is applied to traffic on the fabric. (In addition, fabric interfaces use LLFC to ensure lossless transport for lossless traffic flows.) This section describes how the fabric applies CoS scheduling to traffic:

- [Class Groups for Fabric Forwarding Class Sets on page 1244](#)
- [Class Group Scheduling on page 1244](#)
- [QFabric System CoS on page 1246](#)

Class Groups for Fabric Forwarding Class Sets

To transport traffic across the fabric, the QFabric system organizes the fabric fc-sets into three classes called *class groups*. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric_fcset_strict_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes and in any forwarding classes you create on a Node device that consist of **strict-high** priority or **network-control** forwarding class traffic.
- **Unicast**—All traffic in the fabric fc-sets **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes and in any forwarding classes you create on a Node device that consist of best-effort or lossless traffic. If you use any of the hidden no loss fabric fc-sets (**fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric_fcset_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric_fcset_multicast2**, **fabric_fcset_multicast3**, or **fabric_fcset_multicast4**), that traffic is also classified as part of this class group.

Class Group Scheduling

You cannot configure CoS for class groups or for fabric fc-sets (that is, you cannot attach a traffic control profile to a fabric fc-set—you attach traffic control profiles to Node device fc-sets to apply scheduling to the traffic that belongs to the Node device fc-set). By default, the fabric uses weighted round-robin (WRR) scheduling in which each class group receives a portion of the total available fabric bandwidth based on its type of traffic, as shown in [Table 119 on page 1245](#):

Table 119: Class Group Scheduling Properties and Membership

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Strict-high priority	fabric_fcset_strict_high	<ul style="list-style-type: none"> All strict-high priority forwarding classes network-control 	Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.
Unicast	<ul style="list-style-type: none"> fabric_fcset_be fabric_fcset_noloss1 fabric_fcset_noloss2 <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_noloss3 fabric_fcset_noloss4 fabric_fcset_noloss5 fabric_fcset_noloss6 	<ul style="list-style-type: none"> best-effort fcoe no-loss <p>NOTE: No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p>	Traffic in the unicast class group receives an 80% weight in the WRR calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	fabric_fcset_multicast1 <p>Includes the hidden multidestination fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_multicast2 fabric_fcset_multicast3 fabric_fcset_multicast4 	<ul style="list-style-type: none"> mcast <p>NOTE: No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)

The fabric fc-sets within each class group are weighted equally and receive bandwidth using round-robin scheduling. For example:

- If the unicast class group has three member fabric fc-sets, **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**, then each of the three fabric fc-sets receives one-third of the bandwidth available to the unicast class group.
- If the multidestination class group has one member fc-set, **fabric_fcset_multicast1**, then that fc-set receives all of the bandwidth available to the multidestination class group.
- If the multidestination class group has two member fc-sets, **fabric_fcset_multicast1** and **fabric_fcset_multicast2**, then each of the two fabric fc-sets receives one-half of the bandwidth available to the multidestination class group.

QFabric System CoS

When traffic enters and exits the same Node device, CoS works the same as it works on a standalone QFX3500 switch.

However, when traffic enters a Node device, crosses the Interconnect device, and then exits a different Node device, CoS is applied differently:

1. Traffic entering the ingress Node device receives the CoS configured at the Node ingress (packet classification, congestion notification profile for PFC).
2. When traffic goes from the ingress Node device to the Interconnect device, the fabric fc-set CoS is applied as described in the discussion of fabric forwarding class set scheduling.
3. When traffic goes from the Interconnect device to the egress Node device, the egress Node device applies CoS at the egress port (egress queue scheduling, WRED, IEEE 802.1p or DSCP code-point rewrite).

Support for Flow Control and Lossless Transport Across the Fabric

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the fabric interfaces use LLFC to support lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class on the Node devices.
2. The lossless forwarding class must be mapped to a lossless fabric fc-set on the Interconnect device (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**).

When traffic meets the two configuration constraints, the fabric propagates the back pressure from the egress Node device across the fabric to the ingress Node device during periods of congestion. However, to achieve end-to-end lossless transport across the switch, you must also configure a congestion notification profile to enable PFC on the Node device ingress ports.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.

- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



NOTE: Lossless transport across the fabric also must meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is a 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- **MTU size**—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- **Total number of Node device fabric ports connected to the Interconnect device**—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multdestination traffic.

Table 120 on page 1248 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 120 on page 1248, the QFabric system supports six lossless priorities.



NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 120: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5



NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

Viewing Fabric Forwarding Class Set Information

You can display information about fabric fc-sets using the same CLI command you use to display information about Node device fc-sets:

```
user@switch> show class-of-service forwarding-class-set
Forwarding class set: fabric_fcset_be, Type: fabric-type, Forwarding class set
index: 1
  Forwarding class      Index
  best-effort           0

Forwarding class set: fabric_fcset_mcast1, Type: fabric-type, Forwarding class
set index: 5
  Forwarding class      Index
  mcast                 8
```

Forwarding class set: fabric_fcset_mcast2, Type: fabric-type, Forwarding class set index: 6

Forwarding class set: fabric_fcset_mcast3, Type: fabric-type, Forwarding class set index: 7

Forwarding class set: fabric_fcset_mcast4, Type: fabric-type, Forwarding class set index: 8

Forwarding class set: fabric_fcset_noloss1, Type: fabric-type, Forwarding class set index: 2

Forwarding class	Index
fcoe	1

Forwarding class set: fabric_fcset_noloss2, Type: fabric-type, Forwarding class set index: 3

Forwarding class	Index
no-loss	2

Forwarding class set: fabric_fcset_noloss3, Type: fabric-type, Forwarding class set index: 9

Forwarding class set: fabric_fcset_noloss4, Type: fabric-type, Forwarding class set index: 10

Forwarding class set: fabric_fcset_noloss5, Type: fabric-type, Forwarding class set index: 11

Forwarding class set: fabric_fcset_noloss6, Type: fabric-type, Forwarding class set index: 12

Forwarding class set: fabric_fcset_strict_high, Type: fabric-type, Forwarding class set index: 4

Forwarding class	Index
network-control	3

[Table 121 on page 1249](#) describes the meaning of the **show class-of-service forwarding-class-set** output fields when you display fabric fc-set information.

Table 121: show class-of-service forwarding-class-set Command Output Fields

Field Name	Field Description
Forwarding class set	Name of the fabric forwarding class set.
Type	Type of forwarding class set: <ul style="list-style-type: none"> Fabric-type—Fabric fc-set Normal-type—Node device fc-set
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of the forwarding class.

Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences

Table 122 on page 1250 summarizes the differences between fabric fc-sets and Node device fc-sets:

Table 122: Summary of Differences Between Fabric fc-sets and Node Device fc-sets

Characteristic	Fabric fc-set	Node device fc-set
Location	QFX3008-I or QFX3600-I Interconnect device (the fabric).	QFabric system Node device.
Global or Node-device specific	Global, valid for the entire fabric.	Local to the Node device on which the fc-set is configured.
Ability to create (define) a new fc-set	No. Use the 12 default fabric fc-sets provided.	Yes.
Ability to configure CoS	Default CoS settings only. CoS is not user-configurable.	User-configurable using traffic control profiles.
Ability to map forwarding classes to an fc-set	Yes. Mapping is global and applies to all forwarding classes across Interconnect device fabric (traffic from all connected Node devices).	Yes. Mapping is local to a Node device and applies only to the forwarding classes on the Node device.

Related Documentation

- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [show class-of-service forwarding-class-set on page 5938](#)

Licenses

- [Junos OS Feature Licenses on page 1250](#)
- [Software Features That Require Licenses on the QFX Series on page 1251](#)
- [Junos OS License Keys on page 1253](#)

Junos OS Feature Licenses

Some Junos OS software features require a license to activate the feature. To enable a licensed feature, you need to purchase, install, manage, and verify a license key that

corresponds to each licensed feature. To conform to Junos OS feature licensing requirements, you must purchase one license per feature per device. The presence of the appropriate software license key on your device determines whether you are eligible to configure and use the licensed feature.

To speed deployment of licensed features, Junos OS software implements an honor-based licensing structure and provides you with a 30-day grace period to use a licensed feature without a license key installed. The grace period begins when you configure the feature and your device uses the licensed feature for the first time, but not necessarily when you install the license. After the grace period expires, the system generates system log messages saying that the feature requires a license. To clear the error message and use the licensed feature properly, you must install and verify the required license.

For information about how to purchase software licenses, contact your Juniper Networks sales representative.

**Related
Documentation**

- *License Enforcement*
- [Junos OS License Keys on page 88](#)
- *Software Feature Licenses*
- [Verifying Junos OS License Installation on page 96](#)

Software Features That Require Licenses on the QFX Series

The following Junos OS features require an Advanced Feature License (AFL) on QFX Series devices:



NOTE: If you try to configure a feature that is not licensed, you will receive syslog messages saying that you are using a feature that is licensable and that you do not possess a license for the feature. If you try to commit configuration changes for a feature that is not licensed, you will receive a commit warning saying that you have exceeded the allowed license limit for the feature.

- Border Gateway Protocol (BGP)
- Intermediate System-to-Intermediate System (IS-IS)
- Multiprotocol Label Switching (MPLS)
- Fibre Channel support

[Table 39 on page 87](#) lists the licenses you can purchase for each QFX Series software feature.

For information about how to purchase a software license, contact your Juniper Networks sales representative.

Table 123: Junos OS Feature Licenses and Model Numbers for QFX Series Devices

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Multiprotocol Label Switching (MPLS)	QFX3500, QFX3600, and QFX5100 switch	One per switch	QFX-JSL-EDGE-ADV1
QFX Series advanced feature license for Border Gateway Protocol (BGP)	QFX3100 Director device	One per Node device in a network Node group	QFX-JSL-DRCTR-ADV1
QFX Series advanced feature license for Fibre Channel	QFX3500 switch	One per switch on which fibre channel ports are configured	QFX-JSL-EDGE-FC
QFX Series advanced feature license for Fibre Channel	QFX3100 Director device	One per QFX3500 Node device on which fibre channel ports are configured	QFX-JSL-DRCTR-FC
QFX Series advanced feature license for Fibre Channel - Capacity 16	QFX3100 Director device	One for up to 16 QFX3500 Node devices on which fibre channel ports are configured	QFX-JSL-DRCTR-FC-C16
QFX Series feature license for enabling fabric mode	QFX3500 and QFX3600 device	One per device	QFX3000-JSL-EDGE-FAB
QFX Series feature license for base software for QFX3000-G QFabric system	QFX3100 Director device	One per QFX3000-G QFabric system	QFX3008-JSL-DRCTR-FAB
QFX Series feature license for base software for QFX3000-M QFabric system	QFX3100 Director device	One per QFX3000-M QFabric system	QFX3000M-JSL-DRCTR-FAB

Related Documentation

- [Junos OS Feature Licenses on page 86](#)
- [Junos OS License Keys on page 88](#)
- [Generating the License Keys for a Standalone QFX Series Device on page 90](#)
- [Generating the License Keys for a QFabric System on page 91](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)
- [Deleting a License \(CLI Procedure\) on page 94](#)
- [Saving License Keys on page 95](#)
- [Verifying Junos OS License Installation on page 96](#)

Junos OS License Keys

Licensable Ports on MX5, MX10, and MX40 Routers

Starting with Junos OS Release 12.2, license keys are available to enhance the port capacity on MX5, MX10, and MX40 routers up to the port capacity of an MX80 router. The MX5, MX10, and MX40 routers are derived from the modular MX80 chassis with similar slot and port assignments, and provide all functionality available on an MX80 router, but at a lower capacity. Restricting port capacity is achieved by making a set of MIC slots and ports licensable. MICs without a license are locked, and are unlocked or made usable by installing appropriate upgrade licenses.

The base capacity of a router is identified by the Ideeprom assembly ID (I2C ID), which defines the board type. However, the Junos OS licensing infrastructure allows the use of restricted ports without a license for a grace period of 30 days. After the grace period expires, the router reverts back to the base capacity if no upgrade license is purchased and installed for the locked ports. The I2C ID along with an upgrade license determine the final capacity of an MX5, MX10, or MX40 router.

The MX5, MX10, MX40, and MX80 routers support the following types of MICs:

- A built-in 10-Gigabit Ethernet MIC with four 10-Gigabit Ethernet ports
- Two front-pluggable MICs

A feature ID is assigned to every license upgrade for enhancing port capacity.

[Table 40 on page 89](#) displays the chassis types and their associated port capacity, I2C ID, base capacity, feature ID, feature name, and the final capacity after a license upgrade.

Table 124: Upgrade Licenses for Enhancing Port Capacity

Chassis Type	Port Capacity	I2C ID	Base Capacity	Feature ID and Feature Name	Upgrade Capacity
MX5	20G	0x556	Slot 1 <ul style="list-style-type: none"> • 1/MIC0 	f1—MX5 to MX10 upgrade	Slot 1 and 2 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1
MX10	40G	0x555	Slot 1 and 2 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1 	f2—MX10 to MX40 upgrade	Slot 2 and first 2 ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC1 • First 2 ports on 0/MIC0
MX40	60G	0x554	Slot 1, Slot 2 and first 2 ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC0 • 1/MIC1 • First 2 ports on 0/MIC0 	f3—MX40 to MX80 upgrade	Slot 2 and all ports on Slot 0 <ul style="list-style-type: none"> • 1/MIC1 • All 4 ports on 0/MIC0

When installing an upgrade license for enhancing port capacity on MX5, MX10 and MX40 routers, consider the following:

- To upgrade an MX5 router to MX80 router capacity, licenses for all three features (f1, f2, f3) must be installed. All three features can be provided in a single license key.
- To upgrade an MX10 router to MX40 router capacity, installing a license key with f2 feature is sufficient.
- Non-applicable feature IDs in a license key reject the upgrade license. For example:
 - An f1 feature ID on an MX10 upgrade license key rejects the license.
 - Feature IDs f1 and f2 on an MX40 upgrade license key reject the entire license.

**Related
Documentation**

- [Junos OS Feature Licenses on page 86](#)
- *License Enforcement*
- *Software Feature Licenses*
- [Verifying Junos OS License Installation on page 96](#)

CHAPTER 10

Configuration

- [Initial Setup on page 1255](#)
- [QFabric System Configuration on page 1340](#)
- [Configuration Statements on page 1378](#)

Initial Setup

- [QFabric System Initial and Default Configuration Information on page 1255](#)
- [Converting the Device Mode for a QFabric System Component on page 1258](#)
- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Importing a QFX3000-G QFabric System Control Plane Virtual Chassis Configuration with a USB Flash Drive on page 1308](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- [Importing a QFX3000-M QFabric System Control Plane EX4200 Switch Configuration with a USB Flash Drive on page 1333](#)
- [Generating the MAC Address Range for a QFabric System on page 1334](#)
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)

QFabric System Initial and Default Configuration Information

Once you install the hardware for the QFabric system, you can configure the Junos operating system (Junos OS) to begin using the system. This topic discusses which setup activities you need to perform and which activities are handled automatically by the QFabric system.

The fabric manager Routing Engine in the Director group automatically handles some of the initial setup activities, including:

- Assignment of IP addresses and unique identifiers to each QFabric system component by way of the management control plane
- Inclusion of all QFabric system devices within the default partition
- Establishment of interdevice communication and connectivity through the use of a fabric provisioning protocol and a fabric management protocol

The initial configuration tasks you need to perform to bring up the QFabric system and make it operational include:

- Converting any standalone devices, such as QFX3500 and QFX3600 devices, to Node device mode (see [“Converting the Device Mode for a QFabric System Component” on page 1258](#))
- Setting up the QFabric system control plane cabling, topology, and configuration
 - To set up the control plane cabling, topology, and configuration for the QFX3000-G QFabric system, see [“Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane” on page 1263](#).
 - To set up a copper or fiber-based control plane cabling, topology, and configuration for the QFX3000-M QFabric system, see [“Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane” on page 1309](#).
- Accessing the Director group through a console connection, turning on the devices, and running through the initial setup script (see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#)), which prompts you to:
 - Set IP addresses for the Director devices in the Director group.
 - Set an IP address for the default partition.
 - Add the software serial number for your QFabric system. (Review the e-mail containing the software serial number that you received from Juniper Networks when you purchased your QFabric system.)
 - Set the starting MAC address and the range of MAC addresses for the QFabric system. (See [“Generating the MAC Address Range for a QFabric System” on page 1334](#) for this information.)
 - Set a root password for the Director devices.
 - Set a root password for the QFabric system components, such as Node devices, Interconnect devices, and infrastructure.
- Logging into the default partition by using the IP address you configured when you ran the Director group initial setup script (See [“Gaining Access to the QFabric System Through the Default Partition” on page 1344](#))
- Configuring basic system settings for the default partition, such as time, location, and default gateways



NOTE: Unlike other Juniper Networks devices that run Junos OS, a QFabric system does not have a default factory configuration (containing the basic configuration settings for system logging, interfaces, protocols, and so on) that is loaded when you first install and power on the Director devices. Therefore, you must configure all the settings required for your QFabric system through the default partition CLI.

- Configuring aliases for Node devices (see [“Configuring Aliases for the QFabric System” on page 1353](#))

- Configuring VLANs and interfaces for the QFabric system devices
- Configuring redundant server Node groups to provide resiliency for server and storage connections (see [“Configuring Node Groups for the QFabric System” on page 1363](#))
- Configuring a network Node group to connect the QFabric system to external networks (see [“Configuring Node Groups for the QFabric System” on page 1363](#))
- Configuring the port type on QFX3600 Node devices (see [“Configuring the Port Type on QFX3600 Node Devices” on page 1367](#))
- Configuring routing protocols to run on the network Node group interfaces and reach external networks



NOTE: When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

- Generating and adding the license keys for the QFabric system (see [“Generating the License Keys for a QFabric System” on page 91](#) and [“Adding New Licenses \(CLI Procedure\)” on page 93](#))

Related Documentation

- *QFX3000-G QFabric System Installation Overview*
- *QFX3000-M QFabric System Installation Overview*
- [Converting the Device Mode for a QFabric System Component on page 1258](#)
- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- [Generating the MAC Address Range for a QFabric System on page 1334](#)
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Understanding QFabric System Administration Tasks and Utilities on page 1340](#)
- [Gaining Access to the QFabric System Through the Default Partition on page 1344](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)
- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)
- [Generating the License Keys for a QFabric System on page 91](#)
- [Adding New Licenses \(CLI Procedure\) on page 93](#)

Converting the Device Mode for a QFabric System Component

You can configure some devices to act as a standalone switch or participate in a QFabric system in a particular role. To change the role of your device, you must set the device mode. [Table 125 on page 1258](#) shows the device modes available for various devices.

Table 125: Support for device mode options

Device mode	QFX3500	QFX3600
Interconnect device	N/A	•
Node device	•	•
Standalone	•	•

To convert a device to a different mode, issue the **request chassis device-mode** command and specify the desired device mode. You verify the current and future device mode with the **show chassis device-mode** command.

When you convert a device from standalone mode to either Node device or Interconnect device mode, the software prepares the device to be configured automatically by the QFabric system. However, changing the device mode erases all configuration data on the device.



NOTE: The QFX3600 switch requires Jloader Release 1.1.8 before you can convert the switch to Interconnect device mode. For more information, see: [Jloader 1.1.8 Release for QFX-Series Platforms](#).



CAUTION: We recommend that you back up your device configuration to an external location before converting a device to a different device mode.

The following procedures illustrate the conversion options available when you modify a device mode:

- Convert from standalone switch mode to Node device mode
- Convert from Node device mode to Interconnect device mode
- Convert from Interconnect device mode to Node device mode
- Convert from Node device mode or Interconnect device mode to standalone switch mode

Standalone Switch to Node Device

To convert your device from standalone mode to Node device mode, follow these steps:

1. Connect to your standalone device through the console port and log in as the root user.
2. Upgrade the software on your standalone device to a QFabric system Node and Interconnect device software package that matches the QFabric system complete software package used by your QFabric system. For example, if the complete software package for your QFabric system is named **jinstall-qfabric-11.3X30.6.rpm**, you need to install the **jinstall-qfx-11.3X30.6-domestic-signed.tgz** package on your standalone device. Matching the two software packages ensures a smooth and successful addition of the device to the QFabric system inventory.

```
root@switch# request system software add software-package-name reboot
```

3. Back up your device configuration to an external location.

```
root@switch# save configuration-name external-path
```

4. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Standalone
```

5. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode node-device
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

6. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Node-device
```

7. Reboot the device.

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

```
Shutdown NOW!
[pid 34992]
```

```
root@switch>
```

```
*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```

8. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

9. To enable a converted device to participate in the QFabric system, locate the applicable network cables for your device and connect the device ports to the control plane and data plane.
10. (Optional) If you change the device back from Node device mode to standalone mode, restore the saved backup configuration from your external location.

```
root@switch# load configuration-name external-path
```

Node Device to Interconnect Device

To convert your device from Node device mode to Interconnect device mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to your device through the console port and log in as the root user.

3. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

4. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode interconnect-device
Device mode set to 'interconnect-device' mode.
Please reboot the system to complete the process.
```

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Interconnect-device
```

6. Reboot the device.

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

```
Shutdown NOW!
[pid 34992]
```

```
root@switch>
```

```
*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```

7. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

8. To enable a converted device to participate in the QFabric system in its new role, move the device to a different rack (as needed), locate the applicable network cables for your device, connect the device ports to the control plane and data plane per the design for your specific QFabric system, and reconfigure any aliases for the device at the QFabric default partition CLI prompt.

Interconnect Device to Node Device

To convert your device from Interconnect device mode to Node device mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to your device through the console port and log in as the root user.

3. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

4. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode node-device
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Node-device
```

6. Reboot the device.

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

```
Shutdown NOW!
[pid 34992]
```

```
root@switch>
```

```
*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```

7. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

8. To enable a converted device to participate in the QFabric system in its new role, move the device to a different rack (as needed), locate the applicable network cables for your device, connect the device ports to the control plane and data plane per the design for your specific QFabric system, and reconfigure any aliases for the device at the QFabric default partition CLI prompt.

**QFabric Component
(Interconnect or Node
Device) to Standalone
Switch**

To convert your QFabric component from either Interconnect device mode or Node device mode to standalone switch mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to the desired QFabric component through the console port of the device and log in as the root user.

3. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@node1> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

4. Issue the **request chassis device-mode standalone** command to convert the component to standalone switch mode, while the component is still connected to the QFabric system.

```
root@node1> request chassis device-mode standalone
Device mode set to 'standalone' mode.
Please reboot the system to complete the process.
```



NOTE: Always convert the device mode to standalone before you remove the component from the QFabric system. If you remove the component from the QFabric system before converting the device mode to standalone, the switch might not operate properly. For example, the output of the **show chassis hardware** command might display no FPCs or interfaces for the switch.

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@node1> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Standalone
```

6. Reboot the component to complete the conversion process.

```
root@node1> request system reboot
Reboot the system ? [yes,no] (no) yes
```

```
Shutdown NOW!
[pid 34992]
```

```
root@node1>
```

```
*** FINAL System shutdown message from root@node1 ***
System going down IMMEDIATELY
```

7. Disconnect and remove the component from the QFabric system. You may now operate the device as a standalone switch.

Related Documentation

- [request chassis device-mode on page 1441](#)
- [show chassis device-mode on page 1475](#)

- [Software Installation Overview on page 120](#)
- *Connecting a QFX3500 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network*
- *Connecting a QFX3600 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network*
- *Connecting a QFX3500 Node Device to a QFX3008-I Interconnect Device*
- *Connecting a QFX3600 Node Device to a QFX3008-I Interconnect Device*
- *Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network*
- *Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network*
- *Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network*
- *Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network*
- *Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device*
- *Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device*

Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane

This example shows you how to connect QFabric system components and configure the Virtual Chassis used by the QFX3000-G QFabric system control plane network. Proper wiring of Director devices, Interconnect devices, and Node devices to the Virtual Chassis, combined with a standard configuration, enables you to bring up the internal QFabric system management network and prepare your QFabric system for full operation.

- [Requirements on page 1263](#)
- [Overview on page 1264](#)
- [Configuration on page 1275](#)
- [Verification on page 1291](#)

Requirements

This example uses the following hardware and software components:

- One QFX3000-G QFabric system containing:
 - Two QFX3100 Director devices
 - Two QFX3008-I Interconnect devices
 - Eight QFX3500 Node devices
- Eight EX4200-48T switches, used to make two redundant Virtual Chassis with four members apiece

- Junos OS Release 12.1R1.10 for the EX Series switches used in the Virtual Chassis
- Junos OS Release 12.2 for the QFX Series

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see *Installing and Connecting a QFX3100 Director Device*, *Installing and Connecting a QFX3008-I Interconnect Device*, and *Installing and Connecting a QFX3500 Device*.
- Rack, mount, and install your Virtual Chassis hardware (EX4200 switches). For more information, see *Installing and Connecting an EX4200 Switch*.
- Create two Virtual Chassis of four members each. For more information, see *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.

Overview


The QFX3000-G QFabric system control plane network connects the Director group, Interconnect devices, and Node devices in a QFabric system across a pair of redundant Virtual Chassis. By separating the management control plane from the data plane, the QFabric system can scale efficiently. The control plane network uses Gigabit Ethernet cabling and connections between components, and a 10-Gigabit Ethernet backbone between the redundant Virtual Chassis.

Specific ports have been reserved on the Virtual Chassis to connect to each of the QFabric system device types. Such design simplifies installation and facilitates timely deployment of a QFabric system. It also permits the use of a standard Virtual Chassis configuration included as part of this example. The standard configuration can scale from the minimum topology of eight Node devices shown in this example to the maximum of 128 Node devices for a fully implemented QFX3000-G QFabric system.

Topology

[Figure 26 on page 1265](#) shows the general port ranges where QFabric system devices must be connected to the Virtual Chassis. For each Virtual Chassis member, connect ports 0 through 31 to Node devices, ports 38 and 39 to Interconnect devices, and ports 40 and 41 to Director devices. [Table 126 on page 1265](#) shows the details of the QFabric system device-to-Virtual Chassis port mappings.

Figure 26: QFX3000-G QFabric System Control Plane—Virtual Chassis Port Ranges



g041122



CAUTION:

- The control plane network within a QFabric system should be considered a critical component of the system that should not be shared with other network traffic. In order to scale efficiently, the control plane network must be reserved for the QFabric system and its components. As a result, the ports of the QFabric system control plane must never be used for any purpose other than to transport QFabric system control plane traffic, and we neither recommend nor support the connection of other devices to the QFabric system control plane network.
- Do not install Junos Space and AI-Scripts (AIS) on the control plane network Virtual Chassis in a QFX3000-G QFabric system.



NOTE: Not all port numbers are represented in [Table 126 on page 1265](#), and ports 32 through 37 and ports 42 through 47 are reserved for future uses.

[Table 126 on page 1265](#) shows the specific mappings of QFabric system control plane network ports from the Virtual Chassis to the QFabric system components.

Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments

Member 0	Member 1	Member 2	Member 3	Member Port Number	QFabric System Component
Node0	Node32	Node64	Node96	ge-X/0/0	Node devices
ge-0/0/0	ge-1/0/0	ge-2/0/0	ge-3/0/0		

Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments (*continued*)

Member 0	Member 1	Member 2	Member 3	Member Port Number	QFabric System Component
Node1 ge-0/0/1	Node33 ge-1/0/1	Node65 ge-2/0/1	Node97 ge-3/0/1	ge-X/0/1	Node devices
Node2 ge-0/0/2	Node34 ge-1/0/2	Node66 ge-2/0/2	Node98 ge-3/0/2	ge-X/0/2	Node devices
Node3 ge-0/0/3	Node35 ge-1/0/3	Node67 ge-2/0/3	Node99 ge-3/0/3	ge-X/0/3	Node devices
Node4 ge-0/0/4	Node36 ge-1/0/4	Node68 ge-2/0/4	Node100 ge-3/0/4	ge-X/0/4	Node devices
Node5 ge-0/0/5	Node37 ge-1/0/5	Node69 ge-2/0/5	Node101 ge-3/0/5	ge-X/0/5	Node devices
Node6 ge-0/0/6	Node38 ge-1/0/6	Node70 ge-2/0/6	Node102 ge-3/0/6	ge-X/0/6	Node devices
Node7 ge-0/0/7	Node39 ge-1/0/7	Node71 ge-2/0/7	Node103 ge-3/0/7	ge-X/0/7	Node devices
Node8 ge-0/0/8	Node40 ge-1/0/8	Node72 ge-2/0/8	Node104 ge-3/0/8	ge-X/0/8	Node devices
Node9 ge-0/0/9	Node41 ge-1/0/9	Node73 ge-2/0/9	Node105 ge-3/0/9	ge-X/0/9	Node devices
Node10 ge-0/0/10	Node42 ge-1/0/10	Node74 ge-2/0/10	Node106 ge-3/0/10	ge-X/0/10	Node devices
Node11 ge-0/0/11	Node43 ge-1/0/11	Node75 ge-2/0/11	Node107 ge-3/0/11	ge-X/0/11	Node devices
Node12 ge-0/0/12	Node44 ge-1/0/12	Node76 ge-2/0/12	Node108 ge-3/0/12	ge-X/0/12	Node devices
Node13 ge-0/0/13	Node45 ge-1/0/13	Node77 ge-2/0/13	Node109 ge-3/0/13	ge-X/0/13	Node devices

Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments (*continued*)

Member 0	Member 1	Member 2	Member 3	Member Port Number	QFabric System Component
Node14 ge-0/0/14	Node46 ge-1/0/14	Node78 ge-2/0/14	Node110 ge-3/0/14	ge-X/0/14	Node devices
Node15 ge-0/0/15	Node47 ge-1/0/15	Node79 ge-2/0/15	Node111 ge-3/0/15	ge-X/0/15	Node devices
Node16 ge-0/0/16	Node48 ge-1/0/16	Node80 ge-2/0/16	Node112 ge-3/0/16	ge-X/0/16	Node devices
Node17 ge-0/0/17	Node49 ge-1/0/17	Node81 ge-2/0/17	Node113 ge-3/0/17	ge-X/0/17	Node devices
Node18 ge-0/0/18	Node50 ge-1/0/18	Node82 ge-2/0/18	Node114 ge-3/0/18	ge-X/0/18	Node devices
Node19 ge-0/0/19	Node51 ge-1/0/19	Node83 ge-2/0/19	Node115 ge-3/0/19	ge-X/0/19	Node devices
Node20 ge-0/0/20	Node52 ge-1/0/20	Node84 ge-2/0/20	Node116 ge-3/0/20	ge-X/0/20	Node devices
Node21 ge-0/0/21	Node53 ge-1/0/21	Node85 ge-2/0/21	Node117 ge-3/0/21	ge-X/0/21	Node devices
Node22 ge-0/0/22	Node54 ge-1/0/22	Node86 ge-2/0/22	Node118 ge-3/0/22	ge-X/0/22	Node devices
Node23 ge-0/0/23	Node55 ge-1/0/23	Node87 ge-2/0/23	Node119 ge-3/0/23	ge-X/0/23	Node devices
Node24 ge-0/0/24	Node56 ge-1/0/24	Node88 ge-2/0/24	Node120 ge-3/0/24	ge-X/0/24	Node devices
Node25 ge-0/0/25	Node57 ge-1/0/25	Node89 ge-2/0/25	Node121 ge-3/0/25	ge-X/0/25	Node devices
Node26 ge-0/0/26	Node58 ge-1/0/26	Node90 ge-2/0/26	Node122 ge-3/0/26	ge-X/0/26	Node devices

Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments (*continued*)

Member 0	Member 1	Member 2	Member 3	Member Port Number	QFabric System Component
Node27 ge-0/0/27	Node59 ge-1/0/27	Node91 ge-2/0/27	Node123 ge-3/0/27	ge-X/0/27	Node devices
Node28 ge-0/0/28	Node60 ge-1/0/28	Node92 ge-2/0/28	Node124 ge-3/0/28	ge-X/0/28	Node devices
Node29 ge-0/0/29	Node61 ge-1/0/29	Node93 ge-2/0/29	Node125 ge-3/0/29	ge-X/0/29	Node devices
Node30 ge-0/0/30	Node62 ge-1/0/30	Node94 ge-2/0/30	Node126 ge-3/0/30	ge-X/0/30	Node devices
Node31 ge-0/0/31	Node63 ge-1/0/31	Node95 ge-2/0/31	Node127 ge-3/0/31	ge-X/0/31	Node devices
Reserved ge-0/0/32	Reserved ge-1/0/32	Reserved ge-2/0/32	Reserved ge-3/0/32	ge-X/0/32	Future use
...
Reserved ge-0/0/37	Reserved ge-1/0/37	Reserved ge-2/0/37	Reserved ge-3/0/37	ge-X/0/37	Future use
IC2 CB0 ge-0/0/38	IC2 CB1 ge-1/0/38	IC3 CB0 ge-2/0/38	IC3 CB1 ge-3/0/38	ge-X/0/38	Interconnect devices NOTE: On both Control Boards, use port 0 to connect to VC0, and port 1 to connect to VC1.
IC0 CB0 ge-0/0/39	IC0 CB1 ge-1/0/39	IC1 CB0 ge-2/0/39	IC1 CB1 ge-3/0/39	ge-X/0/39	Interconnect devices NOTE: On both Control Boards, use port 0 to connect to VC0, and port 1 to connect to VC1.

Table 126: QFX3000-G QFabric System Virtual Chassis Control Plane Port Assignments (*continued*)

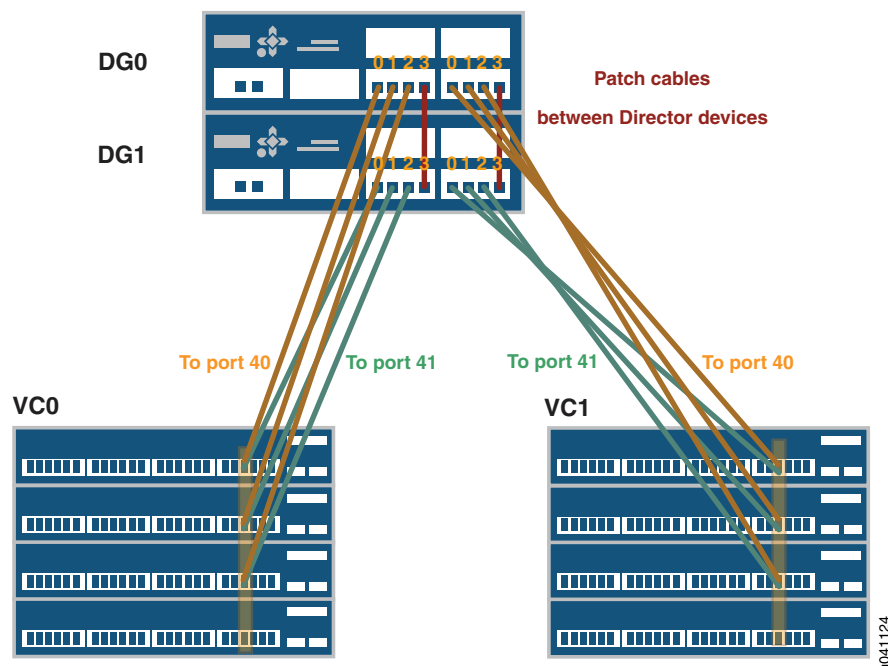
Member 0	Member 1	Member 2	Member 3	Member Port Number	QFabric System Component
DG0 port 0 ge-0/0/40	DG0 port 1 ge-1/0/40	DG0 port 2 ge-2/0/40	Reserved ge-3/0/40	ge-X/0/40	Director device 0
DG1 port 0 ge-0/0/41	DG1 port 1 ge-1/0/41	DG1 port 2 ge-2/0/41	Reserved ge-3/0/41	ge-X/0/41	Director device 1
Reserved ge-0/0/42	Reserved ge-1/0/42	Reserved ge-2/0/42	Reserved ge-3/0/42	ge-X/0/42	Future use
...
Reserved ge-0/0/47	Reserved ge-1/0/47	Reserved ge-2/0/47	Reserved ge-3/0/47	ge-X/0/47	Future use
Inter-VC xe-0/1/0	Inter-VC xe-1/1/0	Inter-VC xe-2/1/0	Inter-VC xe-3/1/0	Inter-VC xe-X/1/0	Inter-Virtual Chassis LAG
Inter-VC xe-0/1/2	Inter-VC xe-1/1/2	Inter-VC xe-2/1/2	Inter-VC xe-3/1/2	Inter-VC xe-X/1/2	Inter-Virtual Chassis LAG

Next, connect the Director devices to the Virtual Chassis. In general, you want to accomplish the following:

- Connect three ports from one network module in a Director device to the first Virtual Chassis, and three ports from the second network module to the second Virtual Chassis. You need to repeat these connections from the second Director device to both Virtual Chassis to provide resiliency for the system.
- Connect the Director devices to each other and create a Director group. You can use either straight-through RJ-45 patch cables or crossover cables, because the Director devices contain autosensing modules. Connect one port from each network module on the first Director device to one port in each network module on the second Director device.

Figure 27 on page 1270 shows the specific ports on the Director group that you must connect to the Virtual Chassis and interconnect between the Director devices.

Figure 27: QFX3000-G QFabric System Control Plane—Director Group to Virtual Chassis Connections



In this specific example, connect ports 0, 1, and 2 from module 0 on Director device DG0 to port 40 on Virtual Chassis VC0 (ge-0/0/40, ge-1/0/40, and ge-2/0/40), and connect ports 0, 1, and 2 from module 1 to port 40 on Virtual Chassis VC1 (ge-0/0/40, ge-1/0/40, and ge-2/0/40).

For Director device DG1, connect ports 0, 1, and 2 from module 0 to port 41 on Virtual Chassis VC0 (ge-0/0/41, ge-1/0/41, and ge-2/0/41), and connect ports 0, 1, and 2 from module 1 to port 41 on Virtual Chassis VC1 (ge-0/0/41, ge-1/0/41, and ge-2/0/41).

To form the Director group, connect module 0, port 3 on Director device DG0 to module 0, port 3 on Director device DG1. Similarly, connect module 1, port 3 on Director device DG0 to module 1, port 3 on Director device DG1. [Table 127 on page 1270](#) shows the port mappings for the Director group in this example.

Table 127: Director Group Port Mappings

Director Device	Virtual Chassis VC0	Virtual Chassis VC1
DG0	<ul style="list-style-type: none"> Module 0, port 0 to ge-0/0/40 on VC0 Module 0, port 1 to ge-1/0/40 on VC0 Module 0, port 2 to ge-2/0/40 on VC0 Module 0, port 3 to module 0, port 3 on DG1 	<ul style="list-style-type: none"> Module 1, port 0 to ge-0/0/40 on VC1 Module 1, port 1 to ge-1/0/40 on VC1 Module 1, port 2 to ge-2/0/40 on VC1 Module 1, port 3 to module 1, port 3 on DG1

Table 127: Director Group Port Mappings (*continued*)

Director Device	Virtual Chassis VC0	Virtual Chassis VC1
DG1	<ul style="list-style-type: none"> Module 0, port 0 to ge-0/0/41 on VC0 Module 0, port 1 to ge-1/0/41 on VC0 Module 0, port 2 to ge-2/0/41 on VC0 Module 0, port 3 to module 0, port 3 on DG0 	<ul style="list-style-type: none"> Module 1, port 0 to ge-0/0/41 on VC1 Module 1, port 1 to ge-1/0/41 on VC1 Module 1, port 2 to ge-2/0/41 on VC1 Module 1, port 3 to module 1, port 3 on DG0


In the software, the ports of each network module are reversed, numbered from right to left, and incremented sequentially across modules. If you issue interface operational commands directly on the Director device, please note the following port mappings as shown in [Table 128 on page 1271](#):

Table 128: Hardware to Software Port Mappings for Director Device Network Modules

Network Module	Port 0	Port 1	Port 2	Port 3
Module 0	eth5	eth4	eth3	eth2
Module 1	eth9	eth8	eth7	eth6

[Figure 28 on page 1272](#) shows the specific ports on the Interconnect devices that you must connect to the Virtual Chassis. In general, connect one port from each Control Board module in an Interconnect device to the first Virtual Chassis, and a second port from each Control Board module to the second Virtual Chassis.

Figure 28: QFX3000-G QFabric System Control Plane—Interconnect Device to Virtual Chassis Connections



In this specific example, for both Interconnect devices IC0 and IC1, connect port 0 from CB0 and CB1 to Virtual Chassis VC0 and port 1 from CB0 and CB1 to Virtual Chassis VC1. Connect the port 0 cables to port 39 on Virtual Chassis VC0 (ge-0/0/39, ge-1/0/39, ge-2/0/39, and ge-3/0/39), and connect the port 1 cables to port 39 on Virtual Chassis VC1 (ge-0/0/39, ge-1/0/39, ge-2/0/39, and ge-3/0/39). [Table 129 on page 1272](#) shows the port mappings for the Interconnect devices in this example.

Table 129: Interconnect Device Port Mappings

Interconnect Device	Virtual Chassis VC0	Virtual Chassis VC1
IC0	<ul style="list-style-type: none"> • CB0, port 0 to ge-0/0/39 • CB1, port 0 to ge-1/0/39 	<ul style="list-style-type: none"> • CB0, port 1 to ge-0/0/39 • CB1, port 1 to ge-1/0/39
IC1	<ul style="list-style-type: none"> • CB0, port 0 to ge-2/0/39 • CB1, port 0 to ge-3/0/39 	<ul style="list-style-type: none"> • CB0, port 1 to ge-2/0/39 • CB1, port 1 to ge-3/0/39

As required, you can extend the number of Interconnect devices from two to four. For additional Interconnect devices IC2 and IC3, connect port 0 from CB0 and CB1 to Virtual


Chassis VC0 and port 1 from CB0 and CB1 to Virtual Chassis VC1. Connect the port 0 cables to port 38 on Virtual Chassis VC0 (ge-0/0/38, ge-1/0/38, ge-2/0/38, and ge-3/0/38), and connect the port 1 cables to port 38 on Virtual Chassis VC1 (ge-0/0/38, ge-1/0/38, ge-2/0/38, and ge-3/0/38). [Table 130 on page 1273](#) shows the port mappings needed to extend the number of Interconnect devices in this example to four devices.

Table 130: Interconnect Device Port Mappings for Two Additional Devices

Interconnect Device	Virtual Chassis VC0	Virtual Chassis VC1
IC2	<ul style="list-style-type: none"> CB0, port 0 to ge-0/0/38 CB1, port 0 to ge-1/0/38 	<ul style="list-style-type: none"> CB0, port 1 to ge-0/0/38 CB1, port 1 to ge-1/0/38
IC3	<ul style="list-style-type: none"> CB0, port 0 to ge-2/0/38 CB1, port 0 to ge-3/0/38 	<ul style="list-style-type: none"> CB0, port 1 to ge-2/0/38 CB1, port 1 to ge-3/0/38

[Figure 29 on page 1273](#) shows the specific ports on the Node devices that you must connect to the Virtual Chassis. In general, connect the first management port from a Node device to the first Virtual Chassis, and the second management port to the second Virtual Chassis.

Figure 29: QFX3000-G QFabric System Control Plane—Node Device to Virtual Chassis Connections



In this specific example, for Node device Node0, connect port C0 (also known as me0) to Virtual Chassis VC0 port ge-0/0/0, and connect port C1 (also known as me1) to Virtual Chassis VC1 port ge-0/0/0.


For the remaining seven Node devices, connect port C0 to the ge-0/0/X port on Virtual Chassis VC0 that matches the Node device number. Similarly, connect port C1 to the port on Virtual Chassis VC1 that matches the Node device number. For example, you would connect Node device Node5 to port ge-0/0/5. [Table 131 on page 1274](#) shows the full set of port mappings for the Node devices in this example.

Table 131: Node Device Port Mappings

Node Device	Virtual Chassis VC0	Virtual Chassis VC1
Node0	C0 to ge-0/0/0	C1 to ge-0/0/0
Node1	C0 to ge-0/0/1	C1 to ge-0/0/1
Node2	C0 to ge-0/0/2	C1 to ge-0/0/2
Node3	C0 to ge-0/0/3	C1 to ge-0/0/3
Node4	C0 to ge-0/0/4	C1 to ge-0/0/4
Node5	C0 to ge-0/0/5	C1 to ge-0/0/5
Node6	C0 to ge-0/0/6	C1 to ge-0/0/6
Node7	C0 to ge-0/0/7	C1 to ge-0/0/7

Figure 30 on page 1274 shows the specific ports on the members of the first Virtual Chassis that you must connect to the members of the second Virtual Chassis. These connections create a link aggregation bundle (LAG) that provides redundancy and resiliency for the Virtual Chassis portion of the control plane. In general, connect each 10-Gigabit Ethernet uplink port from the first Virtual Chassis to the corresponding 10-Gigabit Ethernet uplink port on the second Virtual Chassis.

Figure 30: QFX3000-G QFabric System Control Plane—Inter-Virtual Chassis LAG Connections



In this specific example, for Virtual Chassis VC0, connect port xe-0/1/0 to Virtual Chassis VC1 port xe-0/1/0. For the remaining seven 10-Gigabit Ethernet uplink ports, connect each port from VC0 to the corresponding port on VC1. For example, you would connect the xe-2/1/2 port on VC0 to port xe-2/1/2 on VC1, and so on.

Table 132 on page 1275 shows the full set of port mappings for the Virtual Chassis LAG connections in this example.

Table 132: Virtual Chassis LAG Port Mappings

VCO and VC1	Member 0	Member 1	Member 2	Member 3
Uplink port 0	xe-0/1/0 to xe-0/1/0	xe-1/1/0 to xe-1/1/0	xe-2/1/0 to xe-2/1/0	xe-3/1/0 to xe-3/1/0
Uplink port 2	xe-0/1/2 to xe-0/1/2	xe-1/1/2 to xe-1/1/2	xe-2/1/2 to xe-2/1/2	xe-3/1/2 to xe-3/1/2

Configuration

CLI Quick Configuration

To quickly configure the QFabric system control plane Virtual Chassis, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.



NOTE: The configuration files for a QFabric system control plane network are also available for download from the QFX Series section of the Junos OS software download page at <https://www.juniper.net/support/downloads/junos.html>.

```

set groups qfabric system commit synchronize
set groups qfabric chassis redundancy graceful-switchover
set groups qfabric chassis aggregated-devices ethernet device-count 10
set groups qfabric chassis fpc 0 pic 1 sfppplus pic-mode 10g
set groups qfabric chassis fpc 1 pic 1 sfppplus pic-mode 10g
set groups qfabric chassis fpc 2 pic 1 sfppplus pic-mode 10g
set groups qfabric chassis fpc 3 pic 1 sfppplus pic-mode 10g
set groups qfabric chassis lcd-menu fpc 0 menu-item maintenance-menu disable
set groups qfabric chassis lcd-menu fpc 1 menu-item maintenance-menu disable
set groups qfabric chassis lcd-menu fpc 2 menu-item maintenance-menu disable
set groups qfabric chassis lcd-menu fpc 3 menu-item maintenance-menu disable
set groups qfabric chassis alarm management-ethernet link-down ignore
set groups qfabric protocols rstp interface ae8.0 mode point-to-point
set groups qfabric protocols rstp interface all edge
set groups qfabric protocols rstp interface all no-root-port
set groups qfabric protocols rstp bpdu-block-on-edge
set groups qfabric protocols lldp interface all
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 110
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 111
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_2 loss-priority low code-points 100
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_2 loss-priority high code-points 101
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_0 loss-priority low code-points 010
set groups qfabric class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_0 loss-priority high code-points 001
set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 110

```

```

set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
  forwarding-class class_3 loss-priority low code-points 111
set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
  forwarding-class class_2 loss-priority low code-points 100
set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
  forwarding-class class_2 loss-priority high code-points 101
set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
  forwarding-class class_0 loss-priority low code-points 010
set groups qfabric class-of-service classifiers inet-precedence IP_qfabric_classifier
  forwarding-class class_0 loss-priority high code-points 001
set groups qfabric class-of-service forwarding-classes class class_3 queue-num 7
set groups qfabric class-of-service forwarding-classes class class_2 queue-num 2
set groups qfabric class-of-service forwarding-classes class class_0 queue-num 0
set groups qfabric class-of-service interfaces ge-*/0/* scheduler-map cpe_network_smap
set groups qfabric class-of-service interfaces ge-*/0/* unit 0 classifiers ieee-802.1
  onep_qfabric_classifier
set groups qfabric class-of-service interfaces ge-*/0/* unit 0 classifiers inet-precedence
  IP_qfabric_classifier
set groups qfabric class-of-service interfaces ae* scheduler-map cpe_network_smap
set groups qfabric class-of-service interfaces ae* unit 0 classifiers ieee-802.1
  onep_qfabric_classifier
set groups qfabric class-of-service interfaces ae* unit 0 classifiers inet-precedence
  IP_qfabric_classifier
set groups qfabric class-of-service scheduler-maps cpe_network_smap forwarding-class
  class_3 scheduler scheduler_3
set groups qfabric class-of-service scheduler-maps cpe_network_smap forwarding-class
  class_2 scheduler scheduler_2
set groups qfabric class-of-service scheduler-maps cpe_network_smap forwarding-class
  class_0 scheduler scheduler_0
set groups qfabric class-of-service schedulers scheduler_3 buffer-size percent 30
set groups qfabric class-of-service schedulers scheduler_3 priority strict-high
set groups qfabric class-of-service schedulers scheduler_2 transmit-rate percent 75
set groups qfabric class-of-service schedulers scheduler_2 buffer-size percent 30
set groups qfabric class-of-service schedulers scheduler_2 priority low
set groups qfabric class-of-service schedulers scheduler_0 transmit-rate percent 25
set groups qfabric class-of-service schedulers scheduler_0 buffer-size percent 40
set groups qfabric class-of-service schedulers scheduler_0 priority low
set groups qfabric ethernet-switching-options nonstop-bridging
set groups qfabric ethernet-switching-options storm-control interface all bandwidth
  10000
set groups qfabric vlans qfabric vlan-id 100
set groups qfabric vlans qfabric dot1q-tunneling
set groups qfabric-int interfaces <*> mtu 9216
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching port-mode access
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching vlan members
  qfabric
set groups qfabric-ae interfaces <*> aggregated-ether-options link-speed 1g
set groups qfabric-ae interfaces <*> aggregated-ether-options lacp active
set apply-groups qfabric
set interfaces interface-range Node_Device_Interfaces member "ge-[0-3]/0/[0-31]"
set interfaces interface-range Node_Device_Interfaces apply-groups qfabric-int
set interfaces interface-range Node_Device_Interfaces description "QFabric Node Device"
set interfaces interface-range Interconnect_Device_Interfaces member
  "ge-[0-3]/0/[38-39]"
set interfaces interface-range Interconnect_Device_Interfaces apply-groups qfabric-int

```



```

set interfaces interface-range Interconnect_Device_Interfaces description "QFabric
Interconnect Device"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces member
"ge-[0-3]/0/40"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces description "QFabric
Director Device - DG0"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options 802.3ad
ae0
set interfaces interface-range Director_Device_DG1_LAG_Interfaces member
"ge-[0-3]/0/41"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces description "QFabric
Director Device - DG1"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options 802.3ad
ae1
set interfaces interface-range Director_Device_DG2_LAG_Interfaces member
"ge-[0-3]/0/42"
set interfaces interface-range Director_Device_DG2_LAG_Interfaces description "QFabric
Director Device - DG2"
set interfaces interface-range Director_Device_DG2_LAG_Interfaces ether-options 802.3ad
ae2
set interfaces interface-range Director_Device_DG3_LAG_Interfaces member
"ge-[0-3]/0/43"
set interfaces interface-range Director_Device_DG3_LAG_Interfaces description "QFabric
Director Device - DG3"
set interfaces interface-range Director_Device_DG3_LAG_Interfaces ether-options 802.3ad
ae3
set interfaces interface-range Director_Device_DG4_LAG_Interfaces member
"ge-[0-3]/0/44"
set interfaces interface-range Director_Device_DG4_LAG_Interfaces description "QFabric
Director Device - DG4"
set interfaces interface-range Director_Device_DG4_LAG_Interfaces ether-options 802.3ad
ae4
set interfaces interface-range Director_Device_DG5_LAG_Interfaces member
"ge-[0-3]/0/45"
set interfaces interface-range Director_Device_DG5_LAG_Interfaces description "QFabric
Director Device - DG5"
set interfaces interface-range Director_Device_DG5_LAG_Interfaces ether-options 802.3ad
ae5
set interfaces interface-range Director_Device_DG6_LAG_Interfaces member
"ge-[0-3]/0/46"
set interfaces interface-range Director_Device_DG6_LAG_Interfaces description "QFabric
Director Device - DG6"
set interfaces interface-range Director_Device_DG6_LAG_Interfaces ether-options 802.3ad
ae6
set interfaces interface-range Director_Device_DG7_LAG_Interfaces member
"ge-[0-3]/0/47"
set interfaces interface-range Director_Device_DG7_LAG_Interfaces description "QFabric
Director Device - DG7"
set interfaces interface-range Director_Device_DG7_LAG_Interfaces ether-options 802.3ad
ae7
set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces member
"xe-[0-3]/1/0"
set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces member
"xe-[0-3]/1/2"
set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces description "QFabric
Control Plane (Inter-VC LAG)"

```

```
set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces ether-options
  802.3ad ae8
set interfaces ae0 apply-groups qfabric-int
set interfaces ae0 apply-groups qfabric-ae
set interfaces ae0 description "QFabric Director Device - DG0"
set interfaces ae1 apply-groups qfabric-int
set interfaces ae1 apply-groups qfabric-ae
set interfaces ae1 description "QFabric Director Device - DG1"
set interfaces ae2 apply-groups qfabric-int
set interfaces ae2 apply-groups qfabric-ae
set interfaces ae2 description "QFabric Director Device - DG2"
set interfaces ae3 apply-groups qfabric-int
set interfaces ae3 apply-groups qfabric-ae
set interfaces ae3 description "QFabric Director Device - DG3"
set interfaces ae4 apply-groups qfabric-int
set interfaces ae4 apply-groups qfabric-ae
set interfaces ae4 description "QFabric Director Device - DG4"
set interfaces ae5 apply-groups qfabric-int
set interfaces ae5 apply-groups qfabric-ae
set interfaces ae5 description "QFabric Director Device - DG5"
set interfaces ae6 apply-groups qfabric-int
set interfaces ae6 apply-groups qfabric-ae
set interfaces ae6 description "QFabric Director Device - DG6"
set interfaces ae7 apply-groups qfabric-int
set interfaces ae7 apply-groups qfabric-ae
set interfaces ae7 description "QFabric Director Device - DG7"
set interfaces ae8 description "QFabric Control Plane (Inter-VC LAG)"
set interfaces ae8 mtu 9216
set interfaces ae8 aggregated-ether-options link-speed 10g
set interfaces ae8 aggregated-ether-options lacp active
set interfaces ae8 unit 0 family ethernet-switching vlan members qfabric
set system host-name qfabric-control-plane
set system services ssh
set system services telnet
set system services web-management http
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file messages archive world-readable
set system syslog file messages explicit-priority
set system syslog file interactive-commands interactive-commands any
set system syslog file secure authorization info
set system syslog file default-log-messages any any
set system syslog file default-log-messages structured-data
set system syslog file console any error
set system syslog time-format millisecond
set interfaces vme unit 0 family inet address 192.168.157.26/24
set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1
set virtual-chassis preprovisioned
set virtual-chassis member 0 role routing-engine
set virtual-chassis member 0 serial-number abc123
set virtual-chassis member 1 role routing-engine
set virtual-chassis member 1 serial-number def456
set virtual-chassis member 2 role line-card
set virtual-chassis member 2 serial-number ghi789
set virtual-chassis member 3 role line-card
```

```
set virtual-chassis member 3 serial-number jkl012
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a Virtual Chassis for the QFabric system control plane network:

1. Create a configuration group to define global QFabric system control plane properties. Enable commit synchronization and graceful switchover, set up the number of aggregated Ethernet devices, configure alarm and LCD management, activate loop prevention, nonstop bridging, and storm control, configure Link Layer Discovery Protocol (LLDP), specify a global VLAN (VLAN 100) and 802.1q tunneling, define options for aggregated Ethernet interfaces, and enable the uplink module for 10-Gigabit Ethernet operation.

Enable class of service (CoS) for the QFabric system control plane network. Establish forwarding classes, priorities, scheduler maps, classifiers, and queues for three types of traffic: control traffic, interdevice traffic, and best-effort traffic. Apply the qfabric group settings to the configuration.

[edit]

```
user@switch# set groups qfabric system commit synchronize
user@switch# set groups qfabric chassis redundancy graceful-switchover
user@switch# set groups qfabric chassis aggregated-devices ethernet device-count
10
user@switch# set groups qfabric chassis fpc 0 pic 1 sfppplus pic-mode 10g
user@switch# set groups qfabric chassis fpc 1 pic 1 sfppplus pic-mode 10g
user@switch# set groups qfabric chassis fpc 2 pic 1 sfppplus pic-mode 10g
user@switch# set groups qfabric chassis fpc 3 pic 1 sfppplus pic-mode 10g
user@switch# set groups qfabric chassis lcd-menu fpc 0 menu-item
maintenance-menu disable
user@switch# set groups qfabric chassis lcd-menu fpc 1 menu-item
maintenance-menu disable
user@switch# set groups qfabric chassis lcd-menu fpc 2 menu-item
maintenance-menu disable
user@switch# set groups qfabric chassis lcd-menu fpc 3 menu-item
maintenance-menu disable
user@switch# set groups qfabric chassis alarm management-ethernet link-down
ignore
user@switch# set groups qfabric protocols rstp interface ae8.0 mode point-to-point
user@switch# set groups qfabric protocols rstp interface all edge
user@switch# set groups qfabric protocols rstp interface all no-root-port
user@switch# set groups qfabric protocols rstp bpdu-block-on-edge
user@switch# set groups qfabric protocols lldp interface all
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
onep_qfabric_classifier forwarding-class class_3 loss-priority low code-points 110
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
onep_qfabric_classifier forwarding-class class_3 loss-priority low code-points 111
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
onep_qfabric_classifier forwarding-class class_2 loss-priority low code-points 100
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
onep_qfabric_classifier forwarding-class class_2 loss-priority high code-points
101
```

```
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
  onep_qfabric_classifier forwarding-class class_0 loss-priority low code-points
  010
user@switch# set groups qfabric class-of-service classifiers ieee-802.1
  onep_qfabric_classifier forwarding-class class_0 loss-priority high code-points
  001
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_3 loss-priority low code-points 110
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_3 loss-priority low code-points 111
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_2 loss-priority low code-points 100
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_2 loss-priority high code-points 101
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_0 loss-priority low code-points 010
user@switch# set groups qfabric class-of-service classifiers inet-precedence
  IP_qfabric_classifier forwarding-class class_0 loss-priority high code-points 001
user@switch# set groups qfabric class-of-service forwarding-classes class class_3
  queue-num 7
user@switch# set groups qfabric class-of-service forwarding-classes class class_2
  queue-num 2
user@switch# set groups qfabric class-of-service forwarding-classes class class_0
  queue-num 0
user@switch# set groups qfabric class-of-service interfaces ge-*/0/* scheduler-map
  cpe_network_smap
user@switch# set groups qfabric class-of-service interfaces ge-*/0/* unit 0
  classifiers ieee-802.1 onep_qfabric_classifier
user@switch# set groups qfabric class-of-service interfaces ge-*/0/* unit 0
  classifiers inet-precedence IP_qfabric_classifier
user@switch# set groups qfabric class-of-service interfaces ae* scheduler-map
  cpe_network_smap
user@switch# set groups qfabric class-of-service interfaces ae* unit 0 classifiers
  ieee-802.1 onep_qfabric_classifier
user@switch# set groups qfabric class-of-service interfaces ae* unit 0 classifiers
  inet-precedence IP_qfabric_classifier
user@switch# set groups qfabric class-of-service scheduler-maps
  cpe_network_smap forwarding-class class_3 scheduler scheduler_3
user@switch# set groups qfabric class-of-service scheduler-maps
  cpe_network_smap forwarding-class class_2 scheduler scheduler_2
user@switch# set groups qfabric class-of-service scheduler-maps
  cpe_network_smap forwarding-class class_0 scheduler scheduler_0
user@switch# set groups qfabric class-of-service schedulers scheduler_3 buffer-size
  percent 30
user@switch# set groups qfabric class-of-service schedulers scheduler_3 priority
  strict-high
user@switch# set groups qfabric class-of-service schedulers scheduler_2
  transmit-rate percent 75
user@switch# set groups qfabric class-of-service schedulers scheduler_2 buffer-size
  percent 30
user@switch# set groups qfabric class-of-service schedulers scheduler_2 priority
  low
user@switch# set groups qfabric class-of-service schedulers scheduler_0
  transmit-rate percent 25
user@switch# set groups qfabric class-of-service schedulers scheduler_0 buffer-size
  percent 40
```

```

user@switch# set groups qfabric class-of-service schedulers scheduler_0 priority
low
user@switch# set groups qfabric ethernet-switching-options nonstop-bridging
user@switch# set groups qfabric ethernet-switching-options storm-control interface
all bandwidth 10000
user@switch# set groups qfabric vlans qfabric vlan-id 100
user@switch# set groups qfabric vlans qfabric dot1q-tunneling
user@switch# set groups qfabric-int interfaces <*> mtu 9216
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching
port-mode access
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching
vlan members qfabric
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options
link-speed 1g
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options lacp
active
user@switch# set apply-groups qfabric

```

2. Configure interfaces for the QFabric system control plane network. Set the interface ranges where Node devices (0 through 31), Interconnect devices (38 and 39), and Director devices (40 and 41) connect to the control plane network through the Virtual Chassis. Configure the inter-Virtual Chassis LAG connections for the ae8 interface and apply the ae-interfaces configuration group to the remaining aggregated Ethernet interfaces (ae0 through ae7).

```

[edit]
user@switch# set interfaces interface-range Node_Device_Interfaces member
"ge-[0-3]/0/[0-31]"
user@switch# set interfaces interface-range Node_Device_Interfaces apply-groups
qfabric-int
user@switch# set interfaces interface-range Node_Device_Interfaces description
"QFabric Node Device"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces member
"ge-[0-3]/0/[38-39]"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces
apply-groups qfabric-int
user@switch# set interfaces interface-range Interconnect_Device_Interfaces
description "QFabric Interconnect Device"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
member "ge-[0-3]/0/40"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
description "QFabric Director Device - DG0"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
ether-options 802.3ad ae0
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
member "ge-[0-3]/0/41"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
description "QFabric Director Device - DG1"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
ether-options 802.3ad ae1
user@switch# set interfaces interface-range Director_Device_DG2_LAG_Interfaces
member "ge-[0-3]/0/42"
user@switch# set interfaces interface-range Director_Device_DG2_LAG_Interfaces
description "QFabric Director Device - DG2"
user@switch# set interfaces interface-range Director_Device_DG2_LAG_Interfaces
ether-options 802.3ad ae2

```

```
user@switch# set interfaces interface-range Director_Device_DG3_LAG_Interfaces
member "ge-[0-3]/0/43"
user@switch# set interfaces interface-range Director_Device_DG3_LAG_Interfaces
description "QFabric Director Device - DG3"
user@switch# set interfaces interface-range Director_Device_DG3_LAG_Interfaces
ether-options 802.3ad ae3
user@switch# set interfaces interface-range Director_Device_DG4_LAG_Interfaces
member "ge-[0-3]/0/44"
user@switch# set interfaces interface-range Director_Device_DG4_LAG_Interfaces
description "QFabric Director Device - DG4"
user@switch# set interfaces interface-range Director_Device_DG4_LAG_Interfaces
ether-options 802.3ad ae4
user@switch# set interfaces interface-range Director_Device_DG5_LAG_Interfaces
member "ge-[0-3]/0/45"
user@switch# set interfaces interface-range Director_Device_DG5_LAG_Interfaces
description "QFabric Director Device - DG5"
user@switch# set interfaces interface-range Director_Device_DG5_LAG_Interfaces
ether-options 802.3ad ae5
user@switch# set interfaces interface-range Director_Device_DG6_LAG_Interfaces
member "ge-[0-3]/0/46"
user@switch# set interfaces interface-range Director_Device_DG6_LAG_Interfaces
description "QFabric Director Device - DG6"
user@switch# set interfaces interface-range Director_Device_DG6_LAG_Interfaces
ether-options 802.3ad ae6
user@switch# set interfaces interface-range Director_Device_DG7_LAG_Interfaces
member "ge-[0-3]/0/47"
user@switch# set interfaces interface-range Director_Device_DG7_LAG_Interfaces
description "QFabric Director Device - DG7"
user@switch# set interfaces interface-range Director_Device_DG7_LAG_Interfaces
ether-options 802.3ad ae7
user@switch# set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces
member "xe-[0-3]/1/0"
user@switch# set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces
member "xe-[0-3]/1/2"
user@switch# set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces
description "QFabric Control Plane (Inter-VC LAG)"
user@switch# set interfaces interface-range Control_Plane_Inter_VC_LAG_Interfaces
ether-options 802.3ad ae8
user@switch# set interfaces ae0 apply-groups qfabric-int
user@switch# set interfaces ae0 apply-groups qfabric-ae
user@switch# set interfaces ae0 description "QFabric Director Device - DG0"
user@switch# set interfaces ae1 apply-groups qfabric-int
user@switch# set interfaces ae1 apply-groups qfabric-ae
user@switch# set interfaces ae1 description "QFabric Director Device - DG1"
user@switch# set interfaces ae2 apply-groups qfabric-int
user@switch# set interfaces ae2 apply-groups qfabric-ae
user@switch# set interfaces ae2 description "QFabric Director Device - DG2"
user@switch# set interfaces ae3 apply-groups qfabric-int
user@switch# set interfaces ae3 apply-groups qfabric-ae
user@switch# set interfaces ae3 description "QFabric Director Device - DG3"
user@switch# set interfaces ae4 apply-groups qfabric-int
user@switch# set interfaces ae4 apply-groups qfabric-ae
user@switch# set interfaces ae4 description "QFabric Director Device - DG4"
user@switch# set interfaces ae5 apply-groups qfabric-int
user@switch# set interfaces ae5 apply-groups qfabric-ae
user@switch# set interfaces ae5 description "QFabric Director Device - DG5"
```

```

user@switch# set interfaces ae6 apply-groups qfabric-int
user@switch# set interfaces ae6 apply-groups qfabric-ae
user@switch# set interfaces ae6 description "QFabric Director Device - DG6"
user@switch# set interfaces ae7 apply-groups qfabric-int
user@switch# set interfaces ae7 apply-groups qfabric-ae
user@switch# set interfaces ae7 description "QFabric Director Device - DG7"
user@switch# set interfaces ae8 description "QFabric Control Plane (Inter-VC LAG)"
user@switch# set interfaces ae8 mtu 9216
user@switch# set interfaces ae8 aggregated-ether-options link-speed 10g
user@switch# set interfaces ae8 aggregated-ether-options lacp active
user@switch# set interfaces ae8 unit 0 family ethernet-switching vlan members
qfabric

```

3. Configure settings to enable the Virtual Chassis to interoperate with your management network. Set a hostname, system services (such as Telnet), system log thresholds, management interface parameters, default routes, Virtual Chassis preprovisioning, and any additional preferences you might have.

```

[edit]
user@switch# set system host-name qfabric-control-plane
user@switch# set system services ssh
user@switch# set system services telnet
user@switch# set system services web-management http
user@switch# set system syslog user * any emergency
user@switch# set system syslog file messages any notice
user@switch# set system syslog file messages authorization info
user@switch# set system syslog file messages archive world-readable
user@switch# set system syslog file messages explicit-priority
user@switch# set system syslog file interactive-commands interactive-commands
any
user@switch# set system syslog file secure authorization info
user@switch# set system syslog file default-log-messages any any
user@switch# set system syslog file default-log-messages structured-data
user@switch# set system syslog file console any error
user@switch# set system syslog time-format millisecond
user@switch# set interfaces vme unit 0 family inet address 192.168.157.26/24
user@switch# set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1
user@switch# set virtual-chassis preprovisioned
user@switch# set virtual-chassis member 0 role routing-engine
user@switch# set virtual-chassis member 0 serial-number abc123
user@switch# set virtual-chassis member 1 role routing-engine
user@switch# set virtual-chassis member 1 serial-number def456
user@switch# set virtual-chassis member 2 role line-card
user@switch# set virtual-chassis member 2 serial-number ghi789
user@switch# set virtual-chassis member 3 role line-card
user@switch# set virtual-chassis member 3 serial-number jkl012

```

Results To view the configuration, issue the **show** command in configuration mode or the **show configuration** command in operational mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

The following configuration is the standard configuration that applies universally to both Virtual Chassis in your QFabric system control plane network.

```
[edit]
```

```
groups {
  qfabric {
    system {
      commit {
        synchronize;
      }
    }
  }
  chassis {
    redundancy {
      graceful-switchover;
    }
    aggregated-devices {
      ethernet {
        device-count 10;
      }
    }
    alarm {
      management-ethernet {
        link-down ignore;
      }
    }
  }
  fpc 0 {
    pic 1 {
      sfppplus {
        pic-mode 10g;
      }
    }
  }
  fpc 1 {
    pic 1 {
      sfppplus {
        pic-mode 10g;
      }
    }
  }
  fpc 2 {
    pic 1 {
      sfppplus {
        pic-mode 10g;
      }
    }
  }
  fpc 3 {
    pic 1 {
      sfppplus {
        pic-mode 10g;
      }
    }
  }
  lcd-menu {
    fpc 0 {
      menu-item {
        maintenance-menu disable;
      }
    }
    fpc 1 {
```



```

        menu-item {
            maintenance-menu disable;
        }
    }
    fpc 2 {
        menu-item {
            maintenance-menu disable;
        }
    }
    fpc 3 {
        menu-item {
            maintenance-menu disable;
        }
    }
}
}
protocols {
    rstp {
        interface ae8.0 {
            mode point-to-point;
        }
        interface all {
            edge;
            no-root-port;
        }
        bpdu-block-on-edge;
    }
    lldp {
        interface all;
    }
}
class-of-service {
    classifiers {
        ieee-802.1 onep_qfabric_classifier {
            forwarding-class class_3 {
                loss-priority low code-points [ 110 111 ];
            }
            forwarding-class class_2 {
                loss-priority low code-points 100;
                loss-priority high code-points 101;
            }
            forwarding-class class_0 {
                loss-priority low code-points 010;
                loss-priority high code-points 001;
            }
        }
        inet-precedence IP_qfabric_classifier {
            forwarding-class class_3 {
                loss-priority low code-points [ 110 111 ];
            }
            forwarding-class class_2 {
                loss-priority low code-points 100;
                loss-priority high code-points 101;
            }
            forwarding-class class_0 {
                loss-priority low code-points 010;
            }
        }
    }
}

```

```
        loss-priority high code-points 001;
    }
}
}
forwarding-classes {
    class class_3 queue-num 7;
    class class_2 queue-num 2;
    class class_0 queue-num 0;
}
interfaces {
    ge-*/0/* {
        scheduler-map cpe_network_smap;
        unit 0 {
            classifiers {
                ieee-802.1 onep_qfabric_classifier;
                inet-precedence IP_qfabric_classifier;
            }
        }
    }
}
ae* {
    scheduler-map cpe_network_smap;
    unit 0 {
        classifiers {
            ieee-802.1 onep_qfabric_classifier;
            inet-precedence IP_qfabric_classifier;
        }
    }
}
}
scheduler-maps {
    cpe_network_smap {
        forwarding-class class_3 scheduler scheduler_3;
        forwarding-class class_2 scheduler scheduler_2;
        forwarding-class class_0 scheduler scheduler_0;
    }
}
schedulers {
    scheduler_3 {
        buffer-size percent 30;
        priority strict-high;
    }
    scheduler_2 {
        transmit-rate percent 75;
        buffer-size percent 30;
        priority low;
    }
    scheduler_0 {
        transmit-rate percent 25;
        buffer-size percent 40;
        priority low;
    }
}
}
ethernet-switching-options {
    nonstop-bridging;
    storm-control {
```

```

        interface all {
            bandwidth 10000;
        }
    }
}
vllans {
    qfabric {
        vlan-id 100;
        dot1q-tunneling;
    }
}
}
qfabric-int {
    interfaces {
        <*> {
            mtu 9216;
            unit 0 {
                family ethernet-switching {
                    port-mode access;
                    vlan {
                        members qfabric;
                    }
                }
            }
        }
    }
}
}
qfabric-ae {
    interfaces {
        <*> {
            aggregated-ether-options {
                link-speed 1g;
                lACP {
                    active;
                }
            }
        }
    }
}
}
}
apply-groups [qfabric];
interfaces {
    interface-range Node_Device_Interfaces {
        member "ge-[0-3]/0/[0-31]";
        description "QFabric Node Device";
        apply-groups qfabric-int;
    }
    interface-range Interconnect_Device_Interfaces {
        member "ge-[0-3]/0/[38-39]";
        description "QFabric Interconnect Device";
        apply-groups qfabric-int;
    }
    interface-range Director_Device_DG0_LAG_Interfaces {
        member "ge-[0-3]/0/40";
        description "QFabric Director Device - DG0";
        ether-options {

```

```
        802.3ad ae0;
    }
}
interface-range Director_Device_DG1_LAG_Interfaces {
    member "ge-[0-3]/0/41";
    description "QFabric Director Device - DG1";
    ether-options {
        802.3ad ae1;
    }
}
interface-range Director_Device_DG2_LAG_Interfaces {
    member "ge-[0-3]/0/42";
    description "QFabric Director Device - DG2";
    ether-options {
        802.3ad ae2;
    }
}
interface-range Director_Device_DG3_LAG_Interfaces {
    member "ge-[0-3]/0/43";
    description "QFabric Director Device - DG3";
    ether-options {
        802.3ad ae3;
    }
}
interface-range Director_Device_DG4_LAG_Interfaces {
    member "ge-[0-3]/0/44";
    description "QFabric Director Device - DG4";
    ether-options {
        802.3ad ae4;
    }
}
interface-range Director_Device_DG5_LAG_Interfaces {
    member "ge-[0-3]/0/45";
    description "QFabric Director Device - DG5";
    ether-options {
        802.3ad ae5;
    }
}
interface-range Director_Device_DG6_LAG_Interfaces {
    member "ge-[0-3]/0/46";
    description "QFabric Director Device - DG6";
    ether-options {
        802.3ad ae6;
    }
}
interface-range Director_Device_DG7_LAG_Interfaces {
    member "ge-[0-3]/0/47";
    description "QFabric Director Device - DG7";
    ether-options {
        802.3ad ae7;
    }
}
interface-range Control_Plane_Inter_VC_LAG_Interfaces {
    member "xe-[0-3]/1/0";
    member "xe-[0-3]/1/2";
    description "QFabric Control Plane (Inter-VC LAG)";
```

```

        ether-options {
            802.3ad ae8;
        }
    }
    ae0 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG0";
    }
    ae1 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG1";
    }
    ae2 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG2";
    }
    ae3 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG3";
    }
    ae4 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG4";
    }
    ae5 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG5";
    }
    ae6 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG6";
    }
    ae7 {
        apply-groups [ qfabric-int qfabric-ae ];
        description "QFabric Director Device - DG7";
    }
    ae8 {
        description "QFabric Control Plane (Inter-VC LAG)";
        mtu 9216;
        aggregated-ether-options {
            link-speed 10g;
            lacp {
                active;
            }
        }
        unit 0 {
            family ethernet-switching {
                vlan {
                    members qfabric;
                }
            }
        }
    }
}

```

The following portion of the configuration applies to the specific requirements of your management network. Modify this section to meet the needs of your network.

```
[edit]
system {
  host-name qfabric-control-plane;
  services {
    ssh;
    telnet;
    web-management {
      http;
    }
  }
  syslog {
    user * {
      any emergency;
    }
    file messages {
      any notice;
      authorization info;
      archive world-readable;
      explicit-priority;
    }
    file interactive-commands {
      interactive-commands any;
    }
    file secure {
      authorization info;
    }
    file default-log-messages {
      any any;
      structured-data;
    }
    file console {
      any error;
    }
    time-format millisecond;
  }
}
interfaces {
  vme {
    unit 0 {
      family inet {
        address 192.168.157.26/24;
      }
    }
  }
}
routing-options {
  static {
    route 0.0.0.0/0 next-hop 192.168.157.1;
  }
}
virtual-chassis {
  preprovisioned;
  member 0 {
```

```

    role routing-engine;
    serial-number abc123;
  }
  member 1 {
    role routing-engine;
    serial-number def456;
  }
  member 2 {
    role line-card;
    serial-number ghi789;
  }
  member 3 {
    role line-card;
    serial-number jkl012;
  }
}

```

To verify the syntax of your configuration before committing it, enter **commit check** from configuration mode. If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the Virtual Chassis configuration is working properly.

- [Verifying the QFabric System Control Plane—Virtual Chassis VC0 on page 1291](#)
- [Verifying the QFabric System Control Plane—Virtual Chassis VC1 on page 1299](#)

Verifying the QFabric System Control Plane—Virtual Chassis VC0

Purpose Verify that your first Virtual Chassis is operational.

Action Connect to the Junos OS CLI of Virtual Chassis VC0, either from your management network or from the console port of the master Virtual Chassis member. In operational mode, enter the **show virtual-chassis status** and **show interfaces terse** commands.

Sample Output

```

{master:0}
user@vc0> show virtual-chassis status

Virtual Chassis ID: c809.2c5d.9f7b

```

Member ID	Status	Serial No	Model	Mastership priority	Role	Neighbor List ID	Interface
0 (FPC 0)	Prsnt	BP0210471476	ex4200-48t	128	Master*	1	vcp-1
1 (FPC 1)	Prsnt	BP0210460181	ex4200-48t	128	Backup	0	vcp-0
2 (FPC 2)	Prsnt	BP0210458724	ex4200-48t	128	Linecard	2	vcp-1
3 (FPC 3)	Prsnt	BP0210477189	ex4200-48t	128	Linecard	3	vcp-1
						2	vcp-0

```

Member ID for next new member: 4 (FPC 4)

{master:0}
user@vc0> show interfaces terse

```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	eth-switch		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	eth-switch		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch		
ge-0/0/7	up	up			
ge-0/0/7.0	up	up	eth-switch		
ge-0/0/8	up	down			
ge-0/0/8.0	up	down	eth-switch		
ge-0/0/9	up	down			
ge-0/0/9.0	up	down	eth-switch		
ge-0/0/10	up	down			
ge-0/0/10.0	up	down	eth-switch		
ge-0/0/11	up	down			
ge-0/0/11.0	up	down	eth-switch		
ge-0/0/12	up	down			
ge-0/0/12.0	up	down	eth-switch		
ge-0/0/13	up	down			
ge-0/0/13.0	up	down	eth-switch		
ge-0/0/14	up	down			
ge-0/0/14.0	up	down	eth-switch		
ge-0/0/15	up	down			
ge-0/0/15.0	up	down	eth-switch		
ge-0/0/16	up	down			
ge-0/0/16.0	up	down	eth-switch		
ge-0/0/17	up	down			
ge-0/0/17.0	up	down	eth-switch		
ge-0/0/18	up	down			
ge-0/0/18.0	up	down	eth-switch		
ge-0/0/19	up	down			
ge-0/0/19.0	up	down	eth-switch		
ge-0/0/20	up	down			
ge-0/0/20.0	up	down	eth-switch		
ge-0/0/21	up	down			
ge-0/0/21.0	up	down	eth-switch		
ge-0/0/22	up	down			
ge-0/0/22.0	up	down	eth-switch		
ge-0/0/23	up	down			
ge-0/0/23.0	up	down	eth-switch		
ge-0/0/24	up	down			
ge-0/0/24.0	up	down	eth-switch		
ge-0/0/25	up	down			
ge-0/0/25.0	up	down	eth-switch		
ge-0/0/26	up	down			
ge-0/0/26.0	up	down	eth-switch		
ge-0/0/27	up	down			
ge-0/0/27.0	up	down	eth-switch		
ge-0/0/28	up	down			
ge-0/0/28.0	up	down	eth-switch		
ge-0/0/29	up	down			
ge-0/0/29.0	up	down	eth-switch		


```

ge-0/0/30          up    down
ge-0/0/30.0        up    down eth-switch
ge-0/0/31          up    down
ge-0/0/31.0        up    down eth-switch
ge-0/0/32          up    down
ge-0/0/33          up    down
ge-0/0/34          up    down
ge-0/0/35          up    down
ge-0/0/36          up    down
ge-0/0/36.0        up    down eth-switch
ge-0/0/37          up    down
ge-0/0/37.0        up    down eth-switch
ge-0/0/38          up    down
ge-0/0/38.0        up    down eth-switch
ge-0/0/39          up    up
ge-0/0/39.0        up    up eth-switch
ge-0/0/40          up    up
ge-0/0/40.0        up    up aenet --> ae0.0
ge-0/0/41          up    down
ge-0/0/41.0        up    down aenet --> ae1.0
ge-0/0/42          up    down
ge-0/0/42.0        up    down aenet --> ae2.0
ge-0/0/43          up    down
ge-0/0/43.0        up    down aenet --> ae3.0
ge-0/0/44          up    down
ge-0/0/44.0        up    down aenet --> ae4.0
ge-0/0/45          up    down
ge-0/0/45.0        up    down aenet --> ae5.0
ge-0/0/46          up    down
ge-0/0/46.0        up    down aenet --> ae6.0
ge-0/0/47          up    down
ge-0/0/47.0        up    down aenet --> ae7.0
xe-0/1/0          up    up
xe-0/1/0.0        up    up aenet --> ae8.0
xe-0/1/1          up    down
xe-0/1/2          up    up
xe-0/1/2.0        up    up aenet --> ae8.0
xe-0/1/3          up    down
ge-1/0/0          up    down
ge-1/0/0.0        up    down eth-switch
ge-1/0/1          up    down
ge-1/0/1.0        up    down eth-switch
ge-1/0/2          up    down
ge-1/0/2.0        up    down eth-switch
ge-1/0/3          up    down
ge-1/0/3.0        up    down eth-switch
ge-1/0/4          up    down
ge-1/0/4.0        up    down eth-switch
ge-1/0/5          up    down
ge-1/0/5.0        up    down eth-switch
ge-1/0/6          up    down
ge-1/0/6.0        up    down eth-switch
ge-1/0/7          up    down
ge-1/0/7.0        up    down eth-switch
ge-1/0/8          up    down
ge-1/0/8.0        up    down eth-switch
ge-1/0/9          up    down
ge-1/0/9.0        up    down eth-switch
ge-1/0/10         up    down
ge-1/0/10.0       up    down eth-switch
ge-1/0/11         up    down

```

ge-1/0/11.0	up	down	eth-switch
ge-1/0/12	up	down	
ge-1/0/12.0	up	down	eth-switch
ge-1/0/13	up	down	
ge-1/0/13.0	up	down	eth-switch
ge-1/0/14	up	down	
ge-1/0/14.0	up	down	eth-switch
ge-1/0/15	up	down	
ge-1/0/15.0	up	down	eth-switch
ge-1/0/16	up	down	
ge-1/0/16.0	up	down	eth-switch
ge-1/0/17	up	down	
ge-1/0/17.0	up	down	eth-switch
ge-1/0/18	up	down	
ge-1/0/18.0	up	down	eth-switch
ge-1/0/19	up	down	
ge-1/0/19.0	up	down	eth-switch
ge-1/0/20	up	down	
ge-1/0/20.0	up	down	eth-switch
ge-1/0/21	up	down	
ge-1/0/21.0	up	down	eth-switch
ge-1/0/22	up	down	
ge-1/0/22.0	up	down	eth-switch
ge-1/0/23	up	down	
ge-1/0/23.0	up	down	eth-switch
ge-1/0/24	up	down	
ge-1/0/24.0	up	down	eth-switch
ge-1/0/25	up	down	
ge-1/0/25.0	up	down	eth-switch
ge-1/0/26	up	down	
ge-1/0/26.0	up	down	eth-switch
ge-1/0/27	up	down	
ge-1/0/27.0	up	down	eth-switch
ge-1/0/28	up	down	
ge-1/0/28.0	up	down	eth-switch
ge-1/0/29	up	down	
ge-1/0/29.0	up	down	eth-switch
ge-1/0/30	up	down	
ge-1/0/30.0	up	down	eth-switch
ge-1/0/31	up	down	
ge-1/0/31.0	up	down	eth-switch
ge-1/0/32	up	down	
ge-1/0/33	up	down	
ge-1/0/34	up	down	
ge-1/0/35	up	down	
ge-1/0/36	up	down	
ge-1/0/36.0	up	down	eth-switch
ge-1/0/37	up	down	
ge-1/0/37.0	up	down	eth-switch
ge-1/0/38	up	down	
ge-1/0/38.0	up	down	eth-switch
ge-1/0/39	up	up	
ge-1/0/39.0	up	up	eth-switch
ge-1/0/40	up	up	
ge-1/0/40.0	up	up	aenet --> ae0.0
ge-1/0/41	up	down	
ge-1/0/41.0	up	down	aenet --> ae1.0
ge-1/0/42	up	down	
ge-1/0/42.0	up	down	aenet --> ae2.0
ge-1/0/43	up	down	
ge-1/0/43.0	up	down	aenet --> ae3.0

ge-1/0/44	up	down	
ge-1/0/44.0	up	down	aenet --> ae4.0
ge-1/0/45	up	down	
ge-1/0/45.0	up	down	aenet --> ae5.0
ge-1/0/46	up	down	
ge-1/0/46.0	up	down	aenet --> ae6.0
ge-1/0/47	up	down	
ge-1/0/47.0	up	down	aenet --> ae7.0
xe-1/1/0	up	up	
xe-1/1/0.0	up	up	aenet --> ae8.0
xe-1/1/1	up	down	
xe-1/1/2	up	up	
xe-1/1/2.0	up	up	aenet --> ae8.0
xe-1/1/3	up	down	
ge-2/0/0	up	down	
ge-2/0/0.0	up	down	eth-switch
ge-2/0/1	up	down	
ge-2/0/1.0	up	down	eth-switch
ge-2/0/2	up	down	
ge-2/0/2.0	up	down	eth-switch
ge-2/0/3	up	down	
ge-2/0/3.0	up	down	eth-switch
ge-2/0/4	up	down	
ge-2/0/4.0	up	down	eth-switch
ge-2/0/5	up	down	
ge-2/0/5.0	up	down	eth-switch
ge-2/0/6	up	down	
ge-2/0/6.0	up	down	eth-switch
ge-2/0/7	up	down	
ge-2/0/7.0	up	down	eth-switch
ge-2/0/8	up	down	
ge-2/0/8.0	up	down	eth-switch
ge-2/0/9	up	down	
ge-2/0/9.0	up	down	eth-switch
ge-2/0/10	up	down	
ge-2/0/10.0	up	down	eth-switch
ge-2/0/11	up	down	
ge-2/0/11.0	up	down	eth-switch
ge-2/0/12	up	down	
ge-2/0/12.0	up	down	eth-switch
ge-2/0/13	up	down	
ge-2/0/13.0	up	down	eth-switch
ge-2/0/14	up	down	
ge-2/0/14.0	up	down	eth-switch
ge-2/0/15	up	down	
ge-2/0/15.0	up	down	eth-switch
ge-2/0/16	up	down	
ge-2/0/16.0	up	down	eth-switch
ge-2/0/17	up	down	
ge-2/0/17.0	up	down	eth-switch
ge-2/0/18	up	down	
ge-2/0/18.0	up	down	eth-switch
ge-2/0/19	up	down	
ge-2/0/19.0	up	down	eth-switch
ge-2/0/20	up	down	
ge-2/0/20.0	up	down	eth-switch
ge-2/0/21	up	down	
ge-2/0/21.0	up	down	eth-switch
ge-2/0/22	up	down	
ge-2/0/22.0	up	down	eth-switch
ge-2/0/23	up	down	

ge-2/0/23.0	up	down	eth-switch
ge-2/0/24	up	down	
ge-2/0/24.0	up	down	eth-switch
ge-2/0/25	up	down	
ge-2/0/25.0	up	down	eth-switch
ge-2/0/26	up	down	
ge-2/0/26.0	up	down	eth-switch
ge-2/0/27	up	down	
ge-2/0/27.0	up	down	eth-switch
ge-2/0/28	up	down	
ge-2/0/28.0	up	down	eth-switch
ge-2/0/29	up	down	
ge-2/0/29.0	up	down	eth-switch
ge-2/0/30	up	down	
ge-2/0/30.0	up	down	eth-switch
ge-2/0/31	up	down	
ge-2/0/31.0	up	down	eth-switch
ge-2/0/32	up	down	
ge-2/0/33	up	down	
ge-2/0/34	up	down	
ge-2/0/35	up	down	
ge-2/0/36	up	down	
ge-2/0/36.0	up	down	eth-switch
ge-2/0/37	up	down	
ge-2/0/37.0	up	down	eth-switch
ge-2/0/38	up	down	
ge-2/0/38.0	up	down	eth-switch
ge-2/0/39	up	up	
ge-2/0/39.0	up	up	eth-switch
ge-2/0/40	up	up	
ge-2/0/40.0	up	up	aenet --> ae0.0
ge-2/0/41	up	down	
ge-2/0/41.0	up	down	aenet --> ae1.0
ge-2/0/42	up	down	
ge-2/0/42.0	up	down	aenet --> ae2.0
ge-2/0/43	up	down	
ge-2/0/43.0	up	down	aenet --> ae3.0
ge-2/0/44	up	down	
ge-2/0/44.0	up	down	aenet --> ae4.0
ge-2/0/45	up	down	
ge-2/0/45.0	up	down	aenet --> ae5.0
ge-2/0/46	up	down	
ge-2/0/46.0	up	down	aenet --> ae6.0
ge-2/0/47	up	down	
ge-2/0/47.0	up	down	aenet --> ae7.0
xe-2/1/0	up	up	
xe-2/1/0.0	up	up	aenet --> ae8.0
xe-2/1/1	up	down	
xe-2/1/2	up	up	
xe-2/1/2.0	up	up	aenet --> ae8.0
xe-2/1/3	up	down	
ge-3/0/0	up	down	
ge-3/0/0.0	up	down	eth-switch
ge-3/0/1	up	down	
ge-3/0/1.0	up	down	eth-switch
ge-3/0/2	up	down	
ge-3/0/2.0	up	down	eth-switch
ge-3/0/3	up	down	
ge-3/0/3.0	up	down	eth-switch
ge-3/0/4	up	down	
ge-3/0/4.0	up	down	eth-switch

ge-3/0/5	up	down
ge-3/0/5.0	up	down eth-switch
ge-3/0/6	up	down
ge-3/0/6.0	up	down eth-switch
ge-3/0/7	up	down
ge-3/0/7.0	up	down eth-switch
ge-3/0/8	up	down
ge-3/0/8.0	up	down eth-switch
ge-3/0/9	up	down
ge-3/0/9.0	up	down eth-switch
ge-3/0/10	up	down
ge-3/0/10.0	up	down eth-switch
ge-3/0/11	up	down
ge-3/0/11.0	up	down eth-switch
ge-3/0/12	up	down
ge-3/0/12.0	up	down eth-switch
ge-3/0/13	up	down
ge-3/0/13.0	up	down eth-switch
ge-3/0/14	up	down
ge-3/0/14.0	up	down eth-switch
ge-3/0/15	up	down
ge-3/0/15.0	up	down eth-switch
ge-3/0/16	up	down
ge-3/0/16.0	up	down eth-switch
ge-3/0/17	up	down
ge-3/0/17.0	up	down eth-switch
ge-3/0/18	up	down
ge-3/0/18.0	up	down eth-switch
ge-3/0/19	up	down
ge-3/0/19.0	up	down eth-switch
ge-3/0/20	up	down
ge-3/0/20.0	up	down eth-switch
ge-3/0/21	up	down
ge-3/0/21.0	up	down eth-switch
ge-3/0/22	up	down
ge-3/0/22.0	up	down eth-switch
ge-3/0/23	up	down
ge-3/0/23.0	up	down eth-switch
ge-3/0/24	up	down
ge-3/0/24.0	up	down eth-switch
ge-3/0/25	up	down
ge-3/0/25.0	up	down eth-switch
ge-3/0/26	up	down
ge-3/0/26.0	up	down eth-switch
ge-3/0/27	up	down
ge-3/0/27.0	up	down eth-switch
ge-3/0/28	up	down
ge-3/0/28.0	up	down eth-switch
ge-3/0/29	up	down
ge-3/0/29.0	up	down eth-switch
ge-3/0/30	up	down
ge-3/0/30.0	up	down eth-switch
ge-3/0/31	up	down
ge-3/0/31.0	up	down eth-switch
ge-3/0/32	up	down
ge-3/0/33	up	down
ge-3/0/34	up	down
ge-3/0/35	up	down
ge-3/0/36	up	down
ge-3/0/36.0	up	down eth-switch
ge-3/0/37	up	down

ge-3/0/37.0	up	down	eth-switch
ge-3/0/38	up	down	
ge-3/0/38.0	up	down	eth-switch
ge-3/0/39	up	up	
ge-3/0/39.0	up	up	eth-switch
ge-3/0/40	up	down	
ge-3/0/40.0	up	down	aenet --> ae0.0
ge-3/0/41	up	down	
ge-3/0/41.0	up	down	aenet --> ae1.0
ge-3/0/42	up	down	
ge-3/0/42.0	up	down	aenet --> ae2.0
ge-3/0/43	up	down	
ge-3/0/43.0	up	down	aenet --> ae3.0
ge-3/0/44	up	down	
ge-3/0/44.0	up	down	aenet --> ae4.0
ge-3/0/45	up	down	
ge-3/0/45.0	up	down	aenet --> ae5.0
ge-3/0/46	up	down	
ge-3/0/46.0	up	down	aenet --> ae6.0
ge-3/0/47	up	down	
ge-3/0/47.0	up	down	aenet --> ae7.0
xe-3/1/0	up	up	
xe-3/1/0.0	up	up	aenet --> ae8.0
xe-3/1/1	up	down	
xe-3/1/2	up	up	
xe-3/1/2.0	up	up	aenet --> ae8.0
xe-3/1/3	up	down	
vcp-0	up	down	
vcp-0.32768	up	down	
vcp-1	up	up	
vcp-1.32768	up	up	
ae0	up	up	
ae0.0	up	up	eth-switch
ae1	up	down	
ae1.0	up	down	eth-switch
ae2	up	down	
ae2.0	up	down	eth-switch
ae3	up	down	
ae3.0	up	down	eth-switch
ae4	up	down	
ae4.0	up	down	eth-switch
ae5	up	down	
ae5.0	up	down	eth-switch
ae6	up	down	
ae6.0	up	down	eth-switch
ae7	up	down	
ae7.0	up	down	eth-switch
ae8	up	up	
ae8.0	up	up	eth-switch
ae9	up	down	
bme0	up	up	
bme0.32768	up	up	inet 128.0.0.1/2 128.0.0.16/2 128.0.0.32/2
			tnp 0x10
bme0.32770	up	up	eth-switch
bme0.32771	down	up	eth-switch
bme0.32772	down	up	eth-switch
dsc	up	up	
gre	up	up	
ipip	up	up	

```

jsrv                up    up
jsrv.1              up    up    inet    128.0.0.127/2
lo0                  up    up
lo0.0                up    up    inet    10.255.195.96    --> 0/0
                                iso
47.0005.80ff.f800.0000.0108.0001.0102.5519.5096
                                inet6    abcd::10:255:195:96
                                fe80::2ac0:da0f:fc31:1e80

lsi                  up    up
me0                  up    up
me0.0                up    up    inet    10.94.195.96/24
mtun                  up    up
pimd                  up    up
pime                  up    up
tap                  up    up
vlan                  up    up
vme                  up    up    inet    192.168.157.26/24

```

Meaning In the output of the **show virtual-chassis status** command, if all four members appear, the Virtual Chassis is operational.

In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as up (such as ge-0/0/39, ge-1/0/39, ge-2/0/39, and ge-3/0/39 for the Interconnect devices; ge-0/0/40, ge-1/0/40, and ge-2/0/40 for the Director devices; ge-0/0/0 through ge-0/0/7 for the Node devices; and xe-0/1/0, xe-0/1/2, xe-1/1/0, xe-1/1/2, xe-2/1/0, xe-2/1/2, xe-3/1/0, and xe-3/1/2 for the inter-Virtual Chassis connections), the control plane is properly connected.

Verifying the QFabric System Control Plane—Virtual Chassis VC1

Purpose Verify that your second Virtual Chassis is operational.

Action Connect to the Junos OS CLI of Virtual Chassis VC1, either from your management network or from the console port of the master Virtual Chassis member. In operational mode, enter the **show virtual-chassis status** and **show interfaces terse** commands.

Sample Output

```

{master:0}
user@vc1> show virtual-chassis status

Virtual Chassis ID: c809.2c5d.9f8a

Member ID  Status  Serial No  Model  Mastership  Role  Neighbor List
0 (FPC 0)  Prsnt    BP0210471477  ex4200-48t  128  Master*  1 vcp-1
1 (FPC 1)  Prsnt    BP0210460182  ex4200-48t  128  Backup   0 vcp-0
2 (FPC 2)  Prsnt    BP0210458725  ex4200-48t  128  Linecard 1 vcp-0
3 (FPC 3)  Prsnt    BP0210477180  ex4200-48t  128  Linecard 2 vcp-0
2 vcp-1
3 vcp-1

Member ID for next new member: 4 (FPC 4)

{master:0}
user@vc1> show interfaces terse

Interface          Admin Link Proto  Local  Remote
ge-0/0/0            up    up

```

ge-0/0/0.0	up	up	eth-switch
ge-0/0/1	up	up	
ge-0/0/1.0	up	up	eth-switch
ge-0/0/2	up	up	
ge-0/0/2.0	up	up	eth-switch
ge-0/0/3	up	up	
ge-0/0/3.0	up	up	eth-switch
ge-0/0/4	up	up	
ge-0/0/4.0	up	up	eth-switch
ge-0/0/5	up	up	
ge-0/0/5.0	up	up	eth-switch
ge-0/0/6	up	up	
ge-0/0/6.0	up	up	eth-switch
ge-0/0/7	up	up	
ge-0/0/7.0	up	up	eth-switch
ge-0/0/8	up	down	
ge-0/0/8.0	up	down	eth-switch
ge-0/0/9	up	down	
ge-0/0/9.0	up	down	eth-switch
ge-0/0/10	up	down	
ge-0/0/10.0	up	down	eth-switch
ge-0/0/11	up	down	
ge-0/0/11.0	up	down	eth-switch
ge-0/0/12	up	down	
ge-0/0/12.0	up	down	eth-switch
ge-0/0/13	up	down	
ge-0/0/13.0	up	down	eth-switch
ge-0/0/14	up	down	
ge-0/0/14.0	up	down	eth-switch
ge-0/0/15	up	down	
ge-0/0/15.0	up	down	eth-switch
ge-0/0/16	up	down	
ge-0/0/16.0	up	down	eth-switch
ge-0/0/17	up	down	
ge-0/0/17.0	up	down	eth-switch
ge-0/0/18	up	down	
ge-0/0/18.0	up	down	eth-switch
ge-0/0/19	up	down	
ge-0/0/19.0	up	down	eth-switch
ge-0/0/20	up	down	
ge-0/0/20.0	up	down	eth-switch
ge-0/0/21	up	down	
ge-0/0/21.0	up	down	eth-switch
ge-0/0/22	up	down	
ge-0/0/22.0	up	down	eth-switch
ge-0/0/23	up	down	
ge-0/0/23.0	up	down	eth-switch
ge-0/0/24	up	down	
ge-0/0/24.0	up	down	eth-switch
ge-0/0/25	up	down	
ge-0/0/25.0	up	down	eth-switch
ge-0/0/26	up	down	
ge-0/0/26.0	up	down	eth-switch
ge-0/0/27	up	down	
ge-0/0/27.0	up	down	eth-switch
ge-0/0/28	up	down	
ge-0/0/28.0	up	down	eth-switch
ge-0/0/29	up	down	
ge-0/0/29.0	up	down	eth-switch
ge-0/0/30	up	down	
ge-0/0/30.0	up	down	eth-switch

ge-0/0/31	up	down	
ge-0/0/31.0	up	down	eth-switch
ge-0/0/32	up	down	
ge-0/0/33	up	down	
ge-0/0/34	up	down	
ge-0/0/35	up	down	
ge-0/0/36	up	down	
ge-0/0/36.0	up	down	eth-switch
ge-0/0/37	up	down	
ge-0/0/37.0	up	down	eth-switch
ge-0/0/38	up	down	
ge-0/0/38.0	up	down	eth-switch
ge-0/0/39	up	up	
ge-0/0/39.0	up	up	eth-switch
ge-0/0/40	up	up	
ge-0/0/40.0	up	up	aenet --> ae0.0
ge-0/0/41	up	down	
ge-0/0/41.0	up	down	aenet --> ae1.0
ge-0/0/42	up	down	
ge-0/0/42.0	up	down	aenet --> ae2.0
ge-0/0/43	up	down	
ge-0/0/43.0	up	down	aenet --> ae3.0
ge-0/0/44	up	down	
ge-0/0/44.0	up	down	aenet --> ae4.0
ge-0/0/45	up	down	
ge-0/0/45.0	up	down	aenet --> ae5.0
ge-0/0/46	up	down	
ge-0/0/46.0	up	down	aenet --> ae6.0
ge-0/0/47	up	down	
ge-0/0/47.0	up	down	aenet --> ae7.0
xe-0/1/0	up	up	
xe-0/1/0.0	up	up	aenet --> ae8.0
xe-0/1/1	up	down	
xe-0/1/2	up	up	
xe-0/1/2.0	up	up	aenet --> ae8.0
xe-0/1/3	up	down	
ge-1/0/0	up	down	
ge-1/0/0.0	up	down	eth-switch
ge-1/0/1	up	down	
ge-1/0/1.0	up	down	eth-switch
ge-1/0/2	up	down	
ge-1/0/2.0	up	down	eth-switch
ge-1/0/3	up	down	
ge-1/0/3.0	up	down	eth-switch
ge-1/0/4	up	down	
ge-1/0/4.0	up	down	eth-switch
ge-1/0/5	up	down	
ge-1/0/5.0	up	down	eth-switch
ge-1/0/6	up	down	
ge-1/0/6.0	up	down	eth-switch
ge-1/0/7	up	down	
ge-1/0/7.0	up	down	eth-switch
ge-1/0/8	up	down	
ge-1/0/8.0	up	down	eth-switch
ge-1/0/9	up	down	
ge-1/0/9.0	up	down	eth-switch
ge-1/0/10	up	down	
ge-1/0/10.0	up	down	eth-switch
ge-1/0/11	up	down	
ge-1/0/11.0	up	down	eth-switch
ge-1/0/12	up	down	

ge-1/0/12.0	up	down	eth-switch
ge-1/0/13	up	down	
ge-1/0/13.0	up	down	eth-switch
ge-1/0/14	up	down	
ge-1/0/14.0	up	down	eth-switch
ge-1/0/15	up	down	
ge-1/0/15.0	up	down	eth-switch
ge-1/0/16	up	down	
ge-1/0/16.0	up	down	eth-switch
ge-1/0/17	up	down	
ge-1/0/17.0	up	down	eth-switch
ge-1/0/18	up	down	
ge-1/0/18.0	up	down	eth-switch
ge-1/0/19	up	down	
ge-1/0/19.0	up	down	eth-switch
ge-1/0/20	up	down	
ge-1/0/20.0	up	down	eth-switch
ge-1/0/21	up	down	
ge-1/0/21.0	up	down	eth-switch
ge-1/0/22	up	down	
ge-1/0/22.0	up	down	eth-switch
ge-1/0/23	up	down	
ge-1/0/23.0	up	down	eth-switch
ge-1/0/24	up	down	
ge-1/0/24.0	up	down	eth-switch
ge-1/0/25	up	down	
ge-1/0/25.0	up	down	eth-switch
ge-1/0/26	up	down	
ge-1/0/26.0	up	down	eth-switch
ge-1/0/27	up	down	
ge-1/0/27.0	up	down	eth-switch
ge-1/0/28	up	down	
ge-1/0/28.0	up	down	eth-switch
ge-1/0/29	up	down	
ge-1/0/29.0	up	down	eth-switch
ge-1/0/30	up	down	
ge-1/0/30.0	up	down	eth-switch
ge-1/0/31	up	down	
ge-1/0/31.0	up	down	eth-switch
ge-1/0/32	up	down	
ge-1/0/33	up	down	
ge-1/0/34	up	down	
ge-1/0/35	up	down	
ge-1/0/36	up	down	
ge-1/0/36.0	up	down	eth-switch
ge-1/0/37	up	down	
ge-1/0/37.0	up	down	eth-switch
ge-1/0/38	up	down	
ge-1/0/38.0	up	down	eth-switch
ge-1/0/39	up	up	
ge-1/0/39.0	up	up	eth-switch
ge-1/0/40	up	up	
ge-1/0/40.0	up	up	aenet --> ae0.0
ge-1/0/41	up	down	
ge-1/0/41.0	up	down	aenet --> ae1.0
ge-1/0/42	up	down	
ge-1/0/42.0	up	down	aenet --> ae2.0
ge-1/0/43	up	down	
ge-1/0/43.0	up	down	aenet --> ae3.0
ge-1/0/44	up	down	
ge-1/0/44.0	up	down	aenet --> ae4.0

ge-1/0/45	up	down	
ge-1/0/45.0	up	down	aenet --> ae5.0
ge-1/0/46	up	down	
ge-1/0/46.0	up	down	aenet --> ae6.0
ge-1/0/47	up	down	
ge-1/0/47.0	up	down	aenet --> ae7.0
xe-1/1/0	up	up	
xe-1/1/0.0	up	up	aenet --> ae8.0
xe-1/1/1	up	down	
xe-1/1/2	up	up	
xe-1/1/2.0	up	up	aenet --> ae8.0
xe-1/1/3	up	down	
ge-2/0/0	up	down	
ge-2/0/0.0	up	down	eth-switch
ge-2/0/1	up	down	
ge-2/0/1.0	up	down	eth-switch
ge-2/0/2	up	down	
ge-2/0/2.0	up	down	eth-switch
ge-2/0/3	up	down	
ge-2/0/3.0	up	down	eth-switch
ge-2/0/4	up	down	
ge-2/0/4.0	up	down	eth-switch
ge-2/0/5	up	down	
ge-2/0/5.0	up	down	eth-switch
ge-2/0/6	up	down	
ge-2/0/6.0	up	down	eth-switch
ge-2/0/7	up	down	
ge-2/0/7.0	up	down	eth-switch
ge-2/0/8	up	down	
ge-2/0/8.0	up	down	eth-switch
ge-2/0/9	up	down	
ge-2/0/9.0	up	down	eth-switch
ge-2/0/10	up	down	
ge-2/0/10.0	up	down	eth-switch
ge-2/0/11	up	down	
ge-2/0/11.0	up	down	eth-switch
ge-2/0/12	up	down	
ge-2/0/12.0	up	down	eth-switch
ge-2/0/13	up	down	
ge-2/0/13.0	up	down	eth-switch
ge-2/0/14	up	down	
ge-2/0/14.0	up	down	eth-switch
ge-2/0/15	up	down	
ge-2/0/15.0	up	down	eth-switch
ge-2/0/16	up	down	
ge-2/0/16.0	up	down	eth-switch
ge-2/0/17	up	down	
ge-2/0/17.0	up	down	eth-switch
ge-2/0/18	up	down	
ge-2/0/18.0	up	down	eth-switch
ge-2/0/19	up	down	
ge-2/0/19.0	up	down	eth-switch
ge-2/0/20	up	down	
ge-2/0/20.0	up	down	eth-switch
ge-2/0/21	up	down	
ge-2/0/21.0	up	down	eth-switch
ge-2/0/22	up	down	
ge-2/0/22.0	up	down	eth-switch
ge-2/0/23	up	down	
ge-2/0/23.0	up	down	eth-switch
ge-2/0/24	up	down	

ge-2/0/24.0	up	down	eth-switch
ge-2/0/25	up	down	
ge-2/0/25.0	up	down	eth-switch
ge-2/0/26	up	down	
ge-2/0/26.0	up	down	eth-switch
ge-2/0/27	up	down	
ge-2/0/27.0	up	down	eth-switch
ge-2/0/28	up	down	
ge-2/0/28.0	up	down	eth-switch
ge-2/0/29	up	down	
ge-2/0/29.0	up	down	eth-switch
ge-2/0/30	up	down	
ge-2/0/30.0	up	down	eth-switch
ge-2/0/31	up	down	
ge-2/0/31.0	up	down	eth-switch
ge-2/0/32	up	down	
ge-2/0/33	up	down	
ge-2/0/34	up	down	
ge-2/0/35	up	down	
ge-2/0/36	up	down	
ge-2/0/36.0	up	down	eth-switch
ge-2/0/37	up	down	
ge-2/0/37.0	up	down	eth-switch
ge-2/0/38	up	down	
ge-2/0/38.0	up	down	eth-switch
ge-2/0/39	up	up	
ge-2/0/39.0	up	up	eth-switch
ge-2/0/40	up	up	
ge-2/0/40.0	up	up	aenet --> ae0.0
ge-2/0/41	up	down	
ge-2/0/41.0	up	down	aenet --> ae1.0
ge-2/0/42	up	down	
ge-2/0/42.0	up	down	aenet --> ae2.0
ge-2/0/43	up	down	
ge-2/0/43.0	up	down	aenet --> ae3.0
ge-2/0/44	up	down	
ge-2/0/44.0	up	down	aenet --> ae4.0
ge-2/0/45	up	down	
ge-2/0/45.0	up	down	aenet --> ae5.0
ge-2/0/46	up	down	
ge-2/0/46.0	up	down	aenet --> ae6.0
ge-2/0/47	up	down	
ge-2/0/47.0	up	down	aenet --> ae7.0
xe-2/1/0	up	up	
xe-2/1/0.0	up	up	aenet --> ae8.0
xe-2/1/1	up	down	
xe-2/1/2	up	up	
xe-2/1/2.0	up	up	aenet --> ae8.0
xe-2/1/3	up	down	
ge-3/0/0	up	down	
ge-3/0/0.0	up	down	eth-switch
ge-3/0/1	up	down	
ge-3/0/1.0	up	down	eth-switch
ge-3/0/2	up	down	
ge-3/0/2.0	up	down	eth-switch
ge-3/0/3	up	down	
ge-3/0/3.0	up	down	eth-switch
ge-3/0/4	up	down	
ge-3/0/4.0	up	down	eth-switch
ge-3/0/5	up	down	
ge-3/0/5.0	up	down	eth-switch

ge-3/0/6	up	down
ge-3/0/6.0	up	down eth-switch
ge-3/0/7	up	down
ge-3/0/7.0	up	down eth-switch
ge-3/0/8	up	down
ge-3/0/8.0	up	down eth-switch
ge-3/0/9	up	down
ge-3/0/9.0	up	down eth-switch
ge-3/0/10	up	down
ge-3/0/10.0	up	down eth-switch
ge-3/0/11	up	down
ge-3/0/11.0	up	down eth-switch
ge-3/0/12	up	down
ge-3/0/12.0	up	down eth-switch
ge-3/0/13	up	down
ge-3/0/13.0	up	down eth-switch
ge-3/0/14	up	down
ge-3/0/14.0	up	down eth-switch
ge-3/0/15	up	down
ge-3/0/15.0	up	down eth-switch
ge-3/0/16	up	down
ge-3/0/16.0	up	down eth-switch
ge-3/0/17	up	down
ge-3/0/17.0	up	down eth-switch
ge-3/0/18	up	down
ge-3/0/18.0	up	down eth-switch
ge-3/0/19	up	down
ge-3/0/19.0	up	down eth-switch
ge-3/0/20	up	down
ge-3/0/20.0	up	down eth-switch
ge-3/0/21	up	down
ge-3/0/21.0	up	down eth-switch
ge-3/0/22	up	down
ge-3/0/22.0	up	down eth-switch
ge-3/0/23	up	down
ge-3/0/23.0	up	down eth-switch
ge-3/0/24	up	down
ge-3/0/24.0	up	down eth-switch
ge-3/0/25	up	down
ge-3/0/25.0	up	down eth-switch
ge-3/0/26	up	down
ge-3/0/26.0	up	down eth-switch
ge-3/0/27	up	down
ge-3/0/27.0	up	down eth-switch
ge-3/0/28	up	down
ge-3/0/28.0	up	down eth-switch
ge-3/0/29	up	down
ge-3/0/29.0	up	down eth-switch
ge-3/0/30	up	down
ge-3/0/30.0	up	down eth-switch
ge-3/0/31	up	down
ge-3/0/31.0	up	down eth-switch
ge-3/0/32	up	down
ge-3/0/33	up	down
ge-3/0/34	up	down
ge-3/0/35	up	down
ge-3/0/36	up	down
ge-3/0/36.0	up	down eth-switch
ge-3/0/37	up	down
ge-3/0/37.0	up	down eth-switch
ge-3/0/38	up	down

ge-3/0/38.0	up	down	eth-switch	
ge-3/0/39	up	up		
ge-3/0/39.0	up	up	eth-switch	
ge-3/0/40	up	down		
ge-3/0/40.0	up	down	aenet	--> ae0.0
ge-3/0/41	up	down		
ge-3/0/41.0	up	down	aenet	--> ae1.0
ge-3/0/42	up	down		
ge-3/0/42.0	up	down	aenet	--> ae2.0
ge-3/0/43	up	down		
ge-3/0/43.0	up	down	aenet	--> ae3.0
ge-3/0/44	up	down		
ge-3/0/44.0	up	down	aenet	--> ae4.0
ge-3/0/45	up	down		
ge-3/0/45.0	up	down	aenet	--> ae5.0
ge-3/0/46	up	down		
ge-3/0/46.0	up	down	aenet	--> ae6.0
ge-3/0/47	up	down		
ge-3/0/47.0	up	down	aenet	--> ae7.0
xe-3/1/0	up	up		
xe-3/1/0.0	up	up	aenet	--> ae8.0
xe-3/1/1	up	down		
xe-3/1/2	up	up		
xe-3/1/2.0	up	up	aenet	--> ae8.0
xe-3/1/3	up	down		
vcp-0	up	down		
vcp-0.32768	up	down		
vcp-1	up	up		
vcp-1.32768	up	up		
ae0	up	up		
ae0.0	up	up	eth-switch	
ae1	up	down		
ae1.0	up	down	eth-switch	
ae2	up	down		
ae2.0	up	down	eth-switch	
ae3	up	down		
ae3.0	up	down	eth-switch	
ae4	up	down		
ae4.0	up	down	eth-switch	
ae5	up	down		
ae5.0	up	down	eth-switch	
ae6	up	down		
ae6.0	up	down	eth-switch	
ae7	up	down		
ae7.0	up	down	eth-switch	
ae8	up	up		
ae8.0	up	up	eth-switch	
ae9	up	down		
bme0	up	up		
bme0.32768	up	up	inet	128.0.0.1/2 128.0.0.16/2 128.0.0.32/2 tnp 0x10
bme0.32770	up	up	eth-switch	
bme0.32771	down	up	eth-switch	
bme0.32772	down	up	eth-switch	
dsc	up	up		
gre	up	up		
ipip	up	up		
jsrv	up	up		
jsrv.1	up	up	inet	128.0.0.127/2

```

lo0          up    up
lo0.0        up    up    inet    10.255.195.97    --> 0/0
                                iso
47.0005.80ff.f800.0000.0108.0001.0102.5519.5097
                                inet6   abcd::10:255:195:97
                                                fe80::2ac0:da0f:fc31:1e81

lsi          up    up
me0          up    up
me0.0        up    up    inet    10.94.195.97/24
mtun         up    up
pimd         up    up
pime         up    up
tap          up    up
vlan         up    up
vme          up    up    inet    192.168.157.27/24

```

Meaning In the output of the **show virtual-chassis status** command, if all four members appear, the Virtual Chassis is operational.

In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as up (such as ge-0/0/39, ge-1/0/39, ge-2/0/39, and ge-3/0/39 for the Interconnect devices; ge-0/0/40, ge-1/0/40, and ge-2/0/40 for the Director devices; ge-0/0/0 through ge-0/0/7 for the Node devices; and xe-0/1/0, xe-0/1/2, xe-1/1/0, xe-1/1/2, xe-2/1/0, xe-2/1/2, xe-3/1/0, and xe-3/1/2 for the inter-Virtual Chassis connections), the control plane is properly connected.

Related Documentation

- *QFX3000-G QFabric System Installation Overview*
- *Installing and Connecting a QFX3100 Director Device*
- *Installing and Connecting a QFX3008-I Interconnect Device*
- *Installing and Connecting a QFX3500 Device*
- *Installing and Connecting an EX4200 Switch*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- [Understanding the QFabric System Control Plane on page 1196](#)

Importing a QFX3000-G QFabric System Control Plane Virtual Chassis Configuration with a USB Flash Drive

There are two methods of importing the configuration file to the QFabric control plane Virtual Chassis. You can load the configuration file onto a USB flash drive from the Juniper Networks software download site before inserting the USB flash drive into the Virtual Chassis USB port, or you can copy and paste the configuration from the following example (see “[Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane](#)” on page 1263).

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see *Installing and Connecting a QFX3100 Director Device*, *Installing and Connecting a QFX3008-I Interconnect Device*, and *Installing and Connecting a QFX3500 Device*.
- Rack, mount, and install your Virtual Chassis hardware (EX4200 switches). For more information, see *Installing and Connecting an EX4200 Switch*.
- Create two Virtual Chassis of four members each. For more information, see *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- Select a USB flash drive that meets the QFabric control plane Virtual Chassis USB port specifications. See *USB Port Specifications for an EX Series Switch*.
- Use a computer or other device to load the configuration file from the Internet and copy it to the USB flash drive.

To import the Virtual Chassis configuration file into a USB flash drive:

1. In a browser, go to <http://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. Click **QFX3100** in the **QFX Series** section.

The QFX3100 Download Software page appears.

3. From the **Release** list, select the number of the software version for which you want to download the Virtual Chassis configuration file.
4. Select the **Software** tab and then click **QFX3000-G QFabric System - Control Plane Virtual Chassis Configuration** in the **QFabric System Install Package and Media** section.
A login page appears.
5. Enter your user ID and password and press **Enter**.

6. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
7. Save the Virtual Chassis configuration file onto the USB flash drive using your computer or other device.
8. Remove the USB flash drive from the computer or other device.
9. Insert the USB flash drive into the USB port on the EX4200 switch.
10. Save the file to `/var/home/username`.
11. Load the configuration file into the switch.

```
user@switch# load override filename
```
12. Commit the configuration.

```
user@switch# commit
Load complete
```
13. Remove the USB flash drive from the switch.

**Related
Documentation**

- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)

Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane

This example shows you how to connect QFabric system components and configure the EX4200 switches used for the copper-based QFX3000-M QFabric system control plane network. Proper wiring of Director devices, Interconnect devices, and Node devices to the EX4200 switches, combined with a standard configuration, enables you to bring up the internal QFabric system management network and prepare your QFabric system for full operation.



NOTE: The EX4200 switch configuration is the same for both the copper-based and fiber-based QFX3000-M QFabric system control plane networks. Hence, a separate example for configuring EX4200 switches for the fiber-based control plane network is not provided. However, before you use this example to configure a fiber-based control plane network, ensure that you have installed and wired the QFabric system hardware and EX4200 switches as required for a fiber-based control plane network (see *QFX3000-M QFabric System Installation Overview*).

- [Requirements on page 1310](#)
- [Overview on page 1310](#)
- [Configuration on page 1317](#)
- [Verification on page 1329](#)

Requirements

This example uses the following hardware and software components:

- One QFX3000-M QFabric system containing:
 - Two QFX3100 Director devices with 1000BASE-T network modules installed
 - Two QFX3600-I Interconnect devices
 - Eight QFX3500 Node devices with a 1000BASE-T management board installed
- Two EX4200-24T switches with SFP+ uplink module installed
- Junos OS Release 12.2 for the QFabric system
- Junos OS Release 12.1R1.10 for the EX Series switches

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see *Installing and Connecting a QFX3100 Director Device*, *Installing and Connecting a QFX3600 or QFX3600-I Device*, and *Installing and Connecting a QFX3500 Device*.
- Rack, mount, and install your EX4200 switches. For more information, see *Installing and Connecting an EX4200 Switch*.

Overview

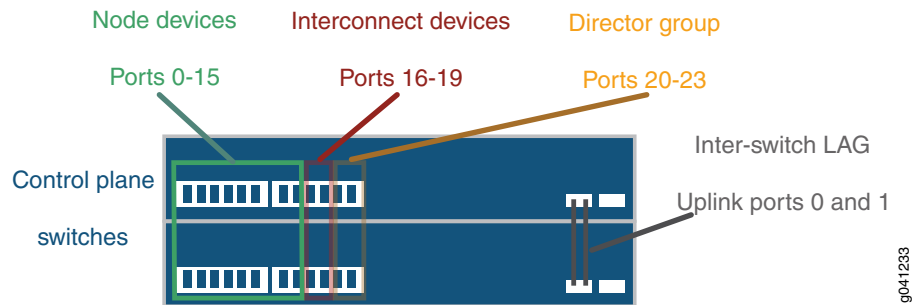
The QFX3000-M QFabric system control plane network connects the Director group, Interconnect devices, and Node devices in a QFabric system across a pair of redundant EX4200 switches. By separating the management control plane from the data plane, the QFabric system can scale efficiently. The copper-based control plane network uses Gigabit Ethernet cabling and connections between components, and two 1-Gigabit Ethernet connections configured in a link aggregation group (LAG) between the redundant EX4200 switches.

Specific ports have been reserved on the EX4200 switches to connect to each of the QFabric system device types. Such design simplifies installation and facilitates timely deployment of a QFabric system. It also permits the use of a standard EX4200 switch configuration included as part of this example. The standard configuration can scale from the 8 Node devices shown in this example to a maximum of 16 Node devices.

Topology

[Figure 31 on page 1311](#) shows the general port ranges where QFabric system devices must be connected to the EX4200 switches. For each EX4200 switch, connect ports **0** through **15** to Node devices, ports **16** through **19** to Interconnect devices, and ports **20** through **23** to Director devices. [Table 133 on page 1311](#) shows the details of the QFabric system component-to-EX4200 switch port mappings.

Figure 31: QFX3000-M QFabric System Control Plane—EX4200 Switch Port Ranges



CAUTION:

- The control plane network within a QFabric system is a critical component of the system that should not be shared with other network traffic. In order to scale efficiently, the control plane network must be reserved for the QFabric system and its components. As a result, the ports of the QFabric system control plane must never be used for any purpose other than to transport QFabric system control plane traffic, and we neither recommend nor support the connection of other devices to the QFabric system control plane network.
- Do not install Junos Space and AI-Scripts (AIS) on the control plane network EX4200 switches in a QFX3000-M QFabric system.

Table 133 on page 1311 shows the specific mappings of QFabric system control plane network ports from the QFabric system components to the EX4200 switches.



NOTE: The uplink ports 2 and 3 on the EX4200 switches are reserved for future use.

Table 133: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX4200 Switch Port Mappings

EX4200 Switch 1 (EX0)	EX4200 Switch 2 (EX1)	QFabric System Component
Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)	Node device 0
Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)	Node device 1
Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)	Node device 2
Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)	Node device 3

Table 133: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX4200 Switch Port Mappings (*continued*)

EX4200 Switch 1 (EX0)	EX4200 Switch 2 (EX1)	QFabric System Component
Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)	Node device 4
Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)	Node device 5
Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)	Node device 6
Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)	Node device 7
Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)	Node device 8
Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)	Node device 9
Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)	Node device 10
Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)	Node device 11
Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)	Node device 12
Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)	Node device 13
Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)	Node device 14
Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)	Node device 15
IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)	Interconnect device 0
IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)	Interconnect device 1
IC2, management port C0 to port 18 (ge-0/0/18)	IC2, management port C1 to port 18 (ge-0/0/18)	Interconnect device 2
IC3, management port C0 to port 19 (ge-0/0/19)	IC3, management port C1 to port 19 (ge-0/0/19)	Interconnect device 3

Table 133: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX4200 Switch Port Mappings (*continued*)

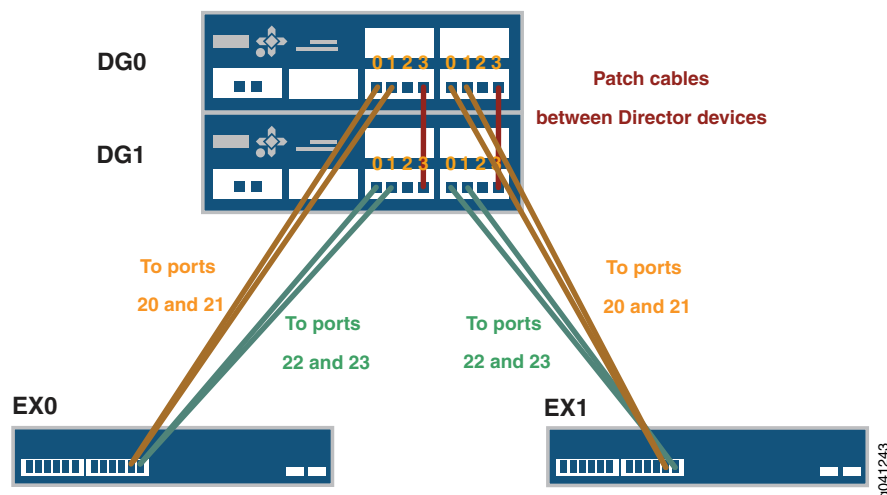
EX4200 Switch 1 (EX0)	EX4200 Switch 2 (EX1)	QFabric System Component
DG0 module 0, port 0 to port 20 (ge-0/0/20)	DG0 module 1, port 0 to port 20 (ge-0/0/20)	Director device 0
DG0 module 0, port 1 to port 21 (ge-0/0/21)	DG0 module 1, port 1 to port 21 (ge-0/0/21)	Director device 0
DG1 module 0, port 0 to port 22 (ge-0/0/22)	DG1 module 1, port 0 to port 22 (ge-0/0/22)	Director device 1
DG1 module 0, port 1 to port 23 (ge-0/0/23)	DG1 module 1, port 1 to port 23 (ge-0/0/23)	Director device 1
EX0, uplink port 0 to EX1, uplink port 0 (ge-0/1/0)	EX1, uplink port 0 to EX0, uplink port 0 (ge-0/1/0)	Inter-EX4200 switch LAG
EX0, uplink port 1 to EX1, uplink port 1 (ge-0/1/1)	EX1, uplink port 1 to EX0, uplink port 1 (ge-0/1/1)	Inter-EX4200 switch LAG
Reserved	Reserved	Future use
Uplink port 2 (ge-0/1/2)	Uplink port 2 (ge-0/1/2)	
Reserved	Reserved	Future use
Uplink port 3 (ge-0/1/3)	Uplink port 3 (ge-0/1/3)	

Next, connect the Director devices to the EX4200 switches. In general, you want to accomplish the following:

- Connect two ports from one network module in a Director device to the first EX4200 switch, and two ports from the second network module to the second EX4200 switch.
- Connect the Director devices to each other and create a Director group. You can use either straight-through RJ-45 patch cables or crossover cables, because the Director devices contain autosensing modules. Connect one port from each network module on the first Director device to one port in each network module on the second Director device.

[Figure 32 on page 1314](#) shows the specific ports on the Director group that you must connect to the EX4200 switches and interconnect between the Director devices.

Figure 32: QFX3000-M QFabric System Control Plane—Director Group to EX4200 Switch Connections



In this specific example, connect ports **0** and **1** from module **0** on Director device DG0 to ports **20** and **21** on EX4200 switch EX0 (ge-0/0/20 and ge-0/0/21), and connect ports **0** and **1** from module **1** to ports **20** and **21** on the *second* EX4200 switch EX1 (ge-0/0/20 and ge-0/0/21).

For Director device DG1, connect ports **0** and **1** from module **0** to ports **22** and **23** on EX4200 switch EX0 (ge-0/0/22 and ge-0/0/23), and connect ports **0** and **1** from module **1** to ports **22** and **23** on the *second* EX4200 switch EX1 (ge-0/0/22 and ge-0/0/23).

To form the Director group, connect port **3** on module **0** on Director device DG0 to port **3** on module **0** on Director device DG1. Similarly, connect port **3** on module **1** on Director device DG0 to port **3** on module **1** on Director device DG1. [Table 134 on page 1314](#) shows the port mappings for the Director group in this example.

Table 134: Director Group Port Mappings

Director Device	EX4200 Switch EX0	EX4200 Switch EX1
DG0	<ul style="list-style-type: none"> DG0 module 0, port 0 to port 20 on EX0 (ge-0/0/20) DG0 module 0, port 1 to port 21 on EX0 (ge-0/0/21) DG0 module 0, port 3 to module 0, port 3 on DG1 	<ul style="list-style-type: none"> DG0 module 1, port 0 to port 20 on EX1 (ge-0/0/20) DG0 module 1, port 1 to port 21 on EX1 (ge-0/0/21) DG0 module 1, port 3 to module 1, port 3 on DG1
DG1	<ul style="list-style-type: none"> DG1 module 0, port 0 to port 22 on EX0 (ge-0/0/22) DG1 module 0, port 1 to port 23 on EX0 (ge-0/0/23) DG1 module 0, port 3 to module 0, port 3 on DG0 	<ul style="list-style-type: none"> DG1 module 1, port 0 to port 22 on EX1 (ge-0/0/22) DG1 module 1, port 1 to port 23 on EX1 (ge-0/0/23) DG1 module 1, port 3 to module 1, port 3 on DG0

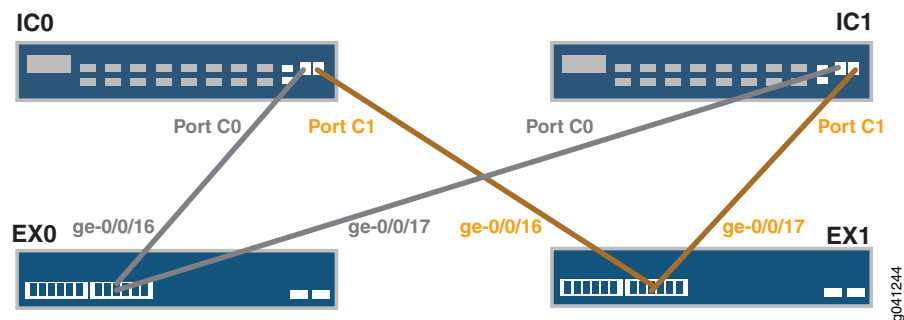
In the software, the ports of each network module on a Director device are reversed, numbered from right to left, and incremented sequentially across modules. If you issue interface operational commands directly on the Director device, note the following port mappings as shown in [Table 128 on page 1271](#):

Table 135: Hardware to Software Port Mappings for Director Device Network Modules

Network Module	Port 0	Port 1	Port 2	Port 3
Module 0	eth5	eth4	eth3	eth2
Module 1	eth9	eth8	eth7	eth6

[Figure 33 on page 1315](#) shows the specific ports on the QFX3600-I Interconnect devices that you must connect to the EX4200 switches. In general, connect the first management port in an Interconnect device to the first EX4200 switch, and the second management port to the second EX4200 switch.

Figure 33: QFX3000-M QFabric System Control Plane—Interconnect Device to EX4200 Switch Connections



In this specific example, for both Interconnect devices IC0 and IC1, connect management port **C0** to EX4200 switches EX0 and EX1 and management port **C1** to EX4200 switches EX0 and EX1. Connect the management port **C0** cables to port 16 on EX4200 switches EX0 and EX1 (ge-0/0/16), and connect the management port **C1** cables to port 17 on EX4200 switches EX0 and EX1 (ge-0/0/17). [Table 136 on page 1315](#) shows the port mappings for the Node devices in this example.


Table 136: Interconnect Device Port Mappings

Interconnect Device	EX4200 Switch EX0	EX4200 Switch EX1
IC0	IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)
IC1	IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)

[Figure 34 on page 1316](#) shows the specific ports on the Node devices that you must connect to the EX4200 switches. In general, connect the first management port from a Node

device to the first EX4200 switch, and the second management port to the second EX4200 switch.

Figure 34: QFX3000-M QFabric System Control Plane—Node Device to EX4200 Switch Connections



In this specific example, for Node device Node0, connect management port **C0** (also known as me5) to EX4200 switch EX0 port **0** (ge-0/0/0), and connect management port **C1** (also known as me6) to the *second* EX4200 switch EX1 port **0** (ge-0/0/0).

For the remaining seven Node devices, connect management port **C0** to the ge-0/0/X port on EX4200 switch EX0 that matches the Node device number. Similarly, connect management port **C1** to the port on the *second* EX4200 switch EX1 that matches the Node device number. For example, you would connect Node device Node5 to port **5** (ge-0/0/5). [Table 131 on page 1274](#) shows the full set of port mappings for the Node devices in this example.

Table 137: Node Device to EX4200 Switch Port Mappings

Node Device	EX4200 Switch EX0	EX4200 Switch EX1
Node0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)
Node4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)

Table 137: Node Device to EX4200 Switch Port Mappings (*continued*)

Node Device	EX4200 Switch EX0	EX4200 Switch EX1
Node7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)

Figure 30 on page 1274 shows the specific uplink ports on the first EX4200 switch that you must connect to the second EX4200 switch. These connections create a link aggregation group (LAG) that provides redundancy and resiliency for the EX4200 switch portion of the control plane. In general, connect each 1-Gigabit Ethernet uplink port from the first EX4200 switch to the corresponding 1-Gigabit Ethernet uplink port on the second EX4200 switch.

Figure 35: QFX3000-M QFabric System Control Plane—Inter-EX4200 Switch LAG Connections



In this specific example, for EX4200 switch EX0, connect uplink port 0 (ge-0/1/0) to EX4200 switch EX1 uplink port 0 (ge-0/1/0). Then connect uplink port 1 (ge-0/1/1) on EX4200 switch EX0 to uplink port 1 (ge-0/1/1) on EX4200 switch EX1.

Table 132 on page 1275 shows the port mappings for the EX4200 switch LAG connections in this example.

Table 138: EX4200 Switch LAG Port Mappings

EX0 and EX1	EX0	EX1
Uplink port 0	ge-0/1/0 to ge-0/1/0	ge-0/1/0 to ge-0/1/0
Uplink port 1	ge-0/1/1 to ge-0/1/1	ge-0/1/1 to ge-0/1/1

Configuration

CLI Quick Configuration

To configure the QFX3000-M QFabric system control plane EX4200 switches quickly, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network, and then copy and paste the commands into the EX4200 switch CLI at the **[edit]** hierarchy level.

```
set groups qfabric chassis aggregated-devices ethernet device-count 3
set groups qfabric chassis alarm management-ethernet link-down ignore
set groups qfabric chassis lcd-menu fpc 0 menu-item maintenance-menu disable
set groups qfabric protocols rstp interface ae2.0 mode point-to-point
set groups qfabric protocols rstp interface all edge
set groups qfabric protocols rstp interface all no-root-port
set groups qfabric protocols rstp bpdu-block-on-edge
set groups qfabric protocols lldp interface all
```

```
set groups qfabric ethernet-switching-options storm-control interface all bandwidth
10000
set groups qfabric vlans qfabric vlan-id 100
set groups qfabric vlans qfabric dot1q-tunneling
set groups qfabric-int interfaces <*> mtu 9216
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching port-mode access
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching vlan members
qfabric
set groups qfabric-ae interfaces <*> aggregated-ether-options link-speed 1g
set groups qfabric-ae interfaces <*> aggregated-ether-options lacp active
set apply-groups qfabric
set chassis fpc 0 pic 1 sfppplus pic-mode 1g
set interfaces interface-range Node_Device_Interfaces member "ge-0/0/[0-15]"
set interfaces interface-range Node_Device_Interfaces description "QFabric Node Device"
set interfaces interface-range Node_Device_Interfaces mtu 9216
set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching
port-mode access
set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching
vlan members qfabric
set interfaces interface-range Interconnect_Device_Interfaces member "ge-0/0/[16-17]"
set interfaces interface-range Interconnect_Device_Interfaces description "QFabric
Interconnect Device"
set interfaces interface-range Interconnect_Device_Interfaces mtu 9216
set interfaces interface-range Interconnect_Device_Interfaces unit 0 family
ethernet-switching port-mode access
set interfaces interface-range Interconnect_Device_Interfaces unit 0 family
ethernet-switching vlan members qfabric
set interfaces interface-range Director_Device_DG0_LAG_Interfaces member
"ge-0/0/[20-21]"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces description "QFabric
Director Device - DG0"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options speed
1g
set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options 802.3ad
ae0
set interfaces interface-range Director_Device_DG1_LAG_Interfaces member
"ge-0/0/[22-23]"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces description "QFabric
Director Device - DG1"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options speed
1g
set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options 802.3ad
ae1
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces member "ge-0/1/[0-1]"
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces description "QFabric
Control Plane (Inter - Switch LAG)"
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces ether-options 802.3ad
ae2
set interfaces ae0 apply-groups qfabric-int
set interfaces ae0 apply-groups qfabric-ae
set interfaces ae0 description "QFabric Director Device - DG0"
set interfaces ae1 apply-groups qfabric-int
set interfaces ae1 apply-groups qfabric-ae
set interfaces ae1 description "QFabric Director Device - DG1"
set interfaces ae2 description "QFabric Control Plane (Inter-Switch LAG)"
set interfaces ae2 mtu 9216
```

```

set interfaces ae2 aggregated-ether-options link-speed 1g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 family ethernet-switching vlan members qfabric
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3
  loss-priority low code-points 110
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3
  loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2
  loss-priority low code-points 100
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2
  loss-priority high code-points 101
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0
  loss-priority low code-points 010
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0
  loss-priority high code-points 001
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_3 loss-priority low code-points 110
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_3 loss-priority low code-points 111
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_2 loss-priority low code-points 100
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_2 loss-priority high code-points 101
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_0 loss-priority low code-points 010
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class
  class_0 loss-priority high code-points 001
set class-of-service forwarding-classes class class_3 queue-num 7
set class-of-service forwarding-classes class class_2 queue-num 2
set class-of-service forwarding-classes class class_0 queue-num 0
set class-of-service interfaces ge-*/0/* scheduler-map cpe_network_smap
set class-of-service interfaces ge-*/0/* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
set class-of-service interfaces ge-*/0/* unit 0 classifiers inet-precedence
  IP_qfabric_classifier
set class-of-service interfaces ae* scheduler-map cpe_network_smap
set class-of-service interfaces ae* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
set class-of-service interfaces ae* unit 0 classifiers inet-precedence IP_qfabric_classifier
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_3
  scheduler scheduler_3
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_2
  scheduler scheduler_2
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_0
  scheduler scheduler_0
set class-of-service schedulers scheduler_3 buffer-size percent 30
set class-of-service schedulers scheduler_3 priority strict-high
set class-of-service schedulers scheduler_2 transmit-rate percent 75
set class-of-service schedulers scheduler_2 buffer-size percent 30
set class-of-service schedulers scheduler_2 priority low
set class-of-service schedulers scheduler_0 transmit-rate percent 25
set class-of-service schedulers scheduler_0 buffer-size percent 40
set class-of-service schedulers scheduler_0 priority low
set system host-name qfabric-control-plane
set system services ssh
set system services telnet
set system services web-management http
set system syslog user * any emergency

```

```
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file messages archive world-readable
set system syslog file messages explicit-priority
set system syslog file interactive-commands interactive-commands any
set system syslog file secure authorization info
set system syslog file default-log-messages any any
set system syslog file default-log-messages structured-data
set system syslog file console any error
set system syslog time-format millisecond
set interfaces me0 unit 0 family inet address 192.168.157.26/24
set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a EX4200 switch for the QFX3000-M QFabric system control plane network:

1. Create a configuration group called `qfabric` to define global QFabric system control plane properties. Set up the number of aggregated Ethernet devices, configure alarm and LCD management, activate loop prevention and storm control, specify a global VLAN (VLAN 100) and 802.1q tunneling, define options for aggregated Ethernet interfaces, and apply the `qfabric` group settings to the configuration.

```
[edit]
user@switch# set groups qfabric chassis aggregated-devices ethernet device-count
3
user@switch# set groups qfabric chassis alarm management-ethernet link-down
ignore
user@switch# set groups qfabric chassis lcd-menu fpc 0 menu-item
maintenance-menu disable
user@switch# set groups qfabric protocols rstp interface ae2.0 mode point-to-point
user@switch# set groups qfabric protocols rstp interface all edge
user@switch# set groups qfabric protocols rstp interface all no-root-port
user@switch# set groups qfabric protocols rstp bpdu-block-on-edge
user@switch# set groups qfabric protocols lldp interface all
user@switch# set groups qfabric ethernet-switching-options storm-control interface
all bandwidth 10000
user@switch# set groups qfabric vlans qfabric vlan-id 100
user@switch# set groups qfabric vlans qfabric dot1q-tunneling
user@switch# set groups qfabric-int interfaces <*> mtu 9216
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching
port-mode access
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching
vlan members qfabric
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options
link-speed 1g
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options lacp
active
user@switch# set apply-groups qfabric
```

2. Configure interfaces for the QFabric system control plane network. Enable the EX4200 switch SFP+ uplink module for 1-Gigabit Ethernet operation. Set the interface ranges where Node devices (0 through 15), Interconnect devices (16 and

17), and Director devices (20 through 23) connect to the control plane network through the EX4200 switches. Configure the inter-EX4200 switch LAG connections for the ae2 interface and apply the qfabric-int and qfabric-ae configuration groups to the aggregated Ethernet interfaces (ae0 and ae1) for the Director devices.

```
[edit]
user@switch# set chassis fpc 0 pic 1 sfppplus pic-mode lg
user@switch# set interfaces interface-range Node_Device_Interfaces member
  "ge-0/0/[0-15]"
user@switch# set interfaces interface-range Node_Device_Interfaces description
  "QFabric Node Device"
user@switch# set interfaces interface-range Node_Device_Interfaces mtu 9216
user@switch# set interfaces interface-range Node_Device_Interfaces unit 0 family
  ethernet-switching port-mode access
user@switch# set interfaces interface-range Node_Device_Interfaces unit 0 family
  ethernet-switching vlan members qfabric
user@switch# set interfaces interface-range Interconnect_Device_Interfaces member
  "ge-0/0/[16-17]"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces
  description "QFabric Interconnect Device"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces mtu
  9216
user@switch# set interfaces interface-range Interconnect_Device_Interfaces unit
  0 family ethernet-switching port-mode access
user@switch# set interfaces interface-range Interconnect_Device_Interfaces unit
  0 family ethernet-switching vlan members qfabric
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
  member "ge-0/0/[20-21]"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
  description "QFabric Director Device - DG0"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
  ether-options speed lg
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces
  ether-options 802.3ad ae0
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
  member "ge-0/0/[22-23]"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
  description "QFabric Director Device - DG1"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
  ether-options speed lg
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces
  ether-options 802.3ad ae1
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces
  member "ge-0/1/[0-1]"
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces
  description "QFabric Control Plane (Inter - Switch LAG)"
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces
  ether-options 802.3ad ae2
user@switch# set interfaces ae0 apply-groups qfabric-int
user@switch# set interfaces ae0 apply-groups qfabric-ae
user@switch# set interfaces ae0 description "QFabric Director Device - DG0"
user@switch# set interfaces ae1 apply-groups qfabric-int
user@switch# set interfaces ae1 apply-groups qfabric-ae
user@switch# set interfaces ae1 description "QFabric Director Device - DG1"
user@switch# set interfaces ae2 description "QFabric Control Plane (Inter-Switch
  LAG)"
```

```

user@switch# set interfaces ae2 mtu 9216
user@switch# set interfaces ae2 aggregated-ether-options link-speed 1g
user@switch# set interfaces ae2 aggregated-ether-options lacp active
user@switch# set interfaces ae2 unit 0 family ethernet-switching vlan members
qfabric

```

3. Enable class of service (CoS) for the QFabric system control plane network. Establish forwarding classes, priorities, scheduler maps, classifiers, and queues for three types of traffic: control traffic, interdevice traffic, and best-effort traffic.

```

[edit]
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 110
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 111
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_2 loss-priority low code-points 100
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_2 loss-priority high code-points 101
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_0 loss-priority low code-points 010
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier
forwarding-class class_0 loss-priority high code-points 001
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 110
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_3 loss-priority low code-points 111
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_2 loss-priority low code-points 100
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_2 loss-priority high code-points 101
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_0 loss-priority low code-points 010
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier
forwarding-class class_0 loss-priority high code-points 001
user@switch# set class-of-service forwarding-classes class class_3 queue-num 7
user@switch# set class-of-service forwarding-classes class class_2 queue-num 2
user@switch# set class-of-service forwarding-classes class class_0 queue-num 0
user@switch# set class-of-service interfaces ge-*/0/* scheduler-map
cpe_network_smap
user@switch# set class-of-service interfaces ge-*/0/* unit 0 classifiers ieee-802.1
onep_qfabric_classifier
user@switch# set class-of-service interfaces ge-*/0/* unit 0 classifiers
inet-precedence IP_qfabric_classifier
user@switch# set class-of-service interfaces ae* scheduler-map cpe_network_smap
user@switch# set class-of-service interfaces ae* unit 0 classifiers ieee-802.1
onep_qfabric_classifier
user@switch# set class-of-service interfaces ae* unit 0 classifiers inet-precedence
IP_qfabric_classifier
user@switch# set class-of-service scheduler-maps cpe_network_smap
forwarding-class class_3 scheduler scheduler_3
user@switch# set class-of-service scheduler-maps cpe_network_smap
forwarding-class class_2 scheduler scheduler_2
user@switch# set class-of-service scheduler-maps cpe_network_smap
forwarding-class class_0 scheduler scheduler_0
user@switch# set class-of-service schedulers scheduler_3 buffer-size percent 30

```

```

user@switch# set class-of-service schedulers scheduler_3 priority strict-high
user@switch# set class-of-service schedulers scheduler_2 transmit-rate percent
75
user@switch# set class-of-service schedulers scheduler_2 buffer-size percent 30
user@switch# set class-of-service schedulers scheduler_2 priority low
user@switch# set class-of-service schedulers scheduler_0 transmit-rate percent
25
user@switch# set class-of-service schedulers scheduler_0 buffer-size percent 40
user@switch# set class-of-service schedulers scheduler_0 priority low

```

4. Configure settings to enable the EX4200 switches to interoperate with your management network. Set a hostname, system services (such as Telnet), system log thresholds, management interface parameters, default routes, and any additional preferences you might have.

```

[edit]
user@switch# set system host-name qfabric-control-plane
user@switch# set system services ssh
user@switch# set system services telnet
user@switch# set system services web-management http
user@switch# set system syslog user * any emergency
user@switch# set system syslog file messages any notice
user@switch# set system syslog file messages authorization info
user@switch# set system syslog file messages archive world-readable
user@switch# set system syslog file messages explicit-priority
user@switch# set system syslog file interactive-commands interactive-commands
any
user@switch# set system syslog file secure authorization info
user@switch# set system syslog file default-log-messages any any
user@switch# set system syslog file default-log-messages structured-data
user@switch# set system syslog file console any error
user@switch# set system syslog time-format millisecond
user@switch# set interfaces me0 unit 0 family inet address 192.168.157.26/24
user@switch# set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1

```

Results To view the configuration, issue the **show** command in configuration mode or the **show configuration** command in operational mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

The following configuration is the standard configuration that applies universally to both EX4200 switches in your QFabric system control plane network.

```

[edit]
groups {
  qfabric {
    chassis {
      aggregated-devices {
        ethernet {
          device-count 3;
        }
      }
    }
    alarm {
      management-ethernet {
        link-down ignore;
      }
    }
  }
}

```

```
    }
    lcd-menu {
      fpc 0 {
        maintenance-menu disable;
      }
    }
  }
  protocols {
    rstp {
      interface ae2.0 {
        mode point-to-point;
      }
      interface all {
        edge;
        no-root-port;
      }
      bpdu-block-on-edge;
    }
    lldp {
      interface all;
    }
  }
  ethernet-switching-options {
    storm-control {
      interface all {
        bandwidth 10000;
      }
    }
  }
  vlans {
    qfabric {
      vlan-id 100;
      dot1q-tunneling;
    }
  }
  qfabric-int {
    interfaces {
      <*> {
        mtu 9216;
        unit 0 {
          family ethernet-switching {
            port-mode access;
            vlan {
              members qfabric;
            }
          }
        }
      }
    }
  }
  qfabric-ae {
    interfaces {
      <*> {
        aggregated-ether-options {
          link-speed 1g;
        }
      }
    }
  }
}
```



```
    member "ge-0/0/[22-23]";
    description "QFabric Director Device - DG1";
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
interface-range Control_Plane_Inter_LAG_Interfaces {
    member "ge-0/1/[0-1]";
    description "QFabric Control Plane (Inter-Switch LAG)";
    ether-options {
        802.3ad ae2;
    }
}
ae0 {
    apply-groups [ qfabric-int qfabric-ae ];
    description "QFabric Director Device - DG0";
}
ae1 {
    apply-groups [ qfabric-int qfabric-ae ];
    description "QFabric Director Device - DG1";
}
ae2 {
    description "QFabric Control Plane (Inter-Switch LAG)";
    mtu 9216;
    aggregated-ether-options {
        link-speed 1g;
        lacp {
            active;
        }
    }
    unit 0 {
        family ethernet-switching {
            vlan {
                members qfabric;
            }
        }
    }
}
}
class-of-service {
    classifiers {
        ieee-802.1 onep_qfabric_classifier {
            forwarding-class class_3 {
                loss-priority low code-points [ 110 111 ];
            }
            forwarding-class class_2 {
                loss-priority low code-points 100;
                loss-priority high code-points 101;
            }
            forwarding-class class_0 {
                loss-priority low code-points 010;
                loss-priority high code-points 001;
            }
        }
    }
}
```

```

    }
    inet-precedence IP_qfabric_classifier {
        forwarding-class class_3 {
            loss-priority low code-points [ 110 111 ];
        }
        forwarding-class class_2 {
            loss-priority low code-points 100;
            loss-priority high code-points 101;
        }
        forwarding-class class_0 {
            loss-priority low code-points 010;
            loss-priority high code-points 001;
        }
    }
}
forwarding-classes {
    class class_3 queue-num 7;
    class class_2 queue-num 2;
    class class_0 queue-num 0;
}
interfaces {
    ge-*/0/* {
        scheduler-map cpe_network_smap;
        unit 0 {
            classifiers {
                ieee-802.1 onep_qfabric_classifier;
                inet-precedence IP_qfabric_classifier;
            }
        }
    }
}
ae* {
    scheduler-map cpe_network_smap;
    unit 0 {
        classifiers {
            ieee-802.1 onep_qfabric_classifier;
            inet-precedence IP_qfabric_classifier;
        }
    }
}
}
scheduler-maps {
    cpe_network_smap {
        forwarding-class class_3 scheduler scheduler_3;
        forwarding-class class_2 scheduler scheduler_2;
        forwarding-class class_0 scheduler scheduler_0;
    }
}
schedulers {
    scheduler_3 {
        buffer-size percent 30;
        priority strict-high;
    }
    scheduler_2 {
        transmit-rate percent 75;
        buffer-size percent 30;
        priority low;
    }
}

```

```
    }
    scheduler_0 {
        transmit-rate percent 25;
        buffer-size percent 40;
        priority low;
    }
}
}
```

The following portion of the configuration applies to the specific requirements of your management network. Modify this section to meet the needs of your network.

```
[edit]
system {
    host-name qfabric-control-plane;
    services {
        ssh;
        telnet;
        web-management {
            http;
        }
    }
}
syslog {
    user * {
        any emergency;
    }
    file messages {
        any notice;
        authorization info;
        archive world-readable;
        explicit-priority;
    }
    file interactive-commands {
        interactive-commands any;
    }
    file secure {
        authorization info;
    }
    file default-log-messages {
        any any;
        structured-data;
    }
    file console {
        any error;
    }
    time-format millisecond;
}
}
interfaces {
    me0 {
        unit 0 {
            family inet {
                address 192.168.157.26/24;
            }
        }
    }
}
```

```

}
routing-options {
  static {
    route 0.0.0.0/0 next-hop 192.168.157.1;
  }
}

```

To verify the syntax of your configuration prior to committing it, enter **commit check** from configuration mode. If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the EX4200 switch configuration is working properly.

- [Verifying the QFX3000-M QFabric System Control Plane—EX4200 Switch EX0 on page 1329](#)
- [Verifying the QFX3000-M QFabric System Control Plane—EX4200 Switch EX1 on page 1331](#)

Verifying the QFX3000-M QFabric System Control Plane—EX4200 Switch EX0

Purpose Verify that the control plane is properly connected on your first EX4200 switch.

Action Connect to the Junos OS CLI of EX4200 switch EX0, either from your management network or from the console port of the switch. In operational mode, enter the **show interfaces terse** command.

Sample Output

```

user@ex0> show interfaces terse

```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	eth-switch		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	eth-switch		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch		
ge-0/0/7	up	up			
ge-0/0/7.0	up	up	eth-switch		
ge-0/0/8	up	down			
ge-0/0/8.0	up	down	eth-switch		
ge-0/0/9	up	down			
ge-0/0/9.0	up	down	eth-switch		
ge-0/0/10	up	down			
ge-0/0/10.0	up	down	eth-switch		
ge-0/0/11	up	down			
ge-0/0/11.0	up	down	eth-switch		
ge-0/0/12	up	down			

ge-0/0/12.0	up	down	eth-switch	
ge-0/0/13	up	down		
ge-0/0/13.0	up	down	eth-switch	
ge-0/0/14	up	down		
ge-0/0/14.0	up	down	eth-switch	
ge-0/0/15	up	down		
ge-0/0/15.0	up	down	eth-switch	
ge-0/0/16	up	up		
ge-0/0/16.0	up	up	eth-switch	
ge-0/0/17	up	up		
ge-0/0/17.0	up	up	eth-switch	
ge-0/0/18	up	down		
ge-0/0/18.0	up	down	eth-switch	
ge-0/0/19	up	down		
ge-0/0/19.0	up	down	eth-switch	
ge-0/0/20	up	up		
ge-0/0/20.0	up	up	aenet	--> ae0.0
ge-0/0/21	up	up		
ge-0/0/21.0	up	up	aenet	--> ae0.0
ge-0/0/22	up	up		
ge-0/0/22.0	up	up	aenet	--> ae1.0
ge-0/0/23	up	up		
ge-0/0/23.0	up	up	aenet	--> ae1.0
ge-0/1/0	up	up		
ge-0/1/0.0	up	up	aenet	--> ae2.0
ge-0/1/1	up	up		
ge-0/1/1.0	up	up	aenet	--> ae2.0
vcp-0	up	down		
vcp-0.32768	up	down		
vcp-1	up	down		
vcp-1.32768	up	down		
ae0	up	up		
ae0.0	up	up	eth-switch	
ae1	up	up		
ae1.0	up	up	eth-switch	
ae2	up	up		
ae2.0	up	up	eth-switch	
bme0	up	up		
bme0.32768	up	up	inet	128.0.0.1/2 128.0.0.16/2 128.0.0.32/2 tnp 0x10
dsc	up	up		
gre	up	up		
ipip	up	up		
lo0	up	up		
lo0.0	up	up	inet	127.0.0.1 --> 0/0
lsi	up	up		
me0	up	up		
me0.0	up	up	inet	192.168.157.26/24
mtun	up	up		
pimd	up	up		
pime	up	up		
tap	up	up		
vlan	up	up		
vme	up	down		

Meaning In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as **up** (such as ge-0/0/16 and ge-0/0/17 for the Interconnect devices; ge-0/0/20 through ge-0/0/23 for the Director devices; ge-0/0/0

through ge-0/0/7 for the Node devices; and ge-0/1/0 and ge-0/1/1 for the inter-EX4200 switch connections), the control plane is properly connected.

Verifying the QFX3000-M QFabric System Control Plane—EX4200 Switch EX1

Purpose Verify that the control plane is properly connected on your second EX4200 switch.

Action Connect to the Junos OS CLI of EX4200 switch EX1, either from your management network or from the console port of the switch. In operational mode, enter the **show interfaces terse** command.

Sample Output

```
user@ex1> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	eth-switch		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	eth-switch		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch		
ge-0/0/7	up	up			
ge-0/0/7.0	up	up	eth-switch		
ge-0/0/8	up	down			
ge-0/0/8.0	up	down	eth-switch		
ge-0/0/9	up	down			
ge-0/0/9.0	up	down	eth-switch		
ge-0/0/10	up	down			
ge-0/0/10.0	up	down	eth-switch		
ge-0/0/11	up	down			
ge-0/0/11.0	up	down	eth-switch		
ge-0/0/12	up	down			
ge-0/0/12.0	up	down	eth-switch		
ge-0/0/13	up	down			
ge-0/0/13.0	up	down	eth-switch		
ge-0/0/14	up	down			
ge-0/0/14.0	up	down	eth-switch		
ge-0/0/15	up	down			
ge-0/0/15.0	up	down	eth-switch		
ge-0/0/16	up	up			
ge-0/0/16.0	up	up	eth-switch		
ge-0/0/17	up	up			
ge-0/0/17.0	up	up	eth-switch		
ge-0/0/18	up	down			
ge-0/0/18.0	up	down	eth-switch		
ge-0/0/19	up	down			
ge-0/0/19.0	up	down	eth-switch		
ge-0/0/20	up	up			
ge-0/0/20.0	up	up	aenet	--> ae0.0	
ge-0/0/21	up	up			
ge-0/0/21.0	up	up	aenet	--> ae0.0	

```

ge-0/0/22          up    up
ge-0/0/22.0        up    up    aenet    --> ae1.0
ge-0/0/23          up    up
ge-0/0/23.0        up    up    aenet    --> ae1.0
ge-0/1/0           up    up
ge-0/1/0.0         up    up    aenet    --> ae2.0
ge-0/1/1           up    up
ge-0/1/1.0         up    up    aenet    --> ae2.0
vcp-0              up    down
vcp-0.32768        up    down
vcp-1              up    down
vcp-1.32768        up    down
ae0                up    up
ae0.0              up    up    eth-switch
ae1                up    up
ae1.0              up    up    eth-switch
ae2                up    up
ae2.0              up    up    eth-switch
bme0               up    up
bme0.32768         up    up    inet      128.0.0.1/2
                                   128.0.0.16/2
                                   128.0.0.32/2
                                   tnp      0x10
dsc                up    up
gre                up    up
ipip               up    up
lo0                up    up
lo0.0              up    up    inet      127.0.0.1          --> 0/0
lsi                up    up
me0                up    up
me0.0              up    up    inet      192.168.157.26/24
mtun               up    up
pimd               up    up
pime               up    up
tap                up    up
vlan               up    up
vme                up    down

```

Meaning In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as **up** (such as ge-0/0/16 and ge-0/0/17 for the Interconnect devices; ge-0/0/20 through ge-0/0/23 for the Director devices; ge-0/0/0 through ge-0/0/7 for the Node devices; and ge-0/1/0 and ge-0/1/1 for the inter-EX4200 switch connections), the control plane is properly connected.

- Related Documentation**
- *QFX3000-M QFabric System Installation Overview*
 - *Installing and Connecting a QFX3100 Director Device*
 - *Installing and Connecting a QFX3600 or QFX3600-I Device*
 - *Installing and Connecting a QFX3500 Device*
 - *Installing and Connecting an EX4200 Switch*
 - [Understanding the QFabric System Control Plane on page 1196](#)

Importing a QFX3000-M QFabric System Control Plane EX4200 Switch Configuration with a USB Flash Drive

There are two methods of importing the configuration file to the QFX3000-M QFabric system control plane EX4200 switches.

- Download the configuration file onto a USB flash drive from the Juniper Networks software download site before inserting the USB flash drive into the EX4200 switch USB port
- Copy and paste the configuration from “[Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane](#)” on page 1309.



NOTE: The EX4200 switch configuration is the same for both the copper-based and fiber-based QFX3000-M QFabric system control plane networks.

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see *Installing and Connecting a QFX3100 Director Device*, *Installing and Connecting a QFX3600 or QFX3600-I Device*, and *Installing and Connecting a QFX3500 Device*.
- Rack, mount, and install your EX4200 switches for the QFabric system control plane. For more information, see *Installing and Connecting an EX4200 Switch*.
- Select a USB flash drive that meets the EX4200 switch USB port specifications. See *USB Port Specifications for an EX Series Switch*.
- Use a computer or other device to download the configuration file from the Internet and copy it to the USB flash drive.

To import the control plane EX4200 switch configuration file onto a USB flash drive:

1. In a browser, go to <http://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms - Download Software page appears.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. In the QFX Series box, select **QFX3000-M QFabric System**.

The QFX3000-M QFabric System - Download Software page appears.

3. Click the **Software** tab and select the software release number from the **Release** list that appears to the right of the Software tab.

A login screen appears.

4. In the QFabric System Control Plane Network section, select **QFX3000-M Control Plane Network Configuration**.

A login screen appears.

5. Enter your user ID and password and click **Login**.
6. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
7. Save the configuration file onto the USB flash drive using your computer or other device.
8. Insert the USB flash drive into the USB port on the EX4200 switch.
9. Save the file to the `/var/home/username` directory on the EX4200 switch.
10. Load the configuration file into the switch.

```
user@switch# load override filename
```

11. Commit the configuration.

```
user@switch# commit
Load complete
```

12. Remove the USB flash drive from the switch.

- Related Documentation**
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane](#) on page 1309

Generating the MAC Address Range for a QFabric System

Each QFabric system requires a range of reserved MAC addresses that is assigned by Juniper Networks. You must specify the MAC address range when you perform the initial setup of the QFX3100 Director group (see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group”](#) on page 1335). Additionally, refer to [Activate Your QFabric System](#) for more information.

When you purchase a QFabric system, you receive an e-mail containing a software serial number from Juniper Networks. You can use the software serial number to generate the MAC address range for your QFabric system.

To generate the MAC address range for a QFabric system:

1. In a browser, log in to the Juniper Networks License Management System at <https://www.juniper.net/lcrs/license.do>.

The Manage Product Licenses page appears.



NOTE: To access the licensing site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. On the Generate Licenses tab, select **QFX Series Product** from the drop-down list, and click **Go**.

The Generate Licenses - QFX Series Product page appears.

3. Select the **QFX Series Product Fabric** option button, and click **Continue**.

The Generate Licenses - QFX Series Product Fabrics page appears.

4. In the **Software Serial No** field, enter the software serial number for your QFabric system, and press the Tab key.

The starting MAC address and number of MAC addresses for your QFabric system are displayed.

5. (Optional) Click **Download/Email MAC Address** to download or e-mail the MAC address range.

The Download/Email MAC Address page appears.

To download the MAC address range:

- Select the **Download to this computer** option button, and click **OK**.

To e-mail the MAC address range:

- Select the **Send e-mail to e-mail ID** option button, and click **OK**.

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)

Performing the QFabric System Initial Setup on a QFX3100 Director Group

You must perform the initial setup of the QFX3100 Director group through the console port. (Before configuring the QFX3100 Director group, see *Installing and Connecting a QFX3100 Director Device*.)

Before you begin connecting and configuring a QFX3100 Director group, set the following parameter values on the console server or PC:

- Baud Rate—9600
- Flow Control—None

- Data—8
- Parity—None
- Stop Bits—1
- DCD State—Disregard



NOTE: When you use the SecureCRT client to connect to a Director device for the initial setup of a QFabric system, the backspace key does not work. As a workaround, use the Shift+Delete key combination in SecureCRT as a backspace key equivalent or use a different UNIX client to support the backspace key natively.

The initial setup requires that you specify certain values for your QFabric system. These include:

- Software serial number for your QFabric system (found in the e-mail containing the software serial number that you received from Juniper Networks when you purchased your QFabric system)
- IP addresses and a default gateway IP address for your QFabric system default partition
- IP addresses for your Director group device management ports
- Range of reserved MAC addresses for your QFabric system (see [“Generating the MAC Address Range for a QFabric System” on page 1334](#) or [Activate Your QFabric System](#) for this information)
- Root password for your Director group
- Root password for the QFabric system components such as the Node devices, Interconnect devices, and infrastructure
- [Performing an Initial Setup on page 1336](#)
- [Restoring a Backup Configuration on page 1339](#)

Performing an Initial Setup

The initial setup can be performed either manually or by using a previously saved backup configuration.

To connect and configure the QFX3100 Director group manually from the console:

1. Connect the console port of one of the Director devices to a laptop or PC using an RJ-45 to DB-9 rollover cable. An RJ-45 to DB-9 rollover cable is supplied with each QFX3100 Director device. The console (**CONSOLE**) port is located on the front panel of the device.
2. Log in as **root**. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

dg0 login: **root**



NOTE: The prompt is either dg0 login or dg1 login depending on the Director device to which you connected your cable.

3. For manual configuration or for initial installation, enter **no** when prompted to specify the backup file. The current Director device configuration is displayed.

Initial Configuration

Before you can access the QFabric system, you must complete the initial setup of the Director group by using the steps that follow. If the initial setup procedure does not complete successfully, log out of the Director device and then log back in to restart this setup menu.

Continue? [y/n]: **y**

You may enter the configuration manually or restore from a backup.

Specify a backup file? [y/n]: **n**

Existing local configuration:

4. Enter the IP addresses and prefixes for both Director devices.



NOTE: The Director group devices and QFabric system default partition IP addresses must be on the same subnet as your management network.

Please enter the Director Group 0 IP address and prefix: *ip address/prefix*

Please enter the Director Group 1 IP address and prefix: *ip address/prefix*

Please enter the Director Group Subnet Mask: *subnet mask*

5. Enter the gateway IP address for the Director group.

Please enter the Director Group gateway IP address: *gateway ip address*

6. Enter the default partition IP address. (You will use this address to log in to the QFabric system on subsequent connections.)

Please enter the QFabric default partition IP address: *ip address*

- 7.



NOTE: On a QFabric system, even though the option to assign IPv6 addresses to a QFabric management network is visible, this feature is not currently supported.

(Optional) Enter the IPv6 addresses for both Director devices and the gateway IPv6 address for the Director group.

Would you like to input IPv6 addresses for Director Group nodes? (y/n): **n**

8. Enter the MAC address information.

Please enter the starting MAC address: *mac address*

Please enter the number of MAC addresses: *number of mac addresses*



NOTE: The minimum number of MAC addresses accepted is 4000.

9. Enter the QFabric system software serial number.

Please enter the QFabric serial ID: *serial id*

10. Create the Director device root password.

Please enter a Director device root password: *director-device-password*

Please re-enter password: *director-device password*

11. Create a password for the QFabric system components.



NOTE: If you need to change the component password after the QFabric system is operational, issue the device-authentication statement at the [edit system] hierarchy level in the QFabric default partition CLI.

Please enter a password for QFabric components (Node devices, Interconnect devices, and infrastructure): *component-password*

Please re-enter password: *component-password*

Note: please record your passwords for recovery purposes.



CAUTION: Carefully save your passwords for future reference, because some cannot be recovered on a QFabric system.

12. Enter the QFabric system platform type.

Supported platform types:

1. QFX3000-G
2. QFX3000-M

Please select product type: *number corresponding to platform type*

13. Confirm the initial configuration. Ensure that the information is accurate before proceeding.

Does the following configuration appear correct?

Director Group 0 IP/Prefix	[10.49.214.74/24]
Director Group 1 IP/Prefix	[10.49.214.75/24]
Director Group Gateway	[10.49.214.254]
Starting MAC address	[00:11:00:00:00:00]
Number of MAC addresses	[4000]
QFabric Default Partition IP	[10.49.214.150]
QFabric serial ID	[qfsn-123456789]
Director Device Password	[*****]
QFabric component Password	[*****]
Product Type:	[QFX3000-G]

14. Confirm the initial setup.

[y/n]: y



CAUTION: Resetting this initial configuration requires assistance from Juniper Networks customer support or [“Performing a QFabric System Recovery Installation on the Director Group” on page 112](#). As a result, make sure you are certain the values you entered are correct before you enter yes.

15. The director device displays the configuration.

Saving temporary configuration...

Configuring peer...

Configuring local interfaces...

Configuring interface eth0 with [10.49.214.74/24:10.49.214.254]

Configured interface eth0 with [10.49.214.74/24:10.49.214.254]

```

Configuring QFabric software with an initial pool of 4000 MAC addresses
[00:11:00:00:00:00 - 00:11:00:00:0f:3b]
Configuring QFabric address [10.49.214.150]
Reconfiguring QFabric software static configuration
Applying the new Director device password
Applying the QFabric component password
First install initial configuration, generating and sharing SSH keys.
First install initial configuration, generating SSH keys.
Configuration complete. Director Group services will auto start within 30
seconds.

```

Restoring a Backup Configuration

Before you restore a backup configuration for the Director group:

- You must have a backup configuration file. You create the backup file with the [request system software configuration-backup](#) command and save it on an external USB flash drive.
- If you need to reinstall the system software, perform that operation first (see [“Performing a QFabric System Recovery Installation on the Director Group” on page 112](#)).

To connect and configure the Director group with a backup configuration:

1. Log in as **root**. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

```
dg0 login: root
```



NOTE: The prompt is either **dg0 login** or **dg1 login** depending on the Director device to which you connected your cable.

2. To use a previously saved backup configuration, enter **yes** when prompted to specify the backup file and then enter the path and filename of the backup configuration.

```
Specify a back up file? [y/n]: y
```

```
Please specify the full path of the configuration backup file: path/filename
```

3. Confirm the restoration of the configuration from the backup. Ensure that the information is accurate before proceeding.

```
Does the following configuration appear correct?
```

```

Director Group 0 IP/Prefix          [10.49.214.74/24]
Director Group 1 IP/Prefix          [10.49.214.75/24]
Director Group Gateway               [10.49.214.254]
Starting MAC address                 [00:11:00:00:00:00]
Number of MAC addresses               [4000]
QFabric Default Partition IP         [10.49.214.150]
QFabric serial ID                    [qfsn-123456789]
Director Device Password             [*****]
QFabric component Password           [*****]
Product Type:                        [QFX3000-G]

```

4. Confirm the backup restoration.

```
[y/n]: y
```

The Director device displays the configuration.

```
Saving temporary configuration...
```

```
Configuring peer...
Configuring local interfaces...
Configuring interface eth0 with [10.49.214.74/24:10.49.214.254]
Configured interface eth0 with [10.49.214.74/24:10.49.214.254]
Configuring QFabric software with an initial pool of 4000 MAC addresses
[00:11:00:00:00:00 - 00:11:00:00:0f:3b]
Configuring QFabric address [10.49.214.150]
Reconfiguring QFabric software static configuration
Applying the new Director device password
Applying the QFabric component password
Configuration complete. Director Group services will auto start within 30
seconds.
```

**Related
Documentation**

- [Generating the MAC Address Range for a QFabric System on page 1334](#)
- [Gaining Access to the QFabric System Through the Default Partition on page 1344](#)
- [QFabric System Initial and Default Configuration Information on page 1255](#)
- [*Installing and Connecting a QFX3100 Director Device*](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)
- [request system software configuration-backup on page 402](#)
- [device-authentication](#)

QFabric System Configuration

- [Understanding QFabric System Administration Tasks and Utilities on page 1340](#)
- [Gaining Access to the QFabric System Through the Default Partition on page 1344](#)
- [Example: Configuring QFabric System Login Classes on page 1345](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)
- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)
- [Example: Configuring SNMP on page 1370](#)
- [Example: Configuring System Log Messages on page 1372](#)
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)

Understanding QFabric System Administration Tasks and Utilities

The following items describe QFabric system components, common administration tasks that you perform on the QFabric system, or utilities that help you to manage the QFabric system and its components.

- **Converting the device mode (QFX3500 and QFX3600 devices)**—Enables you to convert a QFX3500 or QFX3600 device into a Node device so it can be deployed within a QFabric system. By default, QFX3500 and QFX3600 devices operate in *standalone* mode. Before the devices can participate within a QFabric system environment, you must change the device mode for the switch to *node-device* mode. To convert a QFX3500 or QFX3600 device from standalone mode to Node device mode, connect to the console port of the device, issue the **request chassis device-mode node-device** command, verify the future device mode with the **show chassis device-mode** command, connect the management port of the device to the QFabric system control plane, and reboot the device.

**NOTE:**

- Before you convert the device mode, you must upgrade the software on your standalone device to a QFabric system Node and Interconnect device software package that matches the QFabric system complete software package used by your QFabric system. For example, if the complete software package for your QFabric system is named `jinstall-qfabric-11.3X30.6.rpm`, you need to install the `jinstall-qfx-11.3X30.6-domestic-signed.tgz` package on your standalone device. Matching the two software packages ensures a smooth and successful addition of the device to the QFabric system inventory.
 - Converting the device mode erases the switch configuration. We recommend that you save your configuration to an external server or USB flash drive before executing the device mode conversion commands and rebooting the switch.
-
- **QFabric system control plane Ethernet network (EX4200 switches to support the QFabric system)**—Provides a separate control plane network within the QFabric system to handle management traffic. This design enables the data plane network to focus on efficient, low-latency delivery of data, voice, and video traffic.
 - The QFX3000-G QFabric system control plane uses two sets of four EX4200 switches each, configured as a pair of Virtual Chassis to connect all components within the QFabric system. The dual Virtual Chassis architecture provides redundancy and high availability to ensure reliable QFabric system operation for the Director group, the Interconnect devices, and the Node devices.
 - The QFX3000-M QFabric system control plane uses two EX4200 switches to connect all components within the QFabric system. The two EX4200 switches provide redundancy and high availability to ensure reliable QFabric system operation for the Director group, the Interconnect devices, and the Node devices.

Because the level of detail necessary to fully understand the control plane connections, cabling, topology, and configuration is beyond the scope of this topic, see:

- [“Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane” on page 1263](#) for information about a QFX3000-G QFabric system with a copper-based control plane

- [“Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane” on page 1309](#) for information about a QFX3000-M QFabric system with a copper or fiber-based control plane
- **QFabric system data plane network**—Provides a separate network to handle rapid delivery of data plane traffic. The data plane uses QSFP+ interfaces and fiber-optic cabling to connect QFabric system components at speeds of 40 Gbps. By creating a redundant set of connections between the Node devices and the backplane-like Interconnect devices, the data plane enables the Node devices to appear as if they are directly connected to one another in a single tier. To view the connection status of the QFabric system data plane, issue the **show chassis fabric connectivity** command.
- **Director group (QFX3100 Director devices within a QFabric system)**—Provides a redundant, resilient platform that manages the QFabric system components. Two QFX3100 Director devices work together to ensure high availability of the system and load-balance system processes, such as the command-line interface (CLI) and shared storage. To configure the Director group for operation, install and cable two Director devices as a Director group, connect to the console port of one of the Director devices, and perform the initial setup. The setup script starts automatically the first time you power on the Director device. For more information, see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#). To monitor the status of the Director group, log in to the QFabric system default partition and issue the **show fabric administration inventory director-group status** command.
- **Automatic detection and configuration of QFabric system components**—Enables QFabric system components to join the QFabric system automatically. When you install the QFabric system, activate the control plane and Director group, and power on the Node and Interconnect devices, the Director group recognizes these devices, sends each device its own portion of the Junos OS configuration, and adds them to the QFabric system inventory. By default, each individual Node device is placed into a unique server Node group that contains only that single Node device. No configuration is required for the default assignments. The default settings can be overridden when you add Node devices into a redundant server Node group (containing a pair of Node devices) or a network Node group (that can contain up to eight Node devices, run routing protocols, and connect to external networks).
- **QFabric system Routing Engines**—Support the QFabric system by providing virtual, redundant instances of Junos OS that run on the Director group. The Routing Engines perform fabric management tasks, maintain control of the fabric, and host the operation of routing protocols for network Node groups. Because they are generated in pairs, the Routing Engines provide additional high availability for the QFabric system. No configuration is required. To view the status of the QFabric system Routing Engines, issue the **show fabric administration inventory infrastructure** command.
- **QFabric system command-line interface**—Enables you to configure all components of the QFabric system from a single location by using the Junos OS CLI. To access this central location, you need to log in to the QFabric system default partition (an IP address you specify during the initial setup of the Director group). For more information, see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#).

Most existing Junos OS configuration statements and operational mode commands are supported (for example, interfaces, VLANs, protocols, and firewall filters).

To view QFabric system components and check connectivity of the system, issue the **show fabric administration inventory** commands.

- **Alias configuration for Director devices, Interconnect devices, and Node devices**—Enables you to set user-defined aliases for QFabric system Director devices, Interconnect devices, and Node devices to facilitate usability of the QFabric system as it scales. Aliased names appear in the output of many QFabric system operational commands, such as **show fabric administration inventory**. To map the hardware serial number of a Director device, Interconnect device or Node device to a user-defined name, see [“Configuring Aliases for the QFabric System” on page 1353](#).
- **Node group configuration**—Enables you to cluster several Node devices together to provide redundancy, resiliency, and high availability at the ingress and egress points of the QFabric system. There are two types of Node groups you can configure:
 - **Redundant server Node group**—Enables the grouped Node devices to connect the QFabric system to local servers and storage devices. A redundant server Node group can contain a maximum of two Node devices and supports LAG connections that can span both devices.



NOTE: The Node devices in a redundant server Node group must be of the same type, either a QFX3500 Node or a QFX3600 Node. You cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

- **Network Node group**—Enables the grouped Node devices to connect the QFabric system to external networks and run routing protocols such as BGP and OSPF. A network Node group can contain up to eight Node devices and supports LAG connections.



NOTE:

- The name of the network Node group in the default partition, *NW-NG-0*, is preset. You must use this name when adding Node devices to the network Node group. You cannot specify a different name.
- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

To configure a redundant server Node group, include two Node devices with the **node-device node-device-name** statement at the **[edit fabric resources node-group node-group-name]** hierarchy level.

To configure a network Node group, include the **network-domain** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level. In addition, include between

two and eight Node devices with the **node-device *node-device-name*** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level.

**Related
Documentation**

- [Converting the Device Mode for a QFabric System Component on page 1258](#)
- [Example: Configuring the Virtual Chassis for the QFX3000-G QFabric System Control Plane on page 1263](#)
- [Example: Configuring EX4200 Switches for the Copper-Based QFX3000-M QFabric System Control Plane on page 1309](#)
- [show chassis fabric connectivity on page 1520](#)
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [show fabric administration inventory director-group status on page 1548](#)
- [show fabric administration inventory infrastructure on page 1553](#)
- [show fabric administration inventory on page 1543](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)

Gaining Access to the QFabric System Through the Default Partition

This topic explains how to log in to the QFabric system default partition so you can access the Junos OS command-line interface (CLI) and configure the system.

Before you access the QFabric system default partition:

- Install the QFabric system hardware components, including connecting the network and power cables.
- Convert any QFX3500 and QFX3600 standalone devices to *node-device* mode.
- Connect all components to the control plane Ethernet network.
- Turn on the Director group and run the initial setup script. Remember to write down the IP address of the default partition, which must be on the same subnetwork as your management network.

To access the default partition:

1. Open an SSH connection to the QFabric default partition. Use the IP address you set for the default partition as part of the QFabric initial setup procedure. In your network, you can simplify access to the QFabric system by mapping the default partition IP address to a name.

```
[root@customer ~]# ssh root@192.168.1.49
Last login: Fri Sep  2 21:34:54 2011 from customer
Juniper QFabric Director 11.3.5043 2011-08-26 18:05:21 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg1
root@qfabric>
```



NOTE: The QFabric system is load balanced, so the CLI session might be hosted on either Director device DG0 or DG1.

2. Enter configuration mode (the default mode in the QFabric system is **configure private**), configure a root password and hostname for the default partition, and assign QFabric administrator privileges to the root user.

```
root@qfabric> configure
warning: Using private edit on QF/Director
warning: uncommitted changes will be discarded on exit
Entering configuration mode

[edit]

root@qfabric# set system root-authentication plain-text-password
New password: My-Password
Retype new password: My-Password

root@qfabric# set system root-authentication remote-debug-permission qfabric-admin
root@qfabric# set system host-name my-qfabric

[edit]

root@qfabric# commit
commit complete

[edit]
root@my-qfabric#
```

3. Configure your QFabric system as needed. You can configure routing protocols, interfaces, VLANs, and other features as needed. Keep in mind that interfaces require the four-level interface naming convention (*device-name:fpic/pic/port*).

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [QFabric System Initial and Default Configuration Information on page 1255](#)
- [Understanding Interfaces on the QFabric System on page 1173](#)
-

Example: Configuring QFabric System Login Classes

This example shows you how to assign the correct login class to users so they can access components within a QFabric system.

- [Requirements on page 1346](#)
- [Overview on page 1346](#)
- [Configuration on page 1347](#)
- [Verification on page 1349](#)

Requirements

This example uses the following hardware and software components:

- One QFX3000-G QFabric system containing:
 - Two QFX3100 Director devices
 - Two QFX3008-I Interconnect devices
 - Eight QFX3500 Node devices
 - Junos OS Release 12.2 for these QFX Series components
- Eight EX4200 switches, used to make two redundant Virtual Chassis with four members apiece
- Junos OS Release 12.1R1.9 for the EX Series switches used in the Virtual Chassis

Before you begin:

- Perform the initial setup of the QFabric system on the Director group, which includes the creation of a username and password for the QFabric system components. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#).

Overview

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system (such as Node devices and Interconnect devices). The *qfabric-admin* class provides the ability to log in to individual QFabric system components and manage them. The *qfabric-operator* class enables the user to log in to individual components and view component-level operations and configurations. The *qfabric-user* class prevents access to individual QFabric system components.

You include these classes in your configuration at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level. The key task is to decide which class you should apply to users based on their need to access QFabric system components.



NOTE: To set QFabric system login classes for a root user, include the **remote-debug-permission** statement at the **[edit system root-authentication]** hierarchy level and specify the *qfabric-admin* class.

If you assign the *qfabric-admin* or the *qfabric-operator* class to a user, the QFabric system maps the user to a list of authorized users who are permitted to access components. To facilitate ease of use, the QFabric system uses the component password you specified during the initial setup of the Director group. When users assigned the *qfabric-admin* or the *qfabric-operator* class log in to a component by issuing the **request component login** operational mode command, the QFabric system verifies the class and sends the

username and password to the component. The component accepts these credentials and permits access.



NOTE:

- The three QFabric system login classes give access to the components only. To provide access to the QFabric system as a whole through the default partition command-line interface (CLI), you must configure the usual Junos OS login classes or permissions (such as the *super-user* class). For more information about login classes, see [“Junos OS Login Classes Overview” on page 1683](#).
- If you have completed the QFabric system initial setup and the system is operational, you can change the component password by issuing the device-authentication statement at the `[edit system]` hierarchy level in the QFabric default partition CLI.

Topology

This example defines three users: Adam, Oscar, and Ulf. Adam needs to manage QFabric system components, Oscar needs limited access, and Ulf should not have any access to the components. As a result, assign the `qfabric-admin` class to Adam, the `qfabric-operator` class to Oscar, and the `qfabric-user` class to Ulf. However, all three users should have all permissions to access the QFabric system CLI.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

```
set system login class all-qfabric permissions all
set system login user Adam class all-qfabric
set system login user Adam authentication encrypted-password
"$1$aoYSFkve$G/dYqsTV5iSvVW2sND69U."
set system login user Adam authentication remote-debug-permission qfabric-admin
set system login user Oscar class all-qfabric
set system login user Oscar authentication encrypted-password
"$1$3e.3wJQ8$31SrZV0.efdRbk.ZJncKm0"
set system login user Oscar authentication remote-debug-permission qfabric-operator
set system login user Ulf class all-qfabric
set system login user Ulf authentication encrypted-password
"$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"
set system login user Ulf authentication remote-debug-permission qfabric-user
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To provide the same access to the QFabric system CLI for all users, but different QFabric system component-level access to different users:

1. Define and provide all-qfabric access and passwords to all three users. This administrator-defined class provides full permissions, enabling the users to log in to the QFabric system default partition and use the CLI. Alternatively, you can assign the super-user class to these users to accomplish the same goal.

```
[edit]
user@qfabric# set system login class all-qfabric permissions all
user@qfabric# set system login user Adam authentication encrypted-password
"$1$aoYSFkvE$G/dYqsTV5iSvVW2sND69U."
user@qfabric# set system login user Oscar class all-qfabric
user@qfabric# set system login user Oscar authentication encrypted-password
"$1$3e.3wJQ8$31SrV0.efdRbk.ZJncKm0"
user@qfabric# set system login user Ulf class all-qfabric
user@qfabric# set system login user Ulf authentication encrypted-password
"$1$qt9Ncm0o$okNYSN8O4fvITE/SHBdYj0"
```

2. Provide qfabric-admin component access to Adam so he can manage QFabric system components.

```
[edit]
user@qfabric# set system login user Adam authentication remote-debug-permission
qfabric-admin
```

3. Provide qfabric-operator component access to Oscar so he can view the CLI at the QFabric system components.

```
[edit]
user@qfabric# set system login user Oscar authentication remote-debug-permission
qfabric-operator
```

4. Assign qfabric-user component restrictions to Ulf to prevent him from accessing the QFabric system components.

```
[edit]
user@qfabric# set system login user Ulf authentication remote-debug-permission
qfabric-user
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example.

```
[edit]
system {
  login {
    class all-qfabric {
```



```

        permissions all;
    }
    user Adam {
        class all-qfabric;
        authentication {
            encrypted-password "$1$aoYSFkVE$G/dYqsTV5iSvVW2sND69U."; ##
            SECRET-DATA
            remote-debug-permission qfabric-admin;
        }
    }
    user Oscar {
        class all-qfabric;
        authentication {
            encrypted-password "$1$3e.3wJQ8$31SrZV0.efdBk.ZJncKm0"; ## SECRET-DATA
            remote-debug-permission qfabric-operator;
        }
    }
    user Ulf {
        class all-qfabric;
        authentication {
            encrypted-password "$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"; ##
            SECRET-DATA
            remote-debug-permission qfabric-user;
        }
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the QFabric system and component-level access configuration is working properly for all three users. Adam, Oscar, and Ulf should have equivalent, full-permission access to the QFabric system CLI. Adam should have management-level access to components. Oscar should have read-only access to components. Ulf should have no component-level access.

- [Verifying qfabric-admin Access on page 1349](#)
- [Verifying qfabric-operator Access on page 1351](#)
- [Verifying qfabric-user Access on page 1352](#)

Verifying qfabric-admin Access

Purpose Verify that Adam can access the QFabric system CLI at the default partition and manage QFabric system components.

Action From a management station on your network, issue the **ssh user@qfabric** command and enter the password to open an SSH session for Adam to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Adam has permission to use on the QFabric system default partition.

```
> ssh Adam@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
```

```
Adam@qfabric.network.net's password:
Last login: Sun Nov 20 14:12:29 2011 from 192.168.28.19
Juniper QFabric Director 11.3.5510 2011-10-21 16:31:44 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg0
Adam@qfabric>
```

```
Adam@qfabric> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  help           Provide help information
  load           Load information from file
  op             Invoke an operation script
  ping           Ping remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart software process
  save           Save information to file
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  telnet         Telnet to another host
  test           Perform diagnostic debugging
  traceroute     Trace route to remote host
```

Issue the **request component login ?** command to view the components that Adam can access. Next, issue the **request component login *component-name*** command to log in to a Node device without being prompted for a username or password.

```
Adam@qfabric> request component login ?
Possible completions:
  <[Enter]>      Execute this command
  <node-name>    Inventory name for the remote node
  BBAK0372       Node device
  BBAK0394       Node device
  DRE-0          Diagnostic routing engine
  EE3093         Node device
  FC-0           Fabric control
  FC-1           Fabric control
  FM-0           Fabric manager
  NW-NG-0        Node group
  WS001/RE0      Interconnect device control board
  WS001/RE1      Interconnect device control board
  |             Pipe through a command
```

```
Adam@qfabric> request component login EE3093
Warning: Permanently added 'qfnod-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the **?** command to view the CLI operational mode commands that Adam has the permission to use on the Node device. Notice that the CLI prompt now indicates Adam's component access level (**qfabric-admin**) as the username and the Node device identifier (**EE3093**) as the host.

```
qfabric-admin@EE3093> ?
Possible completions:
  clear          Clear information in the system
```

file	Perform file operations
help	Provide help information
load	Load information from file
monitor	Show real-time debugging information
mtrace	Trace multicast path from source to receiver
op	Invoke an operation script
ping	Ping remote target
quit	Exit the management session
request	Make system-level requests
restart	Restart software process
save	Save information to file
set	Set CLI properties, date/time, craft interface message
show	Show system information
ssh	Start secure shell on another host
start	Start shell
telnet	Telnet to another host
test	Perform diagnostic debugging
traceroute	Trace route to remote host

Meaning The output shows that Adam has received the proper permissions to access the QFabric system CLI and log in to individual components with management-level access.

Verifying qfabric-operator Access

Purpose Verify that Oscar can access the QFabric system CLI at the default partition and view the CLI on the QFabric system components.

Action From a management station on your network, issue the **ssh user@qfabric** command and enter the password to open an SSH session for Oscar to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Oscar has permission to use on the QFabric system default partition. Notice that these permissions are the same as those given to Adam.

```
> ssh Oscar@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Oscar@qfabric.network.net's password:
Last login: Sun Nov 19 19:21:29 2011 from 192.168.28.14
Juniper QFabric Director 11.3.5510 2011-10-22 18:33:41 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg1
Oscar@qfabric>
```

```
Oscar@qfabric> ?
Possible completions:
clear          Clear information in the system
configure      Manipulate software configuration information
file           Perform file operations
help           Provide help information
load           Load information from file
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
```

telnet	Telnet to another host
test	Perform diagnostic debugging
tracert	Trace route to remote host

Issue the **request component login *component-name*** command to log in to a Node device without being prompted for a username or password.

```
Oscar@qfabric> request component login EE3093
Warning: Permanently added 'qfnode-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the **?** command to view the CLI operational mode commands that Oscar has permission to use on the Node device. Notice that the CLI prompt now indicates Oscar's component access level (**qfabric-operator**) as the username and the Node device identifier (**EE3093**) as the host. Additionally, Oscar has fewer CLI commands available than Adam because of Oscar's read-only qfabric-operator login class.

```
qfabric-operator@EE3093> ?
Possible completions:
  file      Perform file operations
  help      Provide help information
  load      Load information from file
  op        Invoke an operation script
  quit      Exit the management session
  request   Make system-level requests
  save      Save information to file
  set       Set CLI properties, date/time, craft interface message
  show      Show system information
  start     Start shell
  test      Perform diagnostic debugging
```

Meaning The output shows that Oscar has full permissions to access the QFabric system CLI, but only read-only access when he logs in to individual components. Oscar's permissions on the QFabric system are the same as Adam's, but Oscar has fewer permissions than Adam on the Node device.

Verifying qfabric-user Access

Purpose Verify that Ulf has full access to the QFabric system CLI at the default partition but cannot access the QFabric system components.

Action From a management station on your network, issue the **ssh *user*@qfabric** command and enter the password to open an SSH session for Ulf to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Ulf has permission to use on the QFabric system default partition. Notice that these permissions are the same as those given to Adam and Oscar.

```
> ssh Ulf@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Ulf@qfabric.network.net's password:
Last login: Sun Nov 17 17:12:24 2011 from 192.168.28.22
Juniper QFabric Director 11.3.5510 2011-10-23 19:23:31 UTC
```

RUNNING ON DIRECTOR DEVICE : dg0

```

Ulf@qfabric>

Ulf@qfabric> ?
Possible completions:
clear          Clear information in the system
configure      Manipulate software configuration information
file           Perform file operations
help           Provide help information
load           Load information from file
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
telnet         Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host

```

When Ulf issues the **request component login *component-name*** command, the Node device denies his access attempt.

```

Ulf@qfabric> request component login EE3093
error: User Ulf does not have sufficient permissions to login to device EE3093

```

Meaning The output shows that Ulf has full permissions to access the QFabric system CLI in the same way as Adam and Oscar. However, unlike Adam and Oscar, Ulf cannot access individual components because of the qfabric-user login class assigned to him.

- Related Documentation**
- [Understanding QFabric System Login Classes on page 1230](#)
 - [remote-debug-permission on page 1401](#)
 - [request component login on page 1444](#)
 - [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
 - [Junos OS Login Classes Overview on page 1683](#)

Configuring Aliases for the QFabric System

This topic explains how to configure aliases for components of the QFabric system, such as Director devices, Interconnect devices, and Node devices. Aliases replace the hardware serial numbers of components, making it easier to identify system devices and simplify configuration tasks.

Before you create aliases in a QFabric system:

- Issue one of the **show fabric administration inventory** commands to view the components that are available for aliasing and their hardware serial numbers.



NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

To create an alias for a Node device:

1. Discover the serial number of the Node device you wish to rename by issuing the **set fabric aliases node-device ?** context-sensitive help command.

```
root@qfabric# set fabric aliases node-device ?
Possible completions:
<aliasable-item-name>  The name of the item to be aliased
BBAK8309               Node device
BBAK8283               Node device
BBAK8891               Node device
BBAK8868               Node device
BBAK8276               Node device
BBAK8273               Node device
[edit]
```

As an alternate way to discover the serial number for a Node device, issue the **show fabric administration inventory node-devices** command. In this case, the serial numbers for the Node devices are **BBAK8309BBAK8283BBAK8891BBAK8868BBAK8276** and **BBAK8273**.

```
root@qfabric> show fabric administration inventory node-devices
```

Item	Identifier	Connection	Configuration
Node device			
BBAK8309		Connected	
BBAK8283		Connected	
BBAK8891		Connected	
BBAK8868		Connected	
BBAK8276		Connected	
BBAK8273		Connected	

2. Specify the serial number of the Node device and the desired alias name by including the **node-device** statement at the **[edit fabric aliases]** hierarchy level.

```
[edit fabric aliases]
root@qfabric# set node-device BBAK8309 Node0
root@qfabric# set node-device BBAK8283 Node1
root@qfabric# set node-device BBAK8891 Node2
root@qfabric# set node-device BBAK8868 Node3
root@qfabric# set node-device BBAK8276 Node4
root@qfabric# set node-device BBAK8273 Node5
```

3. Review your configuration and issue the **commit** command.

```
[edit]
root@qfabric# show fabric
aliases {
    node-device BBAK8309 {
        Node0;
    }
    node-device BBAK8283 {
        Node1;
    }
    node-device BBAK8891 {
        Node2;
    }
    node-device BBAK8868 {
        Node3;
    }
    node-device BBAK8276 {
        Node4;
    }
    node-device BBAK8273 {
        Node5;
    }
}

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

4. To view that your aliases are operational, issue the **show fabric administration inventory node-devices** command.

```
root@qfabric> show fabric administration inventory node-devices
```

Item	Identifier	Connection	Configuration
Node device			
node0	BBAK8309	Connected	
node1	BBAK8283	Connected	
node2	BBAK8891	Connected	
node3	BBAK8868	Connected	
node4	BBAK8276	Connected	
node5	BBAK8273	Connected	



NOTE: If you attempt to commit all configuration settings for a new Node group (such as the Node group itself, aliasing, and other features) at the same time, the commit operation might appear to succeed when it actually has failed. For this reason, we recommend configuring and verifying Node groups and aliases first, followed by configuring and verifying other features. Establishing the Node groups and aliases first enables the QFabric system to reject any potentially unsupported configuration. The resulting commit errors indicate where the configuration problem lies. To verify the establishment of Node groups and aliases before configuring other features, issue the **show fabric administration inventory** command.

To create an alias for a Node group:

- Specify a name for the Node group when you include the **node-group** statement at the **[edit fabric resources]** hierarchy level.



NOTE: You cannot use the **aliases** statement at the **[edit fabric]** hierarchy level to create an aliased name for a Node group.

To create an alias for a Director device:

1. Discover the serial number of the Director device you wish to rename by issuing the **set fabric aliases director-device ?** context-sensitive help command.

```

root@qfabric# set fabric aliases director-device ?
Possible completions:
  <aliasable-item-name>  The name of the item to be aliased
    0281052011000001      Director device
    0281052011000032      Director device
[edit]

```

As an alternate way to discover the serial number for a Director device, issue the **show fabric administration inventory director-group status** command. In this case, the serial number for Director device DG0 is **0281052011000001** and the serial number for Director device DG1 is **0281052011000032**.

```

root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 15:11:26 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online
Partition Manager	online

Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master

FTP Server	online
Syslog	online
Distributed Management	online
SNMP Trap Forwarder	online
SNMP Process	online
Platform Management	online

Interface Link Status

Management Interface	up
Control Plane Bridge	up
Control Plane LAG	up
CP Link [0/2]	up
CP Link [0/1]	up
CP Link [0/0]	up
CP Link [1/2]	down
CP Link [1/1]	down
CP Link [1/0]	down
Crossover LAG	up
CP Link [0/3]	up
CP Link [1/3]	up

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	backup

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:8	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.5G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online
Partition Manager	online
Software Mirroring	online
Shared File System master	online
Secure Shell Process	online

```

Network File System          online
DHCP Server master          online    backup

FTP Server                   online
Syslog                       online
Distributed Management       online
SNMP Trap Forwarder         online
SNMP Process                 online
Platform Management         online

Interface Link Status
-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```

- Specify the serial number of the Director device and the desired alias name by including the **director-device** statement at the **[edit fabric aliases]** hierarchy level.

```

[edit fabric aliases]
root@qfabric# set director-device 0281052011000001 Director0
root@qfabric# set director-device 0281052011000032 Director1

```

- Review your configuration and issue the **commit** command.

```

[edit]
root@qfabric# show fabric
aliases {
    director-device 0281052011000001 {
        Director0;
    }
    director-device 0281052011000032 {
        Director1;
    }
}

[edit]
root@qfabric# commit
commit complete

```

- To view that your aliases are operational, issue the **show fabric administration inventory director-group status** command. In this case, the serial numbers in the **Device Id/Alias** field have been replaced with the **Director0** and **Director1** aliased names.

```

root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 15:11:26 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
-----	-----	-----	-----

```

dg0    Director0          online master

Master Services
-----
Database Server           online
Load Balancer Director    online
QFabric Partition Address online

Director Group Managed Services
-----
Shared File System        online
Network File System       online
Virtual Machine Server    online
Load Balancer/DHCP        online

Hard Drive Status
-----
Volume ID:4               optimal
Physical ID:1             online
Physical ID:0             online
SCSI ID:1                 100%
SCSI ID:0                 100%

Size  Used Avail Used% Mounted on
----  -
423G  5.4G 395G   2%  /
99M   16M  79M   17% /boot
93G   7.3G 86G    8%  /pbdata

Director Group Processes
-----
Director Group Manager    online
Partition Manager         online
Software Mirroring        online
Shared File System master online
Secure Shell Process      online
Network File System       online
DHCP Server master        online    master
FTP Server                online
Syslog                   online
Distributed Management     online
SNMP Trap Forwarder       online
SNMP Process              online
Platform Management       online

Interface Link Status
-----
Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG             up
CP Link [0/3]             up
CP Link [1/3]             up

```

```

Member Device Id/Alias  Status  Role
-----
dg1    Director1         online  backup

Director Group Managed Services
-----
Shared File System      online
Network File System     online
Virtual Machine Server  online
Load Balancer/DHCP      online

Hard Drive Status
-----
Volume ID:8             optimal
Physical ID:1           online
Physical ID:0           online
SCSI ID:1               100%
SCSI ID:0               100%

Size  Used Avail Used% Mounted on
----  -
423G  5.5G 395G   2%  /
99M   16M  79M   17% /boot
93G   7.3G 86G    8%  /pbdata

Director Group Processes
-----
Director Group Manager  online
Partition Manager       online
Software Mirroring       online
Shared File System master online
Secure Shell Process    online
Network File System     online
DHCP Server master      online  backup
FTP Server               online
Syslog                   online
Distributed Management  online
SNMP Trap Forwarder     online
SNMP Process             online
Platform Management     online

Interface Link Status
-----
Management Interface    up
Control Plane Bridge    up
Control Plane LAG       up
CP Link [0/2]           up
CP Link [0/1]           up
CP Link [0/0]           up
CP Link [1/2]           down
CP Link [1/1]           down
CP Link [1/0]           down
Crossover LAG           up
CP Link [0/3]           up
CP Link [1/3]           up

```

To create an alias for an Interconnect device:

1. Discover the serial number of the Interconnect device you wish to rename by issuing the **set fabric aliases interconnect-device ?** context-sensitive help command.

```
root@qfabric# set fabric aliases interconnect-device ?
Possible completions:
  <aliasable-item-name>  The name of the item to be aliased
  IC-F1249               Interconnect device
  IC-F4912               Interconnect device
[edit]
```

As an alternate way to discover the serial number for an Interconnect device, issue the **show fabric administration inventory interconnect-devices** command. In this case, the serial numbers for the Interconnect devices are **IC-F1249** and **IC-F4912**.

```
root@qfabric> show fabric administration inventory interconnect-devices
```

Item	Identifier	Connection	Configuration
Interconnect device			
IC-F1249		Connected	Configured
F1249/RE0		Connected	
IC-F4912		Connected	Configured
F4912/RE0		Connected	

2. Specify the serial number of the Interconnect device and the desired alias name by including the **interconnect-device** statement at the **[edit fabric aliases]** hierarchy level.

```
[edit fabric aliases]
root@qfabric# set interconnect-device IC-F1249 Interconnect0
root@qfabric# set interconnect-device IC-F4912 Interconnect1
```

3. Review your configuration and issue the **commit** command.

```
[edit]
root@qfabric# show fabric
aliases {
  interconnect-device IC-F1249 {
    Interconnect0;
  }
  interconnect-device IC-F4912 {
    Interconnect1;
  }
}

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

4. To view that your aliases are operational, issue the **show fabric administration inventory interconnect-devices** command.

```
root@qfabric> show fabric administration inventory interconnect-devices
```

Item	Identifier	Connection	Configuration
Interconnect device			
Interconnect0	IC-F1249	Connected	Configured
F1249/RE0		Connected	
Interconnect1	IC-F4912	Connected	Configured
F4912/RE0		Connected	

- Related Documentation
- [aliases on page 1380](#)
 - [show fabric administration inventory on page 1543](#)
 - [show fabric administration inventory director-group status on page 1548](#)
 - [show fabric administration inventory interconnect-devices on page 1556](#)
 - [show fabric administration inventory node-devices on page 1558](#)
 - [Understanding the Director Group on page 1180](#)
 - [Understanding Interconnect Devices on page 1183](#)
 - [Understanding Node Devices on page 1186](#)

Configuring Node Groups for the QFabric System

This topic explains how to configure Node groups for Node devices within the QFabric system. Node groups provide redundancy for Node devices and make your QFabric system more resilient.

There are three types of Node groups in a QFabric system:

- **Automatically generated server Node groups**—By default, every Node device that joins the QFabric system is placed within an automatically generated server Node group that contains one Node device (the device itself). Server Node groups connect to servers and storage devices.
- **Network Node groups**—You can assign up to eight Node devices to a network Node group. When grouped together, the Node devices within a network Node group connect to other routers running routing protocols such as OSPF and BGP.
- **Redundant server Node groups**—You can assign two Node devices to a redundant server Node group. When grouped together, you can create link aggregation groups (LAGs) that span the interfaces on both Node devices to provide resiliency and redundancy.

Before you create Node groups in a QFabric system:

- Make sure your QFabric system is operational.
- Issue the **show fabric administration inventory node-devices** command to display the Node devices that are available to add to a Node group.
- Issue the **show fabric administration inventory node-groups** command to display the existing Node groups.



NOTE: The following rules apply to QFabric Node group naming:

- Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of a Node group name is 30 characters.
- Node group names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as a Node group name.



NOTE: If you attempt to commit all configuration settings for a new Node group (such as the Node group itself, aliasing, and other features) at the same time, the commit operation might appear to succeed when it actually has failed. For this reason, we recommend configuring and verifying Node groups and aliases first, followed by configuring and verifying other features. Establishing the Node groups and aliases first enables the QFabric system to reject any potentially unsupported configuration. The resulting commit errors indicate where the configuration problem lies. To verify the establishment of Node groups and aliases before configuring other features, issue the **show fabric administration inventory command**.

To display an automatically generated server Node group:

- Issue the **show fabric administration inventory node-groups** command and look for Node groups containing a single Node device that has the same name or serial number as the server Node group.

```
root@qfabric> show fabric administration inventory node-groups
```

Item	Identifier	Connection	Configuration
Node group			
BBAK8281		Connected	Configured
BBAK8281		Connected	
BBAK8835		Connected	Configured
BBAK8835		Connected	
NW-NG-0		Connected	Configured
Node0	BBAK8309	Connected	
Node1	BBAK8283	Connected	
S1		Connected	Configured
Node2	BBAK8891	Connected	
Node3	BBAK8868	Connected	

To create a network Node group:

1. Specify the Node devices you wish to add to the network Node group by including the **node-device** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level.



NOTE:

- The network Node group must use the predefined name NW-NG-0. You must use this name when adding Node devices to the network Node group. You cannot specify a different name. Also, you can configure only one network Node group per partition.
- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

[edit]

```
root@qfabric# set fabric resources node-group NW-NG-0 node-device Node0
root@qfabric# set fabric resources node-group NW-NG-0 node-device Node1
```

2. To designate the Node group as a network Node group, include the **network-domain** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level.

[edit]

```
root@qfabric# set fabric resources node-group NW-NG-0 network-domain
```

3. Review your configuration and issue the **commit** command.

[edit]

```
root@qfabric# show fabric
resources {
  node-group NW-NG-0 {
    network-domain;
    node-device Node0;
    node-device Node1;
  }
}
```

[edit]

```
root@qfabric# commit
commit complete
```



NOTE: When you add or delete Node devices from a Node group configuration, the corresponding Node devices reboot when you commit the configuration change.

4. To determine if your network Node group is operational, issue the **show fabric administration inventory node-groups** command in operational mode.

```
root@qfabric>show fabric administration inventory node-groups NW-NG-0
Item          Identifier          Connection          Configuration
Node group
```

NW-NG-0		Connected	Configured
Node0	BBAK8309	Connected	
Node1	BBAK8283	Connected	

To create a redundant server Node group:

1. Specify the two Node devices you wish to add to the redundant server Node group by including the **node-device** statement at the **[edit fabric resources node-group node-group-name]** hierarchy level.



NOTE: Ensure that the two Node devices are of the same type, either a QFX3500 Node or a QFX3600 Node. You cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

[edit]

```
root@qfabric# set fabric resources node-group S1 node-device Node2
root@qfabric# set fabric resources node-group S1 node-device Node3
```

2. Review your configuration and issue the **commit** command.

[edit]

```
root@qfabric# show fabric
resources {
  node-group S1 {
    node-device Node2;
    node-device Node3;
  }
}
```

[edit]

```
root@qfabric# commit
commit complete
```



NOTE: When you add or delete Node devices from a Node group configuration, the corresponding Node devices reboot when you commit the configuration change.

3. To determine if your redundant server Node groups are operational, issue the **show fabric administration inventory node-groups redundant-server-node-group-name** command in operational mode.

```
root@qfabric> show fabric administration inventory node-groups S1
```

Item	Identifier	Connection	Configuration
Node group			
S1		Connected	Configured
Node2	BBAK8891	Connected	
Node3	BBAK8868	Connected	

Related Documentation

- [show fabric administration inventory node-groups on page 1560](#)
- [show fabric administration inventory node-devices on page 1558](#)
- [Understanding Node Groups on page 1190](#)

- [node-group \(Resources\)](#) on page 1399

Configuring the Port Type on QFX3600 Node Devices

The QFX3600 Node device provides 16 40-Gbps QSFP+ ports. By default, four ports (labeled **Q0** through **Q3**) operate as 40-gigabit data plane (*fte*) uplink ports for uplink connections between your Node device and your Interconnect devices. Twelve ports (labeled **Q4** through **Q15**) operate as 10-Gigabit Ethernet (*xe*) ports to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure ports Q0 through Q7 to operate as 40-gigabit data plane uplink ports, and ports Q2 through Q15 to operate as 10-Gigabit Ethernet ports.



NOTE: You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other devices.



NOTE: When you delete the port type configuration for an individual port or a block of ports, the ports return to operating in their default port type.

This topic explains how to configure the port type on QFX3600 Node devices.

Before you configure the port type on QFX3600 Node devices:

- Make sure your QFabric system is operational.
- Issue the **show fabric administration inventory node-groups** command to display the existing Node groups and the Node devices in each Node group.



NOTE:

- Only ports Q0 through Q7 can be configured to operate as 40-gigabit data plane (*fte*) uplink ports.
- Only ports Q2 through Q15 can be configured to operate as 10-Gigabit Ethernet (*xe*) ports.



CAUTION: The Packet Forwarding Engine on the QFX3600 Node device is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the Node device.

The following message may be displayed in the system log file when the Packet Forwarding Engine is restarted. You can ignore this message.

Pipe write error: Broken pipe

flush operation failed

The following steps describe how to configure either a block of ports or an individual port to operate as 40-gigabit data plane uplink (fte) ports, as well as how to delete a 40-gigabit data plane uplink (fte) port configuration.

1. To configure a block of ports to operate as 40-gigabit data plane uplink (fte) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set fte port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 40-gigabit data plane uplink ports:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set fte port-range 4 7
```

2. To configure an individual port to operate as a 40-gigabit data plane uplink (fte) port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set fte port port-number
```

For example, to configure port Q4 to operate as a 40-gigabit data plane uplink port:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set fte port 4
```

3. Review your configuration and issue the **commit** command.

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

4. To delete the 40-gigabit data plane uplink (fte) port configuration for a block of ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# delete fte port-range port-range-low port-range-high
```

For example, to delete the 40-gigabit data plane uplink port configuration for ports Q4 through Q7:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# delete fte port-range 4 7
```

5. To delete the 40-gigabit data plane uplink (fte) port configuration for an individual port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# delete fte port port-number
```

For example, to delete the 40-gigabit data plane uplink port configuration for port Q4:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# delete fte port 4
```

The following steps describe how to configure either a block of ports or an individual port to operate as 10-Gigabit Ethernet (xe) ports, as well as how to delete a 10-Gigabit Ethernet (xe) port configuration.

1. To configure a block of ports to operate as 10-Gigabit Ethernet (xe) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# set xe port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# set xe port-range 4 7
```

2. To configure an individual port to operate as a 10-Gigabit Ethernet port, specify a port number:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# set xe port port-number
```

For example, to configure port Q4 to operate as a 10-Gigabit Ethernet port:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# set xe port 4
```

3. Review your configuration and issue the **commit** command.

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

4. To delete the 10-Gigabit Ethernet (xe) port configuration for a block of ports, specify a port range:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# delete xe port-range port-range-low port-range-high
```

For example, to delete the 10-Gigabit Ethernet port configuration for ports Q4 through Q7:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# delete xe port-range 4 7
```

5. To delete the 10-Gigabit Ethernet (xe) port configuration for an individual port, specify a port number:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# delete xe port port-number
```

For example, to delete the 10-Gigabit Ethernet port configuration for port Q4:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# delete xe port 4
```

Related Documentation

- [Understanding Node Devices on page 1186](#)
- [Understanding Interfaces on the QFabric System on page 1173](#)
- [pic on page 1400](#)

Example: Configuring SNMP

By default, SNMP is disabled on devices running Junos OS. This example describes the steps for configuring SNMP on the QFabric system.

- [Requirements on page 1370](#)
- [Overview on page 1370](#)
- [Configuration on page 1370](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- Network management system (NMS) (running the SNMP manager)
- QFabric system (running the SNMP agent) with multiple Node devices

Overview

Because SNMP is disabled by default on devices running Junos OS, you must enable SNMP on your device by including configuration statements at the **[edit snmp]** hierarchy level. At a minimum, you must configure the **community public** statement. The community defined as public grants read-only access to MIB data to any client.

If no **clients** statement is configured, all clients are allowed. We recommend that you always include the **restrict** option to limit SNMP client access to the switch.

The network topology in this example includes an NMS, a QFabric system with four Node devices, and external SNMP servers that are configured for receiving traps.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set snmp name "snmp qfabric" description "qfabric0 switch"
set snmp location "Lab 4 Row 11" contact "qfabric-admin@qfabric0"
set snmp community public authorization read-only
set snmp client-list list0 192.168.0.0/24
set snmp community public client-list-name list0
set snmp community public clients 192.170.0.0/24 restrict
set snmp trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure SNMP on the QFabric system:



NOTE: If the name, description, location, contact, or community name contains spaces, enclose the text in quotation marks (" ").

1. Configure the SNMP system name:

```
[edit snmp]
user@switch# set name "snmp qfabric"
```

2. Specify a description.

```
[edit snmp]
user@switch# set description "qfabric0 system"
```

This string is placed into the MIB II sysDescription object.

3. Specify the physical location of the QFabric system.

```
[edit snmp]
user@switch# set location "Lab 4 Row 11"
```

This string is placed into the MIB II sysLocation object.

4. Specify an administrative contact for the SNMP system.

```
[edit snmp]
user@switch# set contact "qfabric-admin@qfabric0"
```

This name is placed into the MIB II sysContact object.

5. Specify a unique SNMP community name and the read-only authorization level.



NOTE: The read-write option is not supported on the QFabric system.

```
[edit snmp]
user@switch# set community public authorization read-only
```

6. Create a client list with a set of IP addresses that can use the SNMP community.

```
[edit snmp]
user@switch# set client-list list0 192.168.0.0/24
user@switch# set community public client-list-name list0
```

7. Specify IP addresses of clients that are restricted from using the community.

```
[edit snmp]
user@switch# set community public clients 192.170.0.0/24 restrict
```

8. Configure a trap group, destination port, and a target to receive the SNMP traps in the trap group.

```
[edit snmp]
user@switch# set trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```



NOTE: You do not need to include the `destination-port` statement if you use the default port 162.

The trap group `qf-traps` is configured to send traps to 192.168.0.100.

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show
snmp {
  name "snmp qfabric";
  description "qfabric0 system";
  location "Lab 4 Row 11";
  contact "qfabric-admin@qfabric0";
  client-list list0 {
    192.168.0.0/24;
  }
  community public {
    authorization read-only;
    clients {
      197.170.0.0/24 restrict;
    }
  }
  trap-group qf-traps {
    destination-port 155;
    targets {
      192.168.0.100;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
 - [snmp on page 1810](#)

Example: Configuring System Log Messages

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

This example describes how to configure system log messages on the QFabric system.

- [Requirements on page 1373](#)
- [Overview on page 1373](#)
- [Configuration on page 1373](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.1
- QFabric system
- External servers that can be configured as system log message hosts

Overview

Component devices that generate system log message events may include Node devices, Interconnect devices, Director devices, and the control plane switches. The following configuration example includes these components in the QFabric system:

- Director software running on the Director group
- Control plane switches
- Interconnect device
- Multiple Node devices

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog host 10.1.1.12 any error
set system syslog file qflogs structured-data brief
set system syslog file qflogs archive size 1g
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure system messages from the QFabric Director device:

1. Specify a host, any facility, and the **error** severity level.


```
[edit system syslog]
user@switch# set host 10.1.1.12 any error
```



NOTE: You can configure more than one system log message server (host). The QFabric system sends the messages to each server configured.

2. (Optional) Specify a filename to capture log messages.



NOTE: On the QFabric system, a syslog file named `messages` is configured implicitly with facility and severity levels of `any any` and a file size of 100 MBs.

[edit system syslog]

user@switch# **set file qflogs structured-data brief**

3. (Optional) Configure the maximum size of your system log message archive file. This example specifies an archive size of 1 GB.

[edit system syslog]

user@switch# **set file qflogs archive size 1g**

Results From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

[edit]

user@switch# **show system**

```
syslog {
  file qflogs {
    archive {
      size 1g;
    }
    structured-data {
      brief;
    }
  }
  host 10.1.1.12 {
    any error;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related
Documentation**

- *Overview of QFabric System Logs*
- [syslog \(QFabric System\) on page 1403](#)

Configuring Graceful Restart for QFabric Systems

When you configure graceful restart in the QFabric CLI, the QFabric system applies the configuration to the network Node group to participate in graceful restart operations with devices external to the QFabric system. Such configuration preserves routing table state and helps neighboring routing devices to resume routing operations more quickly after a system restart. This also enables the network Node group to resume routing operations rapidly if there is a restart in the QFabric system (such as a software upgrade). As a result, we recommend enabling graceful restart for routing protocols in the QFabric CLI.



NOTE: The QFabric system also uses graceful restart internally within the fabric to facilitate interfabric resiliency and recovery. This internal feature is enabled by default with no configuration required.

- [Enabling Graceful Restart on page 1375](#)
- [Configuring Graceful Restart Options for BGP on page 1376](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 1377](#)
- [Tracking Graceful Restart Events on page 1378](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {
  graceful-restart;
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group group-name graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group group-name neighbor ip-address graceful-restart]** hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
graceful-restart {
  helper-disable <both | restart-signaling | standard>
}
```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.

**NOTE:**

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the [edit protocols ospfv3 graceful-restart] hierarchy level.



TIP: You can also track graceful restart events with the traceoptions statement at the [edit protocols (ospf | ospf3)] hierarchy level. For more information, see “Tracking Graceful Restart Events” on page 1378.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the [edit protocols *protocol* traceoptions flag] hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

**Related
Documentation**

- [Graceful Restart Concepts on page 2269](#)
- [Verifying Graceful Restart Operation on page 2343](#)

Configuration Statements

- [aliases on page 1380](#)
- [archive \(QFabric System\) on page 1381](#)
- [chassis \(QFabric System\) on page 1382](#)
- [director-device \(Aliases\) on page 1384](#)
- [fabric on page 1385](#)
- [file \(QFabric System\) on page 1386](#)

- [graceful-restart \(Enabling Globally\) on page 1387](#)
- [graceful-restart \(Protocols BGP\) on page 1388](#)
- [graceful-restart \(Protocols OSPF\) on page 1389](#)
- [interconnect-device \(Chassis\) on page 1391](#)
- [interconnect-device \(Aliases\) on page 1392](#)
- [multicast \(QFabric Routing Options\) on page 1393](#)
- [network-domain on page 1393](#)
- [no-make-before-break on page 1394](#)
- [node-device \(Aliases\) on page 1395](#)
- [node-device \(Chassis\) on page 1396](#)
- [node-device \(Resources\) on page 1397](#)
- [node-group \(Chassis\) on page 1398](#)
- [node-group \(Resources\) on page 1399](#)
- [pic \(Port\) on page 1400](#)
- [remote-debug-permission on page 1401](#)
- [resources on page 1402](#)
- [routing-options \(QFabric System\) on page 1402](#)
- [syslog \(QFabric System\) on page 1403](#)

aliases

Syntax aliases {
 director-device *director-device-name* {
 assigned-director-device-name;
 }
 interconnect-device *interconnect-device-name* {
 assigned-interconnect-device-name;
 }
 node-device *node-device-name* {
 assigned-node-device-name;
 }

Hierarchy Level [edit fabric]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
Options **director-group** and **interconnect-device** introduced in Junos OS Release 12.2 for the QFX Series.

Description (QFabric systems only) Specify the mapping of user defined names to QFabric system components. You can reassign names for Director devices, Interconnect devices, and Node devices.

The remaining statements are explained separately.



NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
 - The maximum length of an alias name is 30 characters.
 - Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
 - You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.
-

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation • [Configuring Aliases for the QFabric System on page 1353](#)
 • [Understanding Node Devices on page 1186](#)

archive (QFabric System)

Syntax	<pre>archive { size size; }</pre>
Hierarchy Level	[edit system syslog file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the archiving properties for the system message log file.
Options	<p>size <i>size</i>—Maximum amount of system log message data that the QFabric system stores in the log file.</p> <p>Syntax: <i>xk</i> to specify the number of kilobytes, <i>xm</i> for the number of megabytes, or <i>xg</i> for the number of gigabytes</p> <p>Range: 65 KB through 1 GB</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• syslog on page 1403

chassis (QFabric System)

```
Syntax  chassis {
        interconnect-device {
            alarm {
                (ethernet | management-ethernet) {
                    link-down (red | yellow | ignore);
                }
            }
            container-devices {
                device-count number;
            }
            craft-lockout {
                alarm {
                    interface-type {
                        link-down (red | yellow | ignore);
                    }
                }
            }
            container-devices {
                device-count number;
            }
            fpc slot {
                power (on | off);
            }
            routing-engine {
                on-disk-failure {
                    disk-failure-action (halt | reboot);
                }
            }
        }
        fpc slot {
            power (on | off);
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
    node-group name {
        aggregated-devices {
            ethernet {
                device-count number;
            }
        }
        alarm {
            interface-type {
                link-down (ignore | red | yellow);
            }
        }
        container-devices {
            device-count number;
        }
        node-device name {
```

```

fibres-channel {
  port-range {
    port-range-low port-range-high;
  }
}
pic pic-number {
  fte {
    port port-number;
    port-range port-range-low port-range-high;
  }
  xe {
    port port-number;
    port-range port-range-low port-range-high;
  }
}
}
routing-engine {
  on-disk-failure {
    disk-failure-action (halt | reboot);
  }
}
}
}

```

Hierarchy Level [\[edit\]](#)


Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

director-device (Aliases)

Syntax	<code>director-device <i>director-device-name</i> { <i>assigned-director-device-name</i>; }</code>
Hierarchy Level	[edit fabric aliases]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Specify the mapping of user-defined names to QFabric system Director devices.
<hr/>	
<div> NOTE: The following rules apply to QFabric component alias naming:</div> <ul style="list-style-type: none">• Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.• The maximum length of an alias name is 30 characters.• Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.• You cannot use the reserved names <code>all</code>, <code>fabric</code>, or <code>director-group</code> as an alias name. <hr/>	
Options	<p><i>director-device-name</i>—Specify a user-defined name for a QFabric system Director device.</p> <p><i>assigned-director-device-name</i>—Specify the Director device identifier or name that has been provided. Identifiers are usually auto-generated by the Director software, and names are usually provided by the administrator of the default partition.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Aliases for the QFabric System on page 1353• Understanding the Director Group on page 1180

fabric

```

Syntax  fabric
        aliases {
            director-device director-device-name {
                assigned-director-device-name;
            }
            interconnect-device interconnect-device-name {
                assigned-interconnect-device-name;
            }
            node-device node-device-name {
                assigned-node-device-name;
            }
        }
        resources {
            node-group node-group-name {
                node-device node-device-name;
                network-domain;
            }
        }
        routing-options {
            multicast
                fabric-optimized-distribution;
            no-make-before-break;
        }
    }

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description (QFabric systems only) Define resources, routing options, and the mapping of user-defined names to QFabric system components.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [show fabric administration inventory on page 1543](#)
- [Understanding QFabric System Administration Tasks and Utilities on page 1340](#)

file (QFabric System)

Syntax file *filename* {
 archive {
 size *maximum-file-size*;
 }
 explicit-priority;
 facility *severity*;
 match "*regular-expression*";
 structured-data {
 brief;
 }
 }

Hierarchy Level [edit system [syslog](#)]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the logging of system messages to a file.



NOTE: On the QFabric system, a syslog file named **messages** is configured implicitly with facility and severity levels of **any any** and a size of 100 MBs.

Options *facility*—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements.

filename—Filename that you specify with the **show log** command.

Default: Filename **messages**

severity—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities at the specified level and higher are logged.

The remaining statements are explained separately.


Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation • [syslog on page 1403](#)

graceful-restart (Enabling Globally)

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.</p> <p>For VPNs, the graceful-restart statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</p> <p>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p>
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Graceful Restart</i> • Configuring Routing Protocols Graceful Restart on page 2285 • <i>Configuring Graceful Restart for MPLS-Related Protocols</i> • <i>Configuring VPN Graceful Restart</i> • <i>Configuring Logical System Graceful Restart</i> • <i>Graceful Restart Configuration Statements</i> • Configuring Graceful Restart for QFabric Systems on page 1375

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time seconds; stale-routes-time seconds; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div>  <p>NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2285 • Configuring Graceful Restart for QFabric Systems on page 1375 • <i>Junos OS High Availability Configuration Guide</i>

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....
notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Graceful Restart for OSPF on page 3894• Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 3898• Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 3901• Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 3904• Configuring Graceful Restart for QFabric Systems on page 1375• <i>Junos OS High Availability Configuration Guide</i> |
|------------------------------|---|

interconnect-device (Chassis)

```
Syntax interconnect-device {
    alarm {
        (ethernet | management-ethernet) {
            link-down (red | yellow | ignore);
        }
    }
    container-devices {
        device-count number;
    }
    craft-lockout {
        alarm {
            interface-type {
                link-down (red | yellow | ignore);
            }
        }
        container-devices {
            device-count number;
        }
        fpc slot {
            power (on | off);
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
    fpc slot {
        power (on | off);
    }
    routing-engine {
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
    }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure properties specific to a QFabric system Interconnect device.


The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Understanding Interconnect Devices on page 1183](#)

interconnect-device (Aliases)

Syntax	<code>interconnect-device <i>interconnect-device-name</i> { <i>assigned-interconnect-device-name</i>; }</code>
Hierarchy Level	[edit fabric aliases]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Specify the mapping of user-defined names to QFabric system Interconnect devices.
<div> NOTE: The following rules apply to QFabric component alias naming:</div> <ul style="list-style-type: none">• Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.• The maximum length of an alias name is 30 characters.• Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.• You cannot use the reserved names <code>all</code>, <code>fabric</code>, or <code>director-group</code> as an alias name.	
Options	<p><i>interconnect-device-name</i>—Specify a user-defined name for a QFabric system Interconnect device.</p> <p><i>assigned-interconnect-device-name</i>—Specify the Interconnect device identifier or name that has been provided. Identifiers are usually auto-generated by the Interconnect software, and names are usually provided by the administrator of the default partition.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Aliases for the QFabric System on page 1353• Understanding Interconnect Devices on page 1183

multicast (QFabric Routing Options)

Syntax	<pre>multicast { fabric-optimized-distribution no-make-before-break; }</pre>
Hierarchy Level	[edit fabric routing-options]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Set multicast routing options in a QFabric system, such as no-make-before-break .
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding the QFabric System Data Plane on page 1199


network-domain

Syntax	network-domain;
Hierarchy Level	[edit fabric resources node-group <i>node-group-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Designate a Node group to become a <i>network Node group</i> , which is used to route traffic between a QFabric system and external networks. The absence of the network-domain configuration statement implies that the Node group is a <i>server Node group</i> , which is used to group sets of Node devices that are connected to servers or storage devices.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Node Groups for the QFabric System on page 1363 • Understanding Node Groups on page 1190

no-make-before-break

Syntax	no-make-before-break;
Hierarchy Level	[edit fabric routing-options multicast]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Disable the default <i>make-before-break</i> multicast feature on QFabric systems to prevent the duplication of system traffic. This feature increases the speed of data plane traffic, but carries the risk of minor traffic losses when compared with the default make-before-break method of creating new multicast fabric paths before tearing down the old paths. The absence of the no-make-before-break configuration statement implies that the make-before-break default system behavior is in effect.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding the QFabric System Data Plane on page 1199

node-device (Aliases)

Syntax	<code>node-device <i>node-device-name</i> { <i>assigned-node-device-name</i>; }</code>
Hierarchy Level	[edit fabric aliases]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Specify the mapping of user-defined names to QFabric system Node devices.
<div>  NOTE: The following rules apply to QFabric component alias naming: <ul style="list-style-type: none"> Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters. The maximum length of an alias name is 30 characters. Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components. You cannot use the reserved names all, fabric, or director-group as an alias name. </div>	
Options	<p><i>node-device-name</i>—Specify a user-defined name for a QFabric system Node device.</p> <p><i>assigned-node-device-name</i>—Specify the Node device identifier or name that has been provided. Identifiers are usually autogenerated by the Director software, and names are usually provided by the administrator of the default partition.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Aliases for the QFabric System on page 1353 Understanding Node Devices on page 1186

node-device (Chassis)

Syntax node-device *name* {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 pic *pic-number* {
 fte {
 port *port-number*;
 port-range *port-range-low port-range-high*;
 }
 xe {
 port *port-number*;
 port-range *port-range-low port-range-high*;
 }
 }
 }

Hierarchy Level [edit chassis **node-group**]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure properties specific to a Node device in a QFabric system.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Link Aggregation on page 2537](#)

node-device (Resources)

Syntax	<code>node-device <i>node-device-name</i>;</code>
Hierarchy Level	<code>[edit fabric resources node-group <i>node-group-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Assign Node devices to Node groups within a QFabric system.



NOTE: Ensure that the Node devices you assign to a redundant server Node group are of the same type, either a QFX3500 Node or a QFX3600 Node. You cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Node Groups for the QFabric System on page 1363 • Understanding Node Groups on page 1190

node-group (Chassis)

Syntax `node-group name {
 aggregated-devices {
 ethernet {
 device-count number;
 }
 }
 alarm {
 interface-type {
 link-down (ignore | red | yellow);
 }
 }
 container-devices {
 device-count number;
 }
 node-device name {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 }
 pic pic-number {
 fte {
 port port-number;
 port-range port-range-low port-range-high;
 }
 xe {
 port port-number;
 port-range port-range-low port-range-high;
 }
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }`

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure properties specific to a Node group.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Link Aggregation on page 2537](#)

node-group (Resources)

Syntax	<code>node-group <i>node-group-name</i> { <i>network-domain</i>; node-device <i>node-device-name</i>; }</code>
Hierarchy Level	[edit fabric resources]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Define Node groups within a QFabric system.



NOTE: The following rules apply to QFabric Node group naming:

- The network Node group must use the predefined name *NW-NG-0*. You must use this name when adding Node devices to the network Node group. You cannot specify a different name. Also, you can configure only one network Node group per partition.
- Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of a Node group name is 30 characters.
- Node group names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names *all*, *fabric*, or *director-group* as a Node group name.

The remaining statements are explained separately.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Node Groups for the QFabric System on page 1363 • Understanding Node Groups on page 1190

pic (Port)

Syntax `pic pic-number {
 fte {
 (port port-number | port-range port-range-low port-range-high);
 }
 xe {
 (port port-number | port-range port-range-low port-range-high);
 }
}`

Hierarchy Level [edit `chassis interconnect-device name fpc slot`]

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description (QFX3600 Node device only) Configure port types on QFX3600 Node devices.



CAUTION: The Packet Forwarding Engine on the QFX3600 Node device is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the Node device.

Options `pic pic-number`—Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure `xe` (10-Gigabit Ethernet) type ports. Specify **1** to configure `fte` (40-gigabit data plane uplink) type ports.

`fte`—Configure a specific port or a range of ports to operate as 40-gigabit data plane uplink ports.

`xe`—Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet ports.

`port-number`—Port number on which you want to configure the port type. Valid values are 0 through 7 if the port type is `fte`, and 2 through 15 if the port type is `xe`.

`port-range-low`—Lowest-numbered port in the range of ports. The lowest possible value is 0 if the port type is `fte`. The lowest possible value is 2 if the port type is `xe`.

`port-range-high`—Highest-numbered port in the range of ports. The highest possible value is 7 if the port type is `fte`. The highest possible value is 15 if the port type is `xe`.



NOTE:

- By default, ports Q0 through Q3 operate as `fte` type ports, and ports Q4 through Q15 operate as `xe` type ports.
- Only ports Q0 through Q7 can be configured as `fte` type ports.
- Only ports Q2 through Q15 can be configured as `xe` type ports.



NOTE: When you delete the port type configuration for an individual port or a block of ports, the ports return to operating in their default port type.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Port Type on QFX3600 Node Devices on page 1367](#)

remote-debug-permission

Syntax remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);

Hierarchy Level [edit system login user *username* authentication]
[edit system root-authentication]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description (QFabric systems only) Configure authentication classes that permit or deny user access to individual components of the QFabric system.

Default qfabric-user

Options **qfabric-admin**—Permits a user to log in to individual QFabric system components, view operations, and change component configurations.

qfabric-operator—Permits a user to log in to individual QFabric system components and view component operations.

qfabric-user—Prevents a user from logging in to individual QFabric system components.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring QFabric System Login Classes on page 1345](#)
- [request component login on page 1444](#)
- [Understanding QFabric System Login Classes on page 1230](#)

resources

Syntax	<pre>resources { node-group node-group-name { network-domain; node-device node-device-name; } }</pre>
Hierarchy Level	[edit fabric]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>(QFabric systems only) Define resources within a QFabric system and handle a variety of component assignments.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Node Groups for the QFabric System on page 1363• Understanding Node Groups on page 1190

routing-options (QFabric System)

Syntax	<pre>routing-options { multicast { no-make-before-break; } }</pre>
Hierarchy Level	[edit fabric]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Set routing options in a QFabric system.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding the QFabric System Data Plane on page 1199

syslog (QFabric System)

Syntax	<pre> syslog { file <i>filename</i> { archive { size <i>maximum-file-size</i>; } explicit-priority; <i>facility severity</i>; match "<i>regular-expression</i>"; structured-data; } filter all { <i>facility severity</i>; match "<i>regular-expression</i>"; } host <i>hostname</i> { explicit-priority; <i>facility severity</i>; facility-override <i>facility</i>; log-prefix <i>string</i>; match "<i>regular-expression</i>"; structured-data; } } </pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Configure system log messages for the QFabric system.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding the Implementation of System Log Messages on the QFabric System on page 1227

Administration

- [Software Upgrade and Recovery on page 1405](#)
- [Operational Mode Commands on page 1436](#)

Software Upgrade and Recovery

- [Performing a Nonstop Software Upgrade on the QFabric System on page 1405](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [Upgrading Software on a QFabric System on page 1430](#)
- [Performing System Backup and Recovery for a QFabric System on page 1435](#)

Performing a Nonstop Software Upgrade on the QFabric System



NOTE: Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the `request system software add component all` command) does not require that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a nonstop software upgrade to Junos OS Release 12.2X50-D42.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device or one fabric control Routing Engine is always operational
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device
- Specifying an upgrade group if you want all Node devices in a Node group to be upgraded in parallel (which shortens the time of the upgrade)
- Rebooting devices automatically as part of the nonstop upgrade process

When performing a nonstop upgrade, start with the Director group upgrade, then issue the fabric upgrade, and end with the Node group upgrades.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade. For node-groups defined with two node-devices, both must be online in order for upgrade to succeed.



NOTE: Before you install the software, we recommend that you back up your current configuration files by issuing the `request system software configuration-backup` command.



NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the `request system software add component all` command.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 1407](#)
- [Downloading Software Files Using a Browser on page 1407](#)
- [Retrieving Software Files for Download on page 1408](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group on page 1408](#)
- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components on page 1408](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups on page 1409](#)
- [Performing a Nonstop Software Upgrade on a Node Group on page 1410](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@qfabric> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup/media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
 2. Click **Download Software**.
 3. In the **Switching** box, click **Junos OS Platforms**.
 4. In the **QFX Series** section, click the name of the platform for which you want to download software.
 5. Click the **Software** tab and select the release number from the **Release** drop-down list.
 6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device” on page 111](#).
- A login screen appears.
7. Enter your user ID and password and press **Enter**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the `jinstall-qfabric-version.rpm` file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@qfabric> request system software download /path/package-name
```

For example:

```
user@qfabric> request system software download  
ftp://server/files/jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group



NOTE: If you reboot any Node groups or Interconnect devices after you perform a nonstop upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group  
jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To verify that both Director devices are online, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric  
jinstall-qfabric-12.2X50-D10.3.rpm
```

(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node group, create an upgrade group and include the devices you want to reboot at the same time.



NOTE:

- Upgrade groups work best when you build in redundancy so that not all the Node devices within the Node group restart at the same time. We suggest using the upgrade group feature for Node groups that have at least 2 pairs of Node devices (4 or more members), such as the network Node group.
 - If you add Node devices that have links to the same link aggregation group (LAG), there might be traffic loss.
-
- Create the upgrade group by issuing the **set chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name* node-devices** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade1 node-devices
[ node1 node2 ]
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade2 node-devices
[ node3 node4 ]
```

Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.



NOTE: You can upgrade multiple Node groups with this command. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To install software on a Node group:

- Issue the **request system software nonstop-upgrade node-group *node-group-name* *package-name*** command.

To perform a nonstop upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1  
jinstall-qfabric-12.2X50-D10.3.rpm
```

To perform a nonstop upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1  
nodegroup2 nodegroup3] jinstall-qfabric-12.2X50-D10.3.rpm
```

Related Documentation

- [Configuring Graceful Restart for QFabric Systems on page 1375](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [show system software upgrade status on page 1570](#)

Verifying Nonstop Software Upgrade for QFabric Systems

This topic discusses how you can monitor the progress of each of the three steps in a nonstop software upgrade. By identifying the key actions and events that define this process, you can track the status of the upgrade with confidence.



TIP: When performing a nonstop software upgrade, open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a

second session to verify that the QFabric system components respond to operational mode commands as expected.

- [Verifying a Director Group Nonstop Software Upgrade on page 1411](#)
- [Verifying a Fabric Nonstop Software Upgrade on page 1423](#)
- [Verifying a Redundant Server Node Group Nonstop Software Upgrade on page 1425](#)
- [Verifying a Network Node Group Nonstop Software Upgrade on page 1428](#)

Verifying a Director Group Nonstop Software Upgrade

Purpose During the Director group portion of a nonstop software upgrade, you should expect to see the Director device that hosts the CLI session selected as the master device. When mastership of all processes moves to the master, the QFabric system upgrades the backup Director device and this Director device reboots. After the backup Director device comes back online, the master Director device suspends CLI operations for 15 minutes, upgrades itself, and reboots. At this point, the backup becomes the new master Director device and you can issue CLI operational commands. Finally, the former master comes back online as a backup and both devices are operational once again. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In one SSH session to the QFabric CLI, verify the current status of the QFabric system by issuing the **show fabric administration inventory**, **show fabric administration inventory director-group status**, and **show fabric session-host** commands. In this case, Director device DG0 is the master device but DG1 hosts the CLI session.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			

DRE-0 Connected Configured

session1@qfabric> show fabric administration inventory director-group status

Director Group Status Tue Jun 5 15:11:26 UTC 2012

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

```

-----
Management Interface          up
Control Plane Bridge          up
Control Plane LAG              up
CP Link [0/2]                  up
CP Link [0/1]                  up
CP Link [0/0]                  up
CP Link [1/2]                  down
CP Link [1/1]                  down
CP Link [1/0]                  down
Crossover LAG                  up
CP Link [0/3]                  up
CP Link [1/3]                  up

```

```

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032  online  backup

```

Director Group Managed Services

```

-----
Shared File System          online
Network File System         online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Volume ID:8                  optimal
Physical ID:1                online
Physical ID:0                online
SCSI ID:1                    100%
SCSI ID:0                    100%

```

```

Size  Used Avail Used% Mounted on
-----

```

```

423G 5.5G 395G  2%  /
99M 16M  79M  17% /boot
93G  7.3G 86G   8% /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online    backup

FTP Server                  online
Syslog                     online
Distributed Management      online
SNMP Trap Forwarder        online
SNMP Process               online
Platform Management         online

```

Interface Link Status

```

-----
Management Interface        up

```

```

Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up
CP Link [0/1]                 up
CP Link [0/0]                 up
CP Link [1/2]                 down
CP Link [1/1]                 down
CP Link [1/0]                 down
Crossover LAG                 up
CP Link [0/3]                 up
CP Link [1/3]                 up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

- In a second SSH session to the QFabric CLI, issue the request for the Director group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm

```

- If the CLI session is being hosted by the master Director device, skip to step 4. However, if the CLI session is hosted by the backup Director device, the Director group mastership switches to the backup device after you issue the nonstop software upgrade command. In this example, mastership switches to Director device DG1.

```

session1@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 15:12:20 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.49.215.38	8%	31905924k	0	3 days, 21:16 hrs
dg1	online	master	10.49.215.39	6%	18010368k	3	3 days, 21:16 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	backup

Director Group Managed Services

Shared File System	offline
Network File System	offline
Virtual Machine Server	offline
Load Balancer/DHCP	offline

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         offline
DHCP Server master          offline    backup

```

```

FTP Server                  online
Syslog                     online
Distributed Management      offline
SNMP Trap Forwarder        offline
SNMP Process               offline
Platform Management        online

```

Interface Link Status

```

-----
Management Interface       up
Control Plane Bridge       up
Control Plane LAG          up
CP Link [0/2]              up
CP Link [0/1]              up
CP Link [0/0]              up
CP Link [1/2]              down
CP Link [1/1]              down
CP Link [1/0]              down
Crossover LAG              up
CP Link [0/3]              up
CP Link [1/3]              up

```

```

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032  online  master

```

Master Services

```

-----
Database Server           online
Load Balancer Director     online
QFabric Partition Address  online

```

Director Group Managed Services

```

-----
Shared File System        online
Network File System        online
Virtual Machine Server     online
Load Balancer/DHCP         online

```

Hard Drive Status

```

-----
Volume ID:8               optimal
Physical ID:1              online
Physical ID:0              online
SCSI ID:1                  100%
SCSI ID:0                  100%

```

Size Used Avail Used% Mounted on

```

-----
423G 6.0G 395G 2% /
99M 16M 79M 17% /boot

```

93G 7.3G 86G 8% /pbdata

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online      master

FTP Server                  online
Syslog                      online
Distributed Management      online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

```

Interface Link Status

```

-----
Management Interface       up
Control Plane Bridge       up
Control Plane LAG          up
CP Link [0/2]              up
CP Link [0/1]              up
CP Link [0/0]              up
CP Link [1/2]              down
CP Link [1/1]              down
CP Link [1/0]              down
Crossover LAG              up
CP Link [0/3]              up
CP Link [1/3]              up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

4. The Director group nonstop software upgrade process continues by downloading and installing software for the fabric manager Routing Engines and the Director devices.

```

root@qfabric>
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

```

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

```
Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package
jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
```

5. When the system upgrades and reboots the backup Director device DGO, notice how this device is not displayed in the output of the **show fabric administration inventory director-group status** command. Because Director device DG1 appears, this means that the DG1 is operational and acts as the master device.



NOTE: If your second SSH session is being hosted by the rebooting Director device, your session terminates and you need to log back in to establish a new session running on the active Director device.

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 15:41:14 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg1	online	master	10.49.215.39	6%	8372272k	4	3 days, 21:25 hrs

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:8	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	6.0G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up

```

Control Plane LAG          up
CP Link [0/2]              up
CP Link [0/1]              up
CP Link [0/0]              up
CP Link [1/2]              down
CP Link [1/1]              down
CP Link [1/0]              down
Crossover LAG              up
CP Link [0/3]              up
CP Link [1/3]              up

```

6. The upgrade continues with master Director device DG1 suspending CLI services for 15 minutes, transferring mastership to Director device DG0, and then rebooting Director device DG1 (which terminates the CLI session).

```
root@qfabric>
```

```
[Peer Update Status]: Setting peer DG node as the master SFC
```

```
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
```

```
[Peer Update Status]: Check for VMs on dg0
```

```
Triggering Final Stage of Fabric Manager Upgrade:
```

```
Updating FM-0 to Junos version 12.2X50-D10.3
```

```
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
```

```
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
```

```
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
```

```
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
```

```
[DRE-0 2012-06-05 16:12:43]: Retrieving package
```

```
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
```

```
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
```

```
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
```

```
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
```

```
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
```

```
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
```

```
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
```

```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
```

```
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!
```

```
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by
peer
```

```
Connection to qfabric-partition0 closed.
```

7. Upon reopening the SSH session, notice that Director device DG0 is now the master device hosting the session and Director device DG1 does not appear in the QFabric system inventory while it is rebooting.

```
session1@qfabric> show fabric session-host
```

```
Identifier: 0281052011000001
```

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 16:21:23 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	13%	20739560k	3	36:29 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status


```

Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG             up
CP Link [0/3]             up
CP Link [1/3]             up

```

8. When Director device DG1 comes back online, it returns to the QFabric system inventory as a backup Director device and hosts some of the Routing Engine processes (which should appear load balanced between the master and backup Director devices).

```

session1@qfabric> show fabric administration inventory director-group status
root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 16:41:02 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	15%	14759920k	6	56:09 mins
dg1	online	backup	10.49.215.39	8%	31486680k	0	07:51 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master         online      master

FTP Server                  online
Syslog                      online
Distributed Management      online
SNMP Trap Forwarder        online
SNMP Process                online
Platform Management        online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	backup

Director Group Managed Services

```

-----
Shared File System          online
Network File System         online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Volume ID:8                 optimal
Physical ID:1               online
Physical ID:0               online
SCSI ID:1                   100%
SCSI ID:0                   100%

```

Size Used Avail Used% Mounted on

```

-----
423G 5.3G 396G 2% /
99M 16M 79M 17% /boot
93G 7.4G 86G 8% /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online

```

```

Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process         online
Network File System          online
DHCP Server master           online    backup

FTP Server                   online
Syslog                       online
Distributed Management       online
SNMP Trap Forwarder          online
SNMP Process                 online
Platform Management          online

```

Interface Link Status

```
-----
```

```

Management Interface        up
Control Plane Bridge         up
Control Plane LAG            up
CP Link [0/2]                up
CP Link [0/1]                up
CP Link [0/0]                up
CP Link [1/2]                down
CP Link [1/1]                down
CP Link [1/0]                down
Crossover LAG                up
CP Link [0/3]                up
CP Link [1/3]                up

```

```
session1@qfabric> show fabric administration inventory infrastructure
dg0:
```

Routing Engine Type	Hostname	PID	CPU-Use(%)
Fabric control	QFabric_default_FC-1_RE0	27906	2.5
Network Node group	QFabric_default_NW-NG-1_RE1	20421	1.8
Fabric manager	FM-0	4211	1.8
Debug Routing Engine	QFabric_DRE	1575	3.3

```
dg1:
```

Routing Engine Type	Hostname	PID	CPU-Use(%)
Fabric control	QFabric_default_FC-0_RE0	5686	2.3
Network Node group	QFabric_default_NW-NG-0_RE0	5866	1.9
Fabric manager	FM-1	572	1.6

Verifying a Fabric Nonstop Software Upgrade

Purpose During the fabric portion of a nonstop software upgrade, you should expect to see both fabric control Routing Engines upgrade first, followed by the upgrade of each Interconnect

device one at a time. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the fabric nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0      2012-06-05 16:48:53]: Retrieving package
[FC-1      2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0      2012-06-05 16:49:51]: ----- re0: -----
[FC-1      2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active
partition
[Status    2012-06-05 16:58:34]: Load completed with 0 errors...
[Status    2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status    2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status    2012-06-05 16:58:39]: Rebooting FC-0
[Status    2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status    2012-06-05 16:58:44]: Rebooting FC-1
[Status    2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status    2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

2. When the fabric components reboot, they appear as **Disconnected** in the output of the **show fabric administration inventory infrastructure fabric-controls** and **show fabric administration inventory interconnect-devices** commands.

```
session1@qfabric> show fabric administration inventory infrastructure fabric-controls
Item                               Identifier                               Connection                               Configuration
Fabric control
FC-0                               Disconnected
FC-1                               Disconnected

session1@qfabric> show fabric administration inventory interconnect-devices IC-F4912
Item                               Identifier                               Connection                               Configuration
Interconnect device
```

```

IC-F4912
F4912/RE0
Disconnected
Disconnected

```

3. When the fabric components return to full service, they appear as **Connected** in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

Verifying a Redundant Server Node Group Nonstop Software Upgrade

Purpose During the redundant server Node group portion of a nonstop software upgrade, you should expect to see the backup Node device upgrade first, followed by the upgrade of the master Node device. Server Node groups with a single device upgrade the device in the same way as a standalone switch. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the redundant server Node group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): RSNG

```

```

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system

```

```
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last
group :1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed
[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
                    RSNG in progress
```

3. During the redundant server Node group upgrade, the backup Node device (in this case, P1571-C) is upgraded first and appears in the **Disconnected** state in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
	NW-NG-0	Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Disconnected	
Interconnect device			
	IC-F4912	Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
	FM-0	Connected	Configured
Fabric control			
	FC-0	Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			
	DRE-0	Connected	Configured

4. After the backup Node device comes back online, the master Node device (in this case, P1550-C) appears in the **Disconnected** state in the output of the **show fabric administration inventory** command while the master Node device upgrades its software.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
	NW-NG-0	Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Disconnected	
	P1571-C	Connected	
Interconnect device			
	IC-F4912	Connected	Configured
	F4912/RE0	Connected	

Fabric manager FM-0	Connected	Configured
Fabric control FC-0	Connected	Configured
FC-1	Connected	Configured
Diagnostic routing engine DRE-0	Connected	Configured

5. After both Node devices in the redundant server Node group come back online, both Node devices appear as **Connected** to indicate the successful completion of the Node group nonstop software upgrade step.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

Verifying a Network Node Group Nonstop Software Upgrade

Purpose During the network Node group portion of a nonstop software upgrade, you should expect to see the backup network Node group Routing Engine upgrade first, followed by the Node devices within the network Node group upgrading one at a time, and ending with the upgrade of the master network Node group Routing Engine. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.



NOTE: If you configure an upgrade group for Node groups containing 2 or more Node devices, all Node devices within the upgrade group reboot at the same time.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the network Node group nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade node-group NW-NG-0
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): NW-NG-0
```

```
[NW-NG-0 2012-06-01 09:45:06]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[NW-NG-0 2012-06-01 09:45:06]: Retrieving package
[NW-NG-0 2012-06-01 09:46:18]: Pushing bundle to fpc0
[NW-NG-0 2012-06-01 09:46:52]: fpc0: Validate package...
[NW-NG-0 2012-06-01 09:53:26]: ----- fpc0 -----
[NW-NG-0 2012-06-01 09:54:01]: Step 1 of 20 Creating temporary file system
[NW-NG-0 2012-06-01 09:54:01]: Step 2 of 20 Determining installation source
[NW-NG-0 2012-06-01 09:54:02]: Step 3 of 20 Processing format options
[NW-NG-0 2012-06-01 09:54:02]: Step 4 of 20 Determining installation slice
[NW-NG-0 2012-06-01 09:54:02]: Step 5 of 20 Creating and labeling new slices
[NW-NG-0 2012-06-01 09:54:03]: Step 6 of 20 Create and mount new file system
[NW-NG-0 2012-06-01 09:54:10]: Step 7 of 20 Getting OS bundles
[NW-NG-0 2012-06-01 09:54:10]: Step 8 of 20 Updating recovery media
[NW-NG-0 2012-06-01 09:54:31]: Step 9 of 20 Extracting incoming image
[NW-NG-0 2012-06-01 09:55:43]: Step 10 of 20 Unpacking OS packages
[NW-NG-0 2012-06-01 09:55:46]: Step 11 of 20 Mounting jbase package
[NW-NG-0 2012-06-01 09:56:09]: Step 12 of 20 Creating base OS symbolic links
[NW-NG-0 2012-06-01 09:57:05]: Step 13 of 20 Creating fstab
[NW-NG-0 2012-06-01 09:57:05]: Step 14 of 20 Creating new system files
[NW-NG-0 2012-06-01 09:57:05]: Step 15 of 20 Adding jbundle package
[NW-NG-0 2012-06-01 09:59:30]: Step 16 of 20 Backing up system data
[NW-NG-0 2012-06-01 09:59:44]: Step 17 of 20 Setting up shared partition data
[NW-NG-0 2012-06-01 09:59:44]: Step 18 of 20 Checking package sanity in
installation
[NW-NG-0 2012-06-01 09:59:44]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[NW-NG-0 2012-06-01 09:59:47]: Step 20 of 20 Setting da0s1 as new active
partition
[NW-NG-0 2012-06-01 09:59:55]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-dc-re.tgz
[NW-NG-0 2012-06-01 09:59:55]: Retrieving package
[NW-NG-0 2012-06-01 10:01:04]: Pushing bundle to re1
[NW-NG-0 2012-06-01 10:01:35]: re1: Validate package...
[NW-NG-0 2012-06-01 10:02:56]: re0: Validate package...
[NW-NG-0 2012-06-01 10:04:45]: Rebooting Backup RE
[NW-NG-0 2012-06-01 10:08:31]: Initiating Chassis In-Service-Upgrade
[NW-NG-0 2012-06-01 10:08:52]: Upgrading group: 0 fpc: 0
[NW-NG-0 2012-06-01 10:18:33]: Upgrade complete for group:0
[NW-NG-0 2012-06-01 10:18:33]: Finished processing all upgrade groups, last
group :0
[NW-NG-0 2012-06-01 10:18:37]: Preparing for Switchover
[NW-NG-0 2012-06-01 10:18:55]: Switchover Completed
[Status 2012-06-01 10:18:58]: Upgrade completed with 0 errors
Success
```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
NW-NG-0 in progress
```

3. Verify the progress of the upgrade by issuing the **show chassis nonstop-upgrade node-group**, **show fabric administration inventory**, **show fabric administration inventory infrastructure**, and **show fabric administration inventory node-groups NW-NG-0** commands. You should see the backup network Node group Routing Engine reboot first, followed by each Node device within the network Node group, and ending with the reboot of master network Node group Routing Engine. Restarting devices appear as **Disconnected** in the output of the **show fabric administration inventory** command and restarting Routing Engines do not appear in output of the **show fabric administration inventory infrastructure** command until they return to service.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [show chassis nonstop-upgrade node-group on page 817](#)
- [show fabric administration inventory on page 1543](#)
- [show fabric administration inventory director-group status on page 1548](#)
- [show fabric administration inventory infrastructure on page 1553](#)
- [show fabric administration inventory interconnect-devices on page 1556](#)
- [show fabric administration inventory node-groups on page 1560](#)
- [show system software upgrade status on page 1570](#)

Upgrading Software on a QFabric System

The QFabric system software package contains software for all of the different components in the QFabric system, such as the Director group, Interconnect devices, Node devices, and other QFabric system components. You can upgrade the software on all of the QFabric components at the same time using the **request system software add package-name component all reboot** command.



NOTE: Downgrading software on a QFabric system is not supported.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 1430](#)
- [Downloading Software Files Using a Browser on page 1431](#)
- [Retrieving Software Files for Download on page 1432](#)
- [Installing the Software Package on the Entire QFabric System on page 1432](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@switch> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@switch> request system software configuration-backup /media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
 2. Click **Download Software**.
 3. In the **Switching** box, click **Junos OS Platforms**.
 4. In the **QFX Series** section, click the name of the platform for which you want to download software.
 5. Click the **Software** tab and select the release number from the **Release** drop-down list.
 6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device” on page 111](#).
- A login screen appears.
7. Enter your user ID and password and click **Login**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@switch> request system software download /path/package-name
```

For example:

```
user@switch> request system software download  
ftp://server/files/jinstall-qfabric-11.3X30.6.rpm
```

Installing the Software Package on the Entire QFabric System



NOTE: On a QFabric system, a QFX3500 Node device or QFX3600 Node device might not be able to participate as a Node device in the QFabric system if the Node device is running a different version of software from that of the Director group. This mismatch of software versions between the Node device and the Director group can occur when the Node device is introduced into the setup, and both Director devices go offline before the Node device completes its auto-upgrade process to upgrade its software version to the same software version running on the Director group. The workaround is to reboot the QFX3500 or QFX3600 Node device once the Director group comes back online. The QFX3500 or QFX3600 Node device will initiate auto-upgrade and upgrade its software version from the Director group.

1. Issue the **request system software add package-name component all reboot** command.

For example:

```
user@switch> request system software add jinstall-qfabric-11.3X30.6.rpm component all  
reboot
```



NOTE: If you receive an error message after issuing the **request system software add package-name component all reboot** command that says that the configuration file cannot be loaded as is, you will need to enter configuration mode, make any necessary changes to the configuration file, and then commit the changes.



NOTE: The default value for a QFabric system software upgrade is `validate`. The validation step adds up to 10 minutes to the overall software upgrade. If the validation fails, the upgrade does not proceed and the QFabric system automatically issues the request system software rollback command to restore the current software image. If you upgrade more than one component (for example, by issuing the `component all` option), validation failure on one device stops the upgrade process for the other devices. If you do not want to validate the software package against the current configuration, issue the `no-validate` option.

2. After the reboot has finished, verify that the new version of software has been properly installed by issuing the `show version component all` command.

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
```

JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

DRE-0:

-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FM-0:

-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

nodedevice1:

-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

interconnectdevice1:

-
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

- Related Documentation**
- [Software Installation Overview on page 120](#)
 - [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)
 - [Upgrading Jloader Software on QFX Series Devices on page 120](#)
 - [request system software add on page 394](#)
 - *Junos® OS Installation and Upgrade Guide*

Performing System Backup and Recovery for a QFabric System

Many routers and switches require an administrator to recover the software package and the configuration file for the device separately. In the case of a device failure, this means the administrator might need to perform two separate tasks (if neither the software package nor the configuration file can be recovered).

In contrast, the QFabric system uses a unique mechanism that saves the backup and recovery files for both the Junos OS software and the system configuration into a single collection. The following QFabric system backup and recovery mechanism simplifies and streamlines the recovery process so you can return to normal operations as quickly as possible.

To backup and recover your QFabric system:

1. (First time only) Implement the following one-time procedure to prepare your QFabric system to use the system backup and recovery feature:
 - Insert a Juniper Networks software installation USB flash drive into the master Director device. (This drive was provided to you as one of the components of your QFabric system shipment.)
 - Issue the **request system software format-qfabric-backup** command. The contents and format of the USB flash drive are copied to the Director group shared directory and are used as the basis for all future backup and recovery operations.

```
user@qfabric> request system software format-qfabric-backup
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```

 - Remove the Juniper Networks software installation USB drive from the master Director device.
2. Issue the **request system software system-backup** command to backup the software package and configuration file. This command saves the current files necessary to recover the QFabric system. The files are saved to a shared memory directory in the Director group.



NOTE: As you upgrade your system with new software and change the system configuration over time, remember to reissue this command periodically to save the newest files for recovery purposes.

```
user@qfabri c> request system software system-backup
```

user@qfabric>

3. Insert a 4 GB or larger USB flash drive into the master Director device for your Director group, and issue the **request system software system-backup usb-create** command. This command copies the recovery files that have been backed up in the Director group and transfers them to the USB flash drive to create a recovery USB drive.



NOTE: Issuing this command overwrites the contents of the USB flash drive with the QFabric system recovery files.

```
user@qfabric> request system software system-backup usb-create /dev/sdb
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```

4. Remove the recovery USB drive from the Director device, and store it securely in a known location that you will remember when you need to use the recovery USB drive.
5. If the QFabric system fails, power off the Director group, insert the recovery USB drive into the master Director device of your Director group, turn on power to the Director device, and follow the prompts to recover your system. This step restores the software package and the configuration file for your QFabric system.

**Related
Documentation**

- [request system software format-qfabric-backup on page 1460](#)
- [request system software system-backup on page 1467](#)

Operational Mode Commands

- [QFabric System Operational Mode Commands on page 1437](#)
- [Filtering Operational Mode Command Output in a QFabric System on page 1439](#)
- [request chassis device-mode](#)
- [request chassis fabric fpc](#)
- [request component login](#)
- [request fabric administration director-group change-master](#)
- [request fabric administration remove](#)
- [request fabric administration system mac-pool add](#)
- [request fabric administration system mac-pool delete](#)
- [request system halt](#)
- [request system reboot](#)
- [request system software format-qfabric-backup](#)

- `request system software nonstop-upgrade`
- `request system software system-backup`
- `set chassis display message`
- `show chassis device-mode`
- `show chassis ethernet-switch interconnect-device cb`
- `show chassis ethernet-switch interconnect-device fpc`
- `show chassis fabric connectivity`
- `show chassis fabric device`
- `show chassis lcd`
- `show chassis nonstop-upgrade node-group`
- `show fabric administration inventory`
- `show fabric administration inventory director-group status`
- `show fabric administration inventory infrastructure`
- `show fabric administration inventory interconnect-devices`
- `show fabric administration inventory node-devices`
- `show fabric administration inventory node-groups`
- `show fabric administration system mac-pool`
- `show fabric inventory`
- `show fabric session-host`
- `show log`
- `show system software upgrade status`

QFabric System Operational Mode Commands

Table 139 on page 1437 summarizes the command line interface (CLI) commands that you can use to monitor and troubleshoot the QFabric system operations.

Table 139: QFabric System Operational Mode Commands

Task	Command
Select the operating mode for the device.	<code>request chassis device-mode</code>
Set the Interconnect device Flexible PIC Concentrator (FPC) offline or online for the QFabric system.	<code>request chassis fabric fpc</code>
Log in to individual QFabric system components for device level troubleshooting.	<code>request component login</code>
Select a Director device to become the new primary device within a Director group.	<code>request fabric administration director-group change-master</code>
Remove a disconnected component from the QFabric system inventory.	<code>request fabric administration remove</code>

Table 139: QFabric System Operational Mode Commands (*continued*)

Task	Command
Add a MAC range to the MAC pool assigned to the QFabric system.	<code>request fabric administration system mac-pool add</code>
Delete a MAC range from the MAC block assigned to the QFabric system.	<code>request fabric administration system mac-pool delete</code>
Halt a Director device.	<code>request system halt director-device</code>
Reboot QFabric system components.	<code>request system reboot</code>
Upgrade the software version of the QFabric system by using the nonstop software upgrade method (which preserves forwarding functionality during the upgrade and enables components to upgrade on a rotating basis).	<code>request system software nonstop-upgrade</code>
Save the software package and system configuration files to be able to recover the QFabric system in the case of a system failure.	<code>request system software system-backup</code>
Specify information to be displayed on the LCD panel of a QFabric system device.	<code>set chassis display message</code>
Display information about the operating mode of the device.	<code>show chassis device-mode</code>
Display the status of Ethernet switching in the Control Board of an Interconnect device.	<code>show chassis ethernet-switch interconnect-device cb</code>
Display the status of Ethernet switching in the Flexible PIC Controller (FPC) of an Interconnect device.	<code>show chassis ethernet-switch interconnect-device fpc</code>
Display the status of the data plane connections in the QFabric system.	<code>show chassis fabric connectivity</code>
Display the fabric management status of devices in your QFabric system.	<code>show chassis fabric device</code>
Display information shown on the LCD screen of a QFabric system device.	<code>show chassis lcd</code>
Display the status of a nonstop software upgrade for a Node group.	<code>show chassis nonstop-upgrade node-group</code>
Display all devices that belong to the QFabric system.	<code>show fabric administration inventory</code>
Display the Director devices that belong to a QFabric system Director group.	<code>show fabric administration inventory director-group status</code>
Display the services running on the Director group for the QFabric system.	<code>show fabric administration inventory infrastructure</code>

Table 139: QFabric System Operational Mode Commands (*continued*)

Task	Command
Display the Interconnect devices that belong to a QFabric system.	<code>show fabric administration inventory interconnect-devices</code>
Display the Node devices that belong to the QFabric system.	<code>show fabric administration inventory node-devices</code>
Display the Node groups and the corresponding Node devices that belong to the QFabric system.	<code>show fabric administration inventory node-groups</code>
Display all devices that belong to the QFabric system.	<code>show fabric inventory</code>
Display the MAC addresses that belong to a QFabric system Director group.	<code>show fabric administration system mac-pool</code>
Display the Director device that hosts the QFabric CLI session.	<code>show fabric session-host</code>
Display the system log messages in the specified file.	<code>show log</code>
Display the status of a QFabric system software upgrade.	<code>show system software upgrade status</code>

Filtering Operational Mode Command Output in a QFabric System

When you issue an operational mode command in a QFabric system, the output generated can be fairly extensive because of the number of components contained within the system. To make the output more accessible, you can filter the output by appending the **| filter** option to the end of most Junos OS commands.

- To filter operational mode command output and limit it to a Node group, include the **| filter node-group *node-group-name*** option at the end of your Junos OS operational mode command.

```
root@qfabric> show interfaces terse | filter node-group NW-NG-0
```

Interface	Admin	Link	Proto	Local	Remote
NW-NG-0:dsc	up	up			
NW-NG-0:em0	up	up			
NW-NG-0:em1	up	up			
NW-NG-0:gre	up	up			
NW-NG-0:ipip	up	up			
NW-NG-0:lo0	up	up			
NW-NG-0:lo0.16384	up	up	inet	127.0.0.1	--> 0/0
NW-NG-0:lo0.16385	up	up	inet		
NW-NG-0:lsi	up	up			
NW-NG-0:mtun	up	up			
NW-NG-0:pimd	up	up			
NW-NG-0:pime	up	up			
NW-NG-0:tap	up	up			
Node01:ge-0/0/10	up	up			
Node01:ge-0/0/40	up	up			

```
Node01:ge-0/0/41      up    up
vlan                  up    up
```


2. To filter operational mode command output and limit it to a set of Node groups, include the **| filter node-group** option at the end of your Junos OS operational mode command and specify the list of Node group names in brackets.

```
root@qfabric> show ethernet-switching interfaces | filter node-group [NW-NG-0 RSNG-1]
```

Interface	State	VLAN members	Tag	Tagging	Blocking
NW-NG-0:ae0.0	up	v200	200	tagged	unblocked
		v50	50	tagged	unblocked
		v51	51	tagged	unblocked
		v52	52	tagged	unblocked
		v53	53	tagged	unblocked
RSNG-1:ae0.0	up	v200	200	untagged	unblocked
RSNG-1:ae47.0	up	v50	50	tagged	unblocked
		v51	51	tagged	unblocked
		v52	52	tagged	unblocked
		v53	53	tagged	unblocked

- Related Documentation**
- [QFabric System Operational Mode Commands on page 1437](#)
 - *Using the Pipe (|) Symbol to Filter Junos Command Output*

request chassis device-mode

Syntax	request chassis device-mode (node-device standalone)
Release Information	Command introduced in Junos OS Release 11.2 for the QFX Series.
Description	Select the operating mode for the device. For example, a QFX3500 or QFX3600 device can operate either as a single switch in <i>standalone</i> mode or as a QFabric system Node device in <i>node-device</i> mode.
<div>  NOTE: <ul style="list-style-type: none"> Issue the request chassis device-mode command only when your management station is connected directly to the device over a console port connection. Changing the device mode erases all configuration data on the device. When you convert a device from standalone to Node device mode, we recommend that you back up your device configuration to an external location before issuing the request chassis device-mode command. </div>	
Options	<p>node-device—Set the device to operate as a Node device within a QFabric system. To complete the Node device mode conversion process, you must connect the device to the QFabric system management control plane and reboot the device.</p> <p>standalone—Set the device to operate as a standalone switch. If the device is in Node device mode, you must reboot the device to return to standalone mode. Standalone mode is the factory default setting.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> Converting the Device Mode for a QFabric System Component on page 1258 show chassis device-mode on page 1475 Understanding Node Devices on page 1186 Understanding the QFabric System Hardware Architecture on page 1177
List of Sample Output	request chassis device-mode node-device (Starting in Standalone Mode) on page 1442 request chassis device-mode node-device (Starting in Node Device Mode) on page 1442 request chassis device-mode standalone (Starting in Node Device Mode) on page 1442 request chassis device-mode standalone (Starting in Standalone Mode) on page 1442
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis device-mode node-device (Starting in Standalone Mode)

```
user@switch> request chassis device-mode node-device
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

request chassis device-mode node-device (Starting in Node Device Mode)

```
user@switch> request chassis device-mode node-device
Device mode set to 'node-device' mode.
No reboot required.
```

request chassis device-mode standalone (Starting in Node Device Mode)

```
user@switch> request chassis device-mode standalone
Device mode set to 'standalone' mode.
Please reboot the system to complete the process.
```

request chassis device-mode standalone (Starting in Standalone Mode)

```
user@switch> request chassis device-mode standalone
Device mode set to 'standalone' mode.
No reboot required.
```

request chassis fabric fpc

Syntax	<code>request chassis fabric fpc interconnect-device <i>interconnect-device-name</i> slot <i>slot-number</i> (offline online)</code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Set the Interconnect device Flexible PIC Concentrator (FPC) offline or online for the QFabric system. When the FPC is offline, traffic is redirected to other FPCs and is not lost while you remove or install an FPC. After issuing this command, you must issue the request chassis fpc command.
Options	<p>interconnect-device <i>interconnect-device-name</i>—Set the Interconnect device containing the FPC you want to bring either offline or online.</p> <p>slot <i>slot-number</i>—Set the specific FPC slot on the Interconnect device.</p> <p>offline—Set the Interconnect device FPC to offline for removal.</p> <p>online—Set the Interconnect device FPC to online after installation.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request chassis fpc on page 363 • show chassis fabric connectivity on page 1520 • show chassis fabric device on page 1527 • Understanding Interconnect Devices on page 1183
List of Sample Output	request chassis fabric fpc online on page 1443 request chassis fabric fpc offline on page 1443
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fabric fpc online

```
user@qfabric> request chassis fabric fpc interconnect-device IC-YW3781 offline slot 15
Graceful offline of the fabric card has been initiated. Please wait 20 seconds
before offlining or removing the card.
```

request chassis fabric fpc offline

```
user@qfabric> request chassis fabric fpc interconnect-device IC-YW3781 online slot 15
Bring the FPC online by issuing the "request chassis fpc online" command.
```

request component login

Syntax	<code>request component login <i>component-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the request component login command, you must first provide the qfabric-admin or qfabric-operator class privilege to your user (for more information, see: remote-debug-permission).
Options	<i>component-name</i> —Specify the QFabric system component to which you wish to log in.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring QFabric System Login Classes on page 1345 • remote-debug-permission on page 1401 • Understanding QFabric System Login Classes on page 1230
List of Sample Output	request component login (with qfabric-admin Privileges) on page 1444 request component login (with qfabric-operator Privileges) on page 1445 request component login (with qfabric-user Privileges) on page 1445

Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
  clear          Clear information in the system
  file           Perform file operations
  help           Provide help information
  load           Load information from file
  monitor        Show real-time debugging information
  mtrace         Trace multicast path from source to receiver
  op             Invoke an operation script
  ping           Ping remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart software process
  save           Save information to file
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  ssh            Start secure shell on another host
  start          Start shell
```



```

telnet          Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host{master}
qfabric-admin@node-ee3093>

```

request component login (with qfabric-operator Privileges)

```

operator@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-operator@node-ee3093> ?
Possible completions:
file          Perform file operations
help          Provide help information
load          Load information from file
op            Invoke an operation script
quit          Exit the management session
request       Make system-level requests
save          Save information to file
set           Set CLI properties, date/time, craft interface message
show          Show system information
start         Start shell
test          Perform diagnostic debugging
{master}
qfabric-operator@node-ee3093>

```

request component login (with qfabric-user Privileges)

```

user0@qfabric> request component login EE3093
error: User user0 does not have sufficient permissions to login to device ee3093

```

request fabric administration director-group change-master

Syntax	<code>request fabric administration director-group change-master (director-device <i>director-device-name</i>)</code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Select a Director device to become the new primary device within a Director group. The specified device becomes the new master Director device, and the previous master Director device becomes a backup Director device.
Options	<p>none—Change the device that controls the Director group. Assign the current backup Director device as the new master and the current master Director device as the backup.</p> <p>director-device <i>director-device-name</i>—Specify which Director device should become the primary device within the Director group.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none">• show fabric administration inventory director-group status on page 1548• Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335• Understanding the Director Group on page 1180
List of Sample Output	request fabric administration director-group change-master on page 1446


Sample Output

request fabric administration director-group change-master

```
user@qfabric> request fabric administration director-group change-master
Do you intend to switchover mastership? [yes,no] (no) yes

Cluster master successfully switched
```

request fabric administration remove

Syntax	<code>request fabric administration remove (interconnect-device <i>interconnect-device-name</i> node-device <i>node-device-name</i>)</code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Remove a disconnected Interconnect or Node device from the QFabric system inventory so that it does not appear in the output of the show fabric administration inventory command.
<div>  NOTE: <ul style="list-style-type: none"> You cannot remove any devices that appear in the Connected state in the output of the show fabric administration inventory command. For Node devices, you can only remove a device if it belongs to an autogenerated server Node group that contains a single Node device. Node devices contained within redundant server Node groups or network Node groups cannot be removed directly. To remove a Node device that is part of a group, delete the device from the Node group configuration first before attempting to remove the device from the inventory. </div>	
Options	<p>interconnect-device <i>interconnect-device-name</i>—Remove a disconnected Interconnect device from the QFabric system inventory.</p> <p>node-device <i>node-device-name</i>—Remove a disconnected Node device from the QFabric system inventory.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> show fabric administration inventory on page 1543 show fabric administration inventory interconnect-devices on page 1556 show fabric administration inventory node-devices on page 1558 show fabric administration inventory node-groups on page 1560
List of Sample Output	request fabric administration remove interconnect-device on page 1447 request fabric administration remove node-device on page 1448

Sample Output

request fabric administration remove interconnect-device

```
user@qfabric> request fabric administration remove interconnect-device IC1
Device successfully removed
```

Sample Output

`request fabric administration remove node-device`

```
user@qfabri> request fabric administration remove node-device node5
Device successfully removed
```

request fabric administration system mac-pool add


Syntax	<code>request fabric administration system mac-pool add mac-base <i>starting-mac-address</i> count <i>number-of-mac-address</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Add a MAC address pool to expand the initial set of MAC addresses assigned to the QFabric system.
Options	<p>mac-base <i>starting-mac-address</i>—Set the starting MAC address for a pool of addresses assigned to the QFabric system.</p> <p>count <i>number-of-mac-address</i>—Set the total number of MAC addresses in the specified address pool assigned to the QFabric system.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request fabric administration system mac-pool delete on page 1450 • show fabric administration system mac-pool on page 1562
List of Sample Output	request fabric administration system mac-pool add mac-base starting-mac-address count on page 1449
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

`request fabric administration system mac-pool add mac-base starting-mac-address count`

```
user@switch> request fabric administration system mac-pool add mac-base 02:00:00:11:22:00
count 10
```

request fabric administration system mac-pool delete

Syntax	request fabric administration system mac-pool delete mac-base <i>starting-mac-address</i>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Delete a range of MAC addresses assigned manually to the QFabric system.
<div> NOTE: You cannot delete the MAC address range assigned during the initial setup of the QFabric system. Also, you cannot delete a MAC address range if the MAC address block is still in use.</div>	
Options	mac-base <i>starting-mac-address</i> —Specify the starting MAC address for a pool of addresses you wish to remove from the QFabric system.
Additional Information	After you issue the request fabric administration system mac-pool delete command, issue the show fabric administration system mac-pool command to verify that the MAC address range has been deleted.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none">• request fabric administration system mac-pool add on page 1449• show fabric administration system mac-pool on page 1562
List of Sample Output	request fabric administration system mac-pool delete mac-base on page 1450


Sample Output

request fabric administration system mac-pool delete mac-base

```
user@switch> request fabric administration system mac-pool delete mac-base 02:00:00:11:22:00
```

request system halt

List of Syntax	Syntax on page 1451 Syntax (EX Series Switches) on page 1451 Syntax (PTX Series) on page 1451 Syntax (TX Matrix Router) on page 1451 Syntax (TX Matrix Plus Router) on page 1451 Syntax (MX Series Router) on page 1452 Syntax (QFX Series) on page 1452
Syntax	<pre>request system halt <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk removable-compact-flash usb)> <message "<i>text</i>"></pre>
Syntax (EX Series Switches)	<pre>request system halt <all-members> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "<i>text</i>"> <other-routing-engine> <slice <i>slice</i>></pre>
Syntax (PTX Series)	<pre>request system halt <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Router)	<pre>request system halt <all-lcc lcc <i>number</i> scc> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "<i>text</i>"></pre>
Syntax (TX Matrix Plus Router)	<pre>request system halt <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>></pre>

	<pre><at <i>time</i>> <backup-routing-engine> <both-routing-engines> <other-routing-engine> <in <i>minutes</i>> <media (compact-flash disk)> <message "text"></pre>
Syntax (MX Series Router)	<pre>request system halt <all-members> <at <i>time</i>> <backup-routing-engine> <both-routing-engines> <in <i>minutes</i>> <local> <media (external internal)> <member <i>member-id</i>> <message "text"> <other-routing-engine></pre>
Syntax (QFX Series)	<pre>request system halt <all-members> <at <i>time</i>> <director-device <i>director-device-id</i>> <in <i>minutes</i>> <local> <media > <member <i>member-id</i>> <message "text"> <slice <i>slice</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>other-routing-engine option introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>director-device option introduced for QFabric systems in Junos OS Release 12.2.</p> <p>backup-routing-engine option introduced in Junos OS Release 13.1.</p>
Description	<p>Stop the router or switch software.</p>
	<div><p>NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, use the member option from the Node group CLI. You cannot issue this command from the QFabric CLI.</p></div>
Options	<p>none—Stop the router or switch software immediately.</p> <p>all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Halt all chassis.</p>

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 or T4000 routers connected to the TX Matrix Plus router.

all-members—(EX4200 switches and MX Series routers only) (Optional) Halt all members of the Virtual Chassis configuration.

at time —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

backup-routing-engine—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

both-routing-engines—(Optional) Halt both Routing Engines at the same time.

director-device *director-device-id*—(QFabric systems only) Halt a specific Director device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches and MX Series routers only) (Optional) Halt the local Virtual Chassis member.

in *minutes*—(Optional) Number of minutes from now to stop the software. This option is an alias for the **at +minutes** option.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to J Series routers only.)

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

member *member-id*—(EX4200 switches and MX Series routers only) (Optional) Halt the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

message "*text*"—(Optional) Message to display to all system users before stopping the software.

other-routing-engine—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

slice *slice*—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.



NOTE: If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the **request system halt** command. You can also halt both Routing Engines at the same time by issuing the **request system halt both-routing-engines** command.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• clear system reboot on page 332• request system power-off on page 385• Rebooting and Halting a QFX Series Product on page 172• Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	request system halt on page 1456 request system halt (In 2 Hours) on page 1456 request system halt (Immediately) on page 1456 request system halt (At 1:20 AM) on page 1456
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host> request system halt at now
```

request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

request system reboot

Syntax	request system reboot <at time> <in minutes> <media (external internal)> <message "text"> <routing-engine>
Syntax (QFX Series)	request system reboot <all <graceful>> <director-device <i>name</i> > <director-group <graceful>> <fabric <graceful>> <node-group <i>name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Reboot the Junos OS.



NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Reboot requests are recorded in the system log files, which you can view with the **show log messages** command. You can view the process names with the **show system processes** command.

- Options**
- none**—Reboots the software immediately.
 - all**—(QFabric systems only) (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.
 - at time**—(Optional) Time at which to reboot the software, specified in one of the following ways:
 - **+minutes**—Number of minutes from now to reboot the software.
 - **hh:mm**—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
 - **now**—Stop or reboot the software immediately. This is the default.
 - **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
 - director-device *name***—(QFabric systems only) (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group—(QFabric systems only) (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric—(QFabric systems only) (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

graceful—(QFabric systems only) (Optional) Allows the QFabric component to reboot with minimal impact to network traffic. This option is only available for the **all**, **fabric**, and **director-group** options.

in minutes—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at +minutes** option.

media (external | internal)—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

message "text"—(Optional) Message to display to all system users before rebooting the software.

node-group name—(QFabric systems only) (Optional) Reboots the software on a server Node group or a network Node group.

routing-engine—(Optional) Reboot the Routing Engine.

Required Privilege Level maintenance

Related Documentation

- [clear system reboot on page 332](#)
- [Rebooting and Halting a QFX Series Product on page 172](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

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

request system reboot (At 2300)

```
user@switch> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request the system to reboot in 2 hours:

```
user@switch> request system reboot at +120
user@switch> request system reboot in 120
user@switch> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@switch> request system reboot at now
```

request system reboot (At 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@switch> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot director-device

```
user@switch> request system reboot director-device Node1
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group

```
user@switch> request system reboot director-group
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system reboot director-group graceful

```
user@switch> request system reboot director-group graceful
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

request system software format-qfabric-backup

Syntax	request system software format-qfabric-backup
Release Information	Command introduced in Junos OS Release 13.1 for the QFX Series.
Description	(QFabric systems only) Copy the install media files from a USB flash drive to your QFabric system recovery directory on the Director group. You must issue this command before you can use the request system software system-backup and request system software system-backup copy-to-usb commands.
Options	none —Copy the install media files from a USB flash drive to a Director group recovery directory.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Performing System Backup and Recovery for a QFabric System on page 1435• request system software system-backup on page 1467• Performing a QFabric System Recovery Installation on the Director Group on page 112
List of Sample Output	request system software format-qfabric-backup on page 1460
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software format-qfabric-backup

```
user@qfabric> request system software format-qfabric-backup
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```


request system software nonstop-upgrade

Syntax	request system software nonstop-upgrade <i>package-name</i> <fabric > <director-group> <node-group <i>name</i> >
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.
Options	<p><i>package-name</i>—Location from which the software is to be installed. For example:</p> <ul style="list-style-type: none"> • <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following: <ul style="list-style-type: none"> • ftp—File Transfer Protocol. Use <i>ftp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>ftp://<username>:<password>@hostname/pathname/package-name</i>. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed. • http—Hypertext Transfer Protocol. Use <i>http://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>http://<username>:<password>@hostname/pathname/package-name</i>. If a password is required and you omit it, you are prompted for it. • scp—Secure copy (available only for Canada and U.S. version). Use <i>scp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>scp://<username>:<password>@hostname/pathname/package-name</i>.



NOTE:

- The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

director-group—Install software package on the Director group and Fabric managers.

fabric—Install software package on the Interconnect devices and Fabric controls.

node-group *name* —Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Performing a Nonstop Software Upgrade on the QFabric System on page 105• Verifying Nonstop Software Upgrade for QFabric Systems on page 1410• show chassis nonstop-upgrade node-group on page 817
List of Sample Output	request system software nonstop-upgrade director-group on page 1462 request system software nonstop-upgrade fabric on page 1464 request system software nonstop-upgrade node-group (Redundant Server Node Group) on page 1464 request system software nonstop-upgrade node-group (Server Node Group) on page 1466
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software nonstop-upgrade director-group](#)

```
user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
```

```

Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0 2012-06-05 16:12:43]: Retrieving package
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...

```

```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...  
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!  
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer  
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric  
jinstall-qfabric-12.2X50-D10.3.rpm  
[FC-0 2012-06-05 16:48:53]: Retrieving package  
[FC-1 2012-06-05 16:48:53]: Retrieving package  
[IC-F4912 2012-06-05 16:48:59]: Retrieving package  
[FC-0 2012-06-05 16:49:51]: ----- re0: -----  
[FC-1 2012-06-05 16:49:52]: ----- re0: -----  
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----  
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system  
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source  
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options  
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice  
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices  
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system  
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles  
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media  
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image  
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages  
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package  
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links  
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab  
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files  
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package  
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data  
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data  
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in  
installation  
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary  
file systems  
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition  
[Status 2012-06-05 16:58:34]: Load completed with 0 errors...  
[Status 2012-06-05 16:58:34]: Reboot is required to complete upgrade ...  
[Status 2012-06-05 16:58:34]: Trying to Connect to Node: FC-0  
[Status 2012-06-05 16:58:39]: Rebooting FC-0  
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1  
[Status 2012-06-05 16:58:44]: Rebooting FC-1  
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912  
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912  
Success
```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group RSNG  
jinstall-qfabric-12.2X50-D10.3.rpm  
Upgrading target(s): RSNG  
  
[RSNG 2012-06-05 17:26:44]: Starting with package  
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
```

```

[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last group
:1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed

```



[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): P1507-C
```

```
[P1507-C 2012-06-26 14:02:44]: Retrieving package
[P1507-C 2012-06-26 14:03:21]: ----- P1507-C: -----
[P1507-C 2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C 2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C 2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C 2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C 2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C 2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C 2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C 2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
[P1507-C 2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C 2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C 2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C 2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C 2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C 2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C 2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C 2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C 2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C 2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in
installation
[P1507-C 2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[P1507-C 2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status 2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status 2012-06-26 14:10:32]: Rebooting P1507-C
[Status 2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```

request system software system-backup

Syntax	request system software system-backup <usb-create>
Release Information	Command introduced in Junos OS Release 13.1 for the QFX Series.
Description	(QFabric systems only) Save a copy of the current QFabric system configuration file and the current software package for recovery purposes. You can use these saved files to restore your QFabric system to full operation after a system failure or shutdown.
Options	none —Copy the QFabric system software package and system configuration file to a Director group recovery directory.
<div>  <p>NOTE: If this command fails, insert a Juniper Networks software installation USB flash drive into the master Director device and issue the <code>request system software format-qfabric-backup</code> command. For more details about this prerequisite procedure that is required before you can use the QFabric system backup and recovery feature, see “Performing System Backup and Recovery for a QFabric System” on page 1435.</p> </div>	
<p>usb-create—Copy the QFabric system software package and system configuration file from the Director group recovery directory to a USB flash drive. When the files have been copied, you can use the USB flash drive to help your QFabric system recover from a failure condition.</p>	
<div>  <p>NOTE: You must issue the <code>request system software system-backup</code> command (which saves the files to the Director group) before you can issue the <code>request system software system-backup usb-create</code> command.</p> </div>	
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Performing System Backup and Recovery for a QFabric System on page 1435 • request system software format-qfabric-backup on page 1460 • save on page 450 • request system software configuration-backup on page 402 • Performing a QFabric System Recovery Installation on the Director Group on page 112
List of Sample Output	request system software system-backup on page 1468 request system software system-backup usb-create on page 1468

Output Fields When you enter these commands, you are provided feedback on the status of your request.

Sample Output

request system software system-backup

```
user@qfabric> request system software system-backup
```

```
user@qfabric>
```

request system software system-backup usb-create

```
user@qfabric> request system software system-backup usb-create /dev/sdb
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```


set chassis display message

List of Syntax	Syntax on page 1469 Syntax (TX Matrix Router) on page 1469 Syntax (TX Matrix Plus Router) on page 1469 Syntax (QFabric Systems) on page 1469
Syntax	set chassis display message " <i>message</i> " <permanent>
Syntax (TX Matrix Router)	set chassis display message " <i>message</i> " (<i>lcc number</i> <i>scc</i>) <permanent>
Syntax (TX Matrix Plus Router)	set chassis display message " <i>message</i> " (<i>fpc-slot slot-number</i> <i>lcc number</i> <i>sfc number</i>) <permanent>
Syntax (QFabric Systems)	set chassis display (<i>interconnect-device</i> <i>node-device</i>) <i>device-id</i> message " <i>message</i> " <permanent>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option for TX Matrix Plus router introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series standalone switches.</p> <p>Command introduced in Junos OS Release 13.1 for QFabric systems.</p>
Description	Display or stop a text message on the craft interface display, which is on the front of the router, or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages three times, switching between messages every 60 seconds.



NOTE: On T Series routers, when this command is executed with the **permanent** option, the display of the text message alternates with that of the standard craft interface message continuously every 60 seconds.

By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options **"message"**—Message to display. On the craft interface display, if the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. An empty pair of quotation marks (" ") deletes the text message from the craft interface display. On the LCD panel display, the message is limited to 16 characters.

fpc-slot slot-number—(TX Matrix Plus routers and EX4200 and QFX Series only) On the router or switch, display the text message on the craft interface for a specific Flexible PIC Concentrator (FPC). Replace **slot-number** with a value from 0 through 31. On the

switch, display the text message for a specific member of a Virtual Chassis, where **fpc-slot slot-number** corresponds to the member ID. Replace **slot-number** with a value from 0 through 9. On the QFX Series, the **slot-number** is always 0. On a TX Matrix Plus router with 3D SIBs replace **slot-number** with a value from 0 through 63.

interconnect-device device-id—(QFabric systems only) Display the text message on the craft interface display of an Interconnect device within a QFabric system.

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

node-device device-id—(QFabric systems only) Display the text message on the craft interface display of a Node device within a QFabric system.

permanent—(Optional) Display a text message on the craft interface display or LCD panel display permanently.

scc—(TX Matrix routers only) Display the text message on the craft interface display of the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus routers only) Display the text message on the craft interface display of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

clear

Related Documentation

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\)](#)
- [clear chassis display message on page 328](#)
- [show chassis craft-interface](#)
- [show chassis lcd on page 785](#)
- [Understanding the Implementation of System Log Messages on the QFabric System on page 1227](#)

List of Sample Output

[set chassis display message \(Creating\) on page 1471](#)
[set chassis display message \(Deleting\) on page 1471](#)
[set chassis display message \(EX Series and QFX Series\) on page 1471](#)

[set chassis display message \(QFabric Systems - Interconnect Device\) on page 1472](#)

[set chassis display message \(QFabric Systems - Node Device\) on page 1474](#)

Output Fields See *show chassis craft-interface* or [show chassis lcd](#) for an explanation of output fields.

Sample Output

set chassis display message (Creating)

The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "NOC contact Dusty (888) 555-1234"
message sent

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|NOC contact Dusty |
|(888) 555-1234   |
+-----+
```

set chassis display message (Deleting)

The following example shows how to delete the display message and verify that the message is removed:

```
user@host> set chassis display message ""
message sent

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|host           |
|Up: 0+17:05:47 |
|              |
|Temperature OK  |
+-----+
```

set chassis display message (EX Series and QFX Series)

The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "Joe 888-555-4321"
```

message sent

```
user@host> show chassis lcd
Front panel contents for: FPC 0
-----
LCD screen:
00:RE qfx3500s
Joe 888-555-4321
LEDs status:
Status/Beacon LED: Yellow Blinking
```

set chassis display message (QFabric Systems - Interconnect Device)

The following example shows how to set the display message and verify the result:

```
user@qfabric> set chassis display interconnect-device IC-F1012 message "Bob 888-555-1234"
message sent
```

```
user@qfabric> show chassis lcd interconnect-device IC-F1012
Front Panel Module Information
-----
LCD screen:
IC-F1012      3 Alarms active
Bob 888-555-1234
LEDs status:
Status LED: Green
Power LED : Green
Major Alarm LED: off
Minor Alarm LED: Yellow
Fan 0 LED : Green
Fan 1 LED : Green
Fan 2 LED : Green
Fan 3 LED : Green
Fan 4 LED : Green
Fan 5 LED : Green
Fan 6 LED : Green
Fan 7 LED : Green
Fan 8 LED : Green
Fan 9 LED : Green
PEM 0 LED : Green
PEM 1 LED : Green
PEM 2 LED : Green
PEM 3 LED : off
PEM 4 LED : off
PEM 5 LED : off

LED info for: CB - 0
-----
LEDs status:
Status LED: Green
Mastership LED: Green
```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

```
LED info for: CB - 1
```

LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

LED info for: FC 1 FPC - 1

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

LED info for: RC 0 FPC - 8

LEDs status:

Status LED: Green

LED info for: RC 1 FPC - 9

LEDs status:

Status LED: Green

LED info for: RC 2 FPC - 10

LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

LEDs status:

Status LED: Green

LED info for: RC 4 FPC - 12

LEDs status:

```

Status LED: Green

LED info for: RC 5 FPC - 13
-----
LEDs status:
Status LED: Green

LED info for: RC 6 FPC - 14
-----
LEDs status:
Status LED: Green

LED info for: RC 7 FPC - 15
-----
LEDs status:
Status LED: Green

```

set chassis display message (QFabric Systems - Node Device)

The following example shows how to set the display message and verify the result:

```

user@qfabric> set chassis display node-device P3774-C message "Contact Bill T."
message sent

```

```

user@qfabric> show chassis lcd node-device P3774-C
Front panel contents for: P3774-C
-----
LCD screen:
P3774-C
Contact Bill T.

LEDs status:
Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

show chassis device-mode


Syntax	show chassis device-mode
Release Information	Command introduced in Junos OS Release 11.2 for the QFX Series.
Description	Display information about the operating mode of the device. For example, QFX3500 devices operate either as a single switch in standalone mode or as a QFabric system Node device in node-device mode.
	<div>  <p>NOTE: Issue the <code>show chassis device-mode</code> command only when your management station is connected directly to the device over a console port connection.</p> </div>
Options	There are no options for this command.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • Converting the Device Mode for a QFabric System Component on page 1258 • request chassis device-mode on page 1441 • Understanding Interconnect Devices on page 1183 • Understanding Node Devices on page 1186
List of Sample Output	show chassis device-mode (Interconnect Device Mode) on page 1476 show chassis device-mode (Interconnect Device Mode, but Node Device-Ready) on page 1476 show chassis device-mode (Interconnect Device Mode, but Standalone-Ready) on page 1476 show chassis device-mode (Node Device Mode) on page 1476 show chassis device-mode (Node Device Mode, but Interconnect Device-Ready) on page 1476 show chassis device-mode (Node Device Mode, but Standalone-Ready) on page 1476 show chassis device-mode (Standalone Mode) on page 1476 show chassis device-mode (Standalone Mode, but Interconnect Device-Ready) on page 1476 show chassis device-mode (Standalone Mode, but Node Device-Ready) on page 1477
Output Fields	Table 140 on page 1476 lists the output fields for the <code>show chassis device-mode</code> command. Output fields are listed in the approximate order in which they appear.

Table 140: show chassis device-mode Output Fields

Field Name	Field Description
Current device-mode	Existing operational mode for the device. The device can be in Interconnect device mode, Node device mode, or standalone mode.
Future device-mode after reboot	Future operational mode for the device after you reboot it. The device can be set to enter Interconnect device mode, Node device mode, or standalone mode.
	NOTE: To set the future mode of the device, issue the request chassis device-mode command.

Sample Output

show chassis device-mode (Interconnect Device Mode)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Interconnect Device Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Node-device
```

show chassis device-mode (Interconnect Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Standalone
```

show chassis device-mode (Node Device Mode)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

show chassis device-mode (Node Device Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Node Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode)

```
user@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
```



```
Current device-mode : Standalone
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Standalone Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Node-device
```

show chassis ethernet-switch interconnect-device cb

Syntax	<code>show chassis ethernet-switch interconnect-device <i>name</i> cb</code> <code><detail></code> <code><port <i>number</i>></code> <code><slot<i>number</i>></code>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFX3000-G QFabric systems only) Display Ethernet switch information for the Control Board (CB) ports in an Interconnect device.
Options	<p>none—Display Ethernet switch information about each connected port on each online CB in the Interconnect device.</p> <p>detail—(Optional) Display detailed status information for all CBs or for the CB in the specified slot in the Interconnect device.</p> <p>port <i>number</i>—(Optional) Display Ethernet switch information about a specific port on a CB in the Interconnect device.</p> <p>slot <i>number</i>—(Optional) Display Ethernet switch information about a CB in a specific slot in the Interconnect device.</p>
Required Privilege Level	view
List of Sample Output	show chassis ethernet-switch interconnect-device cb on page 1482 show chassis ethernet-switch interconnect-device cb detail on page 1483 show chassis ethernet-switch interconnect-device cb detail slot port on page 1488
Output Fields	Table 52 on page 466 lists the output fields for the show chassis ethernet-switch interconnect-device cb command. Output fields are listed in the approximate order in which they appear.

Table 141: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: <ul style="list-style-type: none"> • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .
Duplex is	Duplex type of the Ethernet link: full or half .

Table 141: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
Flow Control TX is Enabled (or Disabled)	Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
Flow Control RX is Enabled (or Disabled)	Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted packets of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.

Table 141: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.

Table 141: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device cb

```
user@switch> show chassis ethernet-switch interconnect-device IC-WS001 cb
Displaying summary for switch 0
Link is down on XE port 1 connected to device: FPC7
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 2 connected to device: FPC6
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 3 connected to device: FPC5
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 5 connected to device: FPC4
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 7 connected to device: FPC3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 9 connected to device: FPC2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 10 connected to device: FPC1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                326358
  RX Octets                 237947

Link is good on XE port 11 connected to device: FPC0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                548249
  RX Octets                 386013

Link is down on XE port 20 connected to device: SFP3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 21 connected to device: SFP2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 22 connected to device: SFP1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                1
  RX Octets                11704758

Link is good on XE port 23 connected to device: SFP0
  Speed is 1000Mb
```

```

Duplex is full
Autonegotiate is Enabled
TX Octets          1500022
RX Octets          11629453

```

```

Link is good on XE port 24 connected to device: VCCPD
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
TX Octets          23332467
RX Octets          1500023

```

```

Link is good on GE port 25 connected to device: SFI
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
TX Octets          643918
RX Octets          894548

```

show chassis ethernet-switch interconnect-device cb detail

```

user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb detail
Port statistics for CB switch

```

```

Link is down on XE port 1 connected to device: FPC7

```

```

Link is down on XE port 2 connected to device: FPC6

```

```

Link is down on XE port 3 connected to device: FPC5

```

```

Link is down on XE port 5 connected to device: FPC4

```

```

Link is down on XE port 7 connected to device: FPC3

```

```

Link is down on XE port 9 connected to device: FPC2

```

```

Statistics for port 10 connected to device FPC1:

```

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 319293
TX Packets 256-511 Octets 5043
TX Packets 512-1023 Octets 2072
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets              326415
TX Multicast Packets    0
TX Broadcast Packets    1
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets     0
TX FCS Error Counter    0
TX Fragment Counter     0
TX Byte Counter         71659246
TX Packet OK Counter    326415
TX Pause Packet Counter  0
TX Unicast Counter      326414
RX Packets 64 Octets     0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 235428
RX Packets 256-511 Octets 2134

```

```
RX Packets 512-1023 Octets 420
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets 237988
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 55821504
RX Unicast Frame Count 237988
RX Packet OK Count 237988
Statistics for port 11 connected to device FPC0:
TX Packets 64 Octets 0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 535483
TX Packets 256-511 Octets 9289
TX Packets 512-1023 Octets 3564
TX Packets 1024-1518 Octets 5
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets 548342
TX Multicast Packets 0
TX Broadcast Packets 1
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 120498414
TX Packet OK Counter 548342
TX Pause Packet Counter 0
TX Unicast Counter 548341
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 382931
RX Packets 256-511 Octets 2762
RX Packets 512-1023 Octets 386
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets 386082
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
```


RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	90717369
RX Unicast Frame Count	386082
RX Packet OK Count	386082

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	0
TX Packets 128-255 Octets	0
TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	0
TX Packets 1024-1518 Octets	1
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	1
TX Multicast Packets	1
TX Broadcast Packets	0
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1422
RX Packet OK Count	1
RX Packets 64 Octets	230013
RX Packets 65-127 Octets	174529
RX Packets 128-255 Octets	286735
RX Packets 256-511 Octets	343412
RX Packets 512-1023 Octets	172152
RX Packets 1024-1518 Octets	10500065
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11706906
RX Multicast Packets	11672320
RX Broadcast Packets	34460
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0

```
RX Jabbers 0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 2379464164
RX Packet OK Count 11706906
Statistics for port 23 connected to device SFP0:
TX Packets 64 Octets 3
TX Packets 65-127 Octets 484733
TX Packets 128-255 Octets 219112
TX Packets 256-511 Octets 129014
TX Packets 512-1023 Octets 503
TX Packets 1024-1518 Octets 666958
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 0
TX Octets 1500323
TX Multicast Packets 794098
TX Broadcast Packets 1040
TX Single Collision Frames 0
TX Mult. Collision Frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision Frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC Ctrl Frames 0
TX Frame Deferred Xmsns 0
TX Frame Excessive Deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 1065466891
RX Packet OK Count 1500323
RX Packets 64 Octets 341563
RX Packets 65-127 Octets 430810
RX Packets 128-255 Octets 318279
RX Packets 256-511 Octets 347147
RX Packets 512-1023 Octets 184798
RX Packets 1024-1518 Octets 10008993
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 11631590
RX Multicast Packets 10878484
RX Broadcast Packets 33420
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol Errors 0
RX Unsupported Opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
```

```

RX Pause Frame Counter      0
RX Byte Counter             1720484325
RX Packet OK Count          11631591
Statistics for port 24 connected to device VCCPD:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1176546
TX Packets 128-255 Octets   604988
TX Packets 256-511 Octets   690561
TX Packets 512-1023 Octets  356942
TX Packets 1024-1518 Octets 20507438
TX Packets 1519-2047 Octets 278
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 278
TX Octets                   23336753
TX Multicast Packets        22549383
TX Broadcast Packets        67862
TX Single Collision Frames  0
TX Mult. Collision Frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision Frames         0
TX PAUSEMAC Ctrl Frames    0
TX MAC Ctrl Frames         0
TX Frame Deferred Xmsns     0
TX Frame Excessive Deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             4191296788
RX Packet OK Count          23336753
RX Packets 64 Octets        3
RX Packets 65-127 Octets    484673
RX Packets 128-255 Octets   219074
RX Packets 256-511 Octets   129100
RX Packets 512-1023 Octets  516
RX Packets 1024-1518 Octets 666959
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                   1500325
RX Multicast Packets        794099
RX Broadcast Packets        1040
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol Errors            0
RX Unsupported Opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             1071469739
RX Packet OK Count          1500325
Statistics for port 25 connected to device SFI:

```

TX Packets 64 Octets	12
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	618363
TX Packets 256-511 Octets	4896
TX Packets 512-1023 Octets	806
TX Packets 1024-1518 Octets	19950
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	644028
TX Multicast Packets	4
TX Broadcast Packets	19954
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	167039705
RX Packet OK Count	644028
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	854776
RX Packets 256-511 Octets	14332
RX Packets 512-1023 Octets	5636
RX Packets 1024-1518 Octets	19954
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	894698
RX Multicast Packets	0
RX Broadcast Packets	19943
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	212658920
RX Packet OK Count	894698

show chassis ethernet-switch interconnect-device cb detail slot port

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb slot 1 port 1
```

```
re0:
```

```
-----
Port statistics for CB switch
```

```
Link is down on XE port 1 connected to device: FPC7
```

```
Link is down on XE port 2 connected to device: FPC6
```

```
Link is down on XE port 3 connected to device: FPC5
```

```
Link is down on XE port 5 connected to device: FPC4
```

```
Link is down on XE port 7 connected to device: FPC3
```

```
Link is down on XE port 9 connected to device: FPC2
```

```
Statistics for port 10 connected to device FPC1:
```

```
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 319366
TX Packets 256-511 Octets 5043
TX Packets 512-1023 Octets 2072
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets      326488
TX Multicast Packets      0
TX Broadcast Packets      1
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets      0
TX FCS Error Counter      0
TX Fragment Counter      0
TX Byte Counter      71675330
TX Packet OK Counter      326488
TX Pause Packet Counter   0
TX Unicast Counter      326487
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 235481
RX Packets 256-511 Octets 2134
RX Packets 512-1023 Octets 420
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets      238041
RX Multicast Packets      0
RX Broadcast Packets      0
RX FCS Errors      0
RX Fragments      0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets      0
RX Jabbers      0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter      55834224
RX Unicast Frame Count    238041
```

```
RX Packet OK Count          238041
Statistics for port 11 connected to device FPC0:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   535606
TX Packets 256-511 Octets   9289
TX Packets 512-1023 Octets  3564
TX Packets 1024-1518 Octets 5
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   548465
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             120525524
TX Packet OK Counter        548465
TX Pause Packet Counter     0
TX Unicast Counter          548464
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   383018
RX Packets 256-511 Octets   2762
RX Packets 512-1023 Octets  386
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   386169
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             90738249
RX Unicast Frame Count      386169
RX Packet OK Count          386169
```

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

```
TX Packets 64 Octets        0
TX Packets 65-127 Octets    0
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 1
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
```

```

TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 0
TX Octets 1
TX Multicast Packets 1
TX Broadcast Packets 0
TX Single Collision Frames 0
TX Mult. Collision Frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision Frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC Ctrl Frames 0
TX Frame Deferred Xmsns 0
TX Frame Excessive Deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 1422
RX Packet OK Count 1
RX Packets 64 Octets 230071
RX Packets 65-127 Octets 174571
RX Packets 128-255 Octets 286812
RX Packets 256-511 Octets 343500
RX Packets 512-1023 Octets 172203
RX Packets 1024-1518 Octets 10502544
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 11709701
RX Multicast Packets 11675110
RX Broadcast Packets 34465
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol Errors 0
RX Unsupported Opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 2383079858
RX Packet OK Count 11709701
Statistics for port 23 connected to device SFP0:
TX Packets 64 Octets 3
TX Packets 65-127 Octets 485048
TX Packets 128-255 Octets 219200
TX Packets 256-511 Octets 129053
TX Packets 512-1023 Octets 503
TX Packets 1024-1518 Octets 667127
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 0
TX Octets 1500934
TX Multicast Packets 794300

```

TX Broadcast Packets	1040
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1065764997
RX Packet OK Count	1500934
RX Packets 64 Octets	341648
RX Packets 65-127 Octets	431183
RX Packets 128-255 Octets	318367
RX Packets 256-511 Octets	347225
RX Packets 512-1023 Octets	184849
RX Packets 1024-1518 Octets	10011311
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11634583
RX Multicast Packets	10881071
RX Broadcast Packets	33425
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1723893006
RX Packet OK Count	11634583

Statistics for port 24 connected to device VCCPD:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1177102
TX Packets 128-255 Octets	605153
TX Packets 256-511 Octets	690727
TX Packets 512-1023 Octets	357044
TX Packets 1024-1518 Octets	20512235
TX Packets 1519-2047 Octets	278
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	278
TX Octets	23342539
TX Multicast Packets	22554760
TX Broadcast Packets	67872
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0

TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	4198344167
RX Packet OK Count	23342539
RX Packets 64 Octets	3
RX Packets 65-127 Octets	484985
RX Packets 128-255 Octets	219164
RX Packets 256-511 Octets	129139
RX Packets 512-1023 Octets	516
RX Packets 1024-1518 Octets	667128
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	1500935
RX Multicast Packets	794301
RX Broadcast Packets	1040
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1071770147
RX Packet OK Count	1500935

Statistics for port 25 connected to device SFI:

TX Packets 64 Octets	12
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	618503
TX Packets 256-511 Octets	4896
TX Packets 512-1023 Octets	806
TX Packets 1024-1518 Octets	19950
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	644168
TX Multicast Packets	4
TX Broadcast Packets	19954
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0

TX Frame Deferred Xms	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	167073305
RX Packet OK Count	644168
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	854972
RX Packets 256-511 Octets	14332
RX Packets 512-1023 Octets	5636
RX Packets 1024-1518 Octets	19954
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	894894
RX Multicast Packets	0
RX Broadcast Packets	19943
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	212702114
RX Packet OK Count	894894

show chassis ethernet-switch interconnect-device fpc

Syntax	<code>show chassis ethernet-switch interconnect-device <i>name</i> fpc</code> <code><detail></code> <code><port <i>number</i>></code> <code><slot <i>number</i>></code>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFX3000-G QFabric systems only) Display Ethernet switch information for the front card Flexible Port Concentrators (FPCs) in an Interconnect device.
Options	<p>none—Display Ethernet switch information about each connected port on each online FPC in the Interconnect device.</p> <p>detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot in the Interconnect device.</p> <p>port <i>number</i>—(Optional) Display Ethernet switch information about a specific port on an FPC in the Interconnect device.</p> <p>slot <i>number</i>—(Optional) Display Ethernet switch information about an FPC in a specific slot in the Interconnect device.</p>
Required Privilege Level	view
List of Sample Output	show chassis ethernet-switch interconnect-device fpc on page 1498 show chassis ethernet-switch interconnect-device fpc detail on page 1501 show chassis ethernet-switch fpc detail slot on page 1508 show chassis ethernet-switch fpc interconnect-device port on page 1514 show chassis ethernet-switch fpc interconnect-device detail port on page 1516
Output Fields	Table 52 on page 466 lists the output fields for the show chassis ethernet-switch interconnect-device fpc command. Output fields are listed in the approximate order in which they appear.

Table 142: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: <ul style="list-style-type: none"> • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .

Table 142: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
Duplex is	Duplex type of the Ethernet link: full or half .
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted frames of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.

Table 142: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.

Table 142: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device fpc

```

user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full

```

```
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          124638
RX Octets          86496

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82191
RX Octets          58979

Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          145475
RX Octets          206828

Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82290
RX Octets          59443

Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          40900
RX Octets          30013

Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          89456
RX Octets          123189
```

```
Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0
```

```
root@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
```

```
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          124697
RX Octets          86535
```

```
Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82229
RX Octets          59009
```

```
Link is good on XE port 28 connected to device: CB0
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          145544
RX Octets          206925
```

```
Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0
```

```
Summary for switch on FC1
```

```
Link is good on GE port 2 connected to device: FWD-SWITCH-0
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          82327
RX Octets          59472
```

```
Link is good on GE port 4 connected to device: FWD-SWITCH-1
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
```



```

Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          40918
RX Octets          30028

```

Link is good on XE port 28 connected to device: CB0

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          89500
RX Octets          123244

```

Link is good on XE port 29 connected to device: CB1

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

```

show chassis ethernet-switch interconnect-device fpc detail

```
user@host> show chassis ethernet-switch interconnect-device IC-WS001 fpc detail
```

Port statistics for FC0 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 121716
TX Packets 256-511 Octets 2200
TX Packets 512-1023 Octets 823
TX Packets 1024-1518 Octets 2
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                124742
TX Multicast Packets      0
TX Broadcast Packets      1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           27391588
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 85924
RX Packets 256-511 Octets 555
RX Packets 512-1023 Octets 86
RX Packets 1024-1518 Octets 1

```

```
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 86566
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 20380581
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets 0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 80374
TX Packets 256-511 Octets 1347
TX Packets 512-1023 Octets 532
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 82257
TX Multicast Packets 0
TX Broadcast Packets 1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 18146746
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 58410
RX Packets 256-511 Octets 522
RX Packets 512-1023 Octets 96
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 59030
RX Multicast Packets 0
RX Broadcast Packets 0
```

```

RX FCS Errors          0
RX Align Errors        0
RX Fragments           0
RX Symbol errors       0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets   0
RX Oversize Packets    0
RX Jabbers             0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter  0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter        13882179
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets    0
TX Packets 65-127 Octets 0
TX Packets 128-255 Octets 144334
TX Packets 256-511 Octets 1077
TX Packets 512-1023 Octets 182
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets              145596
TX Multicast Packets    0
TX Broadcast Packets    0
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets     0
TX FCS Error Counter    0
TX Fragment Counter     0
TX Byte Counter         34262760
RX Packets 64 Octets    0
RX Packets 65-127 Octets 1
RX Packets 128-255 Octets 202090
RX Packets 256-511 Octets 3547
RX Packets 512-1023 Octets 1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets              206998
RX Multicast Packets    0
RX Broadcast Packets    1
RX FCS Errors           0
RX Fragments            0
RX MAC Control Packets  0
RX Out of Range Length  0
RX Undersize Packets    0
RX Oversize Packets     0
RX Jabbers              0
RX Control Frame Counter 0
RX Pause Frame Counter  0
RX Byte Counter         45538262
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets    0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 0

```

TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	0
TX Packets 1024-1518 Octets	0
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	1
TX Multicast Packets	0
TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

Port statistics for FC1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80560
TX Packets 256-511 Octets	1279
TX Packets 512-1023 Octets	514
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82357
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0

```

TX Frame deferred Xmsns      0
TX Frame excessive deferl     0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              18059906
RX Packets 64 Octets         0
RX Packets 65-127 Octets     0
RX Packets 128-255 Octets    58733
RX Packets 256-511 Octets    639
RX Packets 512-1023 Octets   119
RX Packets 1024-1518 Octets  3
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    59494
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              13994432
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets         0
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    39971
TX Packets 256-511 Octets    668
TX Packets 512-1023 Octets   290
TX Packets 1024-1518 Octets  3
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms 0
TX Octets                    40933
TX Multicast Packets         0
TX Broadcast Packets         1
TX Single Collision frames    0
TX Mult. Collision frames     0
TX Late Collisions            0
TX Excessive Collisions       0
TX Collision frames           0
TX PAUSEMAC Ctrl Frames      0
TX MAC ctrl frames            0
TX Frame deferred Xmsns      0
TX Frame excessive deferl     0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0

```

TX Byte Counter	9050841
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	29767
RX Packets 256-511 Octets	225
RX Packets 512-1023 Octets	44
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	30039
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	7043738

Statistics for port 28 connected to device CB0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	0
TX Packets 128-255 Octets	88500
TX Packets 256-511 Octets	864
TX Packets 512-1023 Octets	163
TX Packets 1024-1518 Octets	6
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	89533
TX Multicast Packets	0
TX Broadcast Packets	0
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	21038170
RX Packets 64 Octets	0
RX Packets 65-127 Octets	1
RX Packets 128-255 Octets	120531
RX Packets 256-511 Octets	1947
RX Packets 512-1023 Octets	804
RX Packets 1024-1518 Octets	6
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	123289
RX Multicast Packets	0
RX Broadcast Packets	1
RX FCS Errors	0

```

RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter              27110675
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                    1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter              72
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   0
RX Packets 256-511 Octets   0
RX Packets 512-1023 Octets  0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                    0
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter              0

```

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7

show chassis ethernet-switch fpc detail slot

```
user@qfabric> show chassis ethernet-switch fpc detail 0  
re0:
```

```
-----  
Port statistics for FC0 switch  
Statistics for port 2 connected to device FWD-SWITCH-0:
```

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	121823
TX Packets 256-511 Octets	2200
TX Packets 512-1023 Octets	823
TX Packets 1024-1518 Octets	2
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	124849
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	27414524
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	85998
RX Packets 256-511 Octets	557
RX Packets 512-1023 Octets	86
RX Packets 1024-1518 Octets	1
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	86642
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0


```

RX False Carrier Errors      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             20398564
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   80443
TX Packets 256-511 Octets   1347
TX Packets 512-1023 Octets  532
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                   82326
TX Multicast Packets        0
TX Broadcast Packets        1
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames    0
TX MAC ctrl frames         0
TX Frame deferred Xmsns     0
TX Frame excessive deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             18161734
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   58460
RX Packets 256-511 Octets   523
RX Packets 512-1023 Octets  96
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                   59081
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter       0

```

```
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             13894171
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    0
TX Packets 128-255 Octets   144458
TX Packets 256-511 Octets   1080
TX Packets 512-1023 Octets  182
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   145723
TX Multicast Packets        0
TX Broadcast Packets        0
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             34292735
RX Packets 64 Octets        0
RX Packets 65-127 Octets    1
RX Packets 128-255 Octets   202266
RX Packets 256-511 Octets   3547
RX Packets 512-1023 Octets  1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   207174
RX Multicast Packets        0
RX Broadcast Packets        1
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             45576186
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets    0
TX Packets 256-511 Octets    0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
```

TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

Port statistics for FC1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80629
TX Packets 256-511 Octets	1279
TX Packets 512-1023 Octets	514
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82426
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18074790
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58785
RX Packets 256-511 Octets	640
RX Packets 512-1023 Octets	119

```
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 59547
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 14006842
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets 0
TX Packets 65-127 Octets 1
TX Packets 128-255 Octets 40004
TX Packets 256-511 Octets 668
TX Packets 512-1023 Octets 290
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 40966
TX Multicast Packets 0
TX Broadcast Packets 1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 9058102
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 29794
RX Packets 256-511 Octets 225
RX Packets 512-1023 Octets 44
RX Packets 1024-1518 Octets 3
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 30066
RX Multicast Packets 0
```

```

RX Broadcast Packets      0
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0
RX Unsupported opcodes    0
RX Out of Range Length    0
RX False Carrier Errors   0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter     0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           7050000
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  0
TX Packets 128-255 Octets 88579
TX Packets 256-511 Octets 865
TX Packets 512-1023 Octets 163
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                 89613
TX Multicast Packets      0
TX Broadcast Packets      0
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           21056842
RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 120633
RX Packets 256-511 Octets 1947
RX Packets 512-1023 Octets 804
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                 123391
RX Multicast Packets      0
RX Broadcast Packets      1
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           27132820
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1

```

```
TX Packets 128-255 Octets 0
TX Packets 256-511 Octets 0
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets 1
TX Multicast Packets 0
TX Broadcast Packets 1
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 72
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 0
RX Packets 256-511 Octets 0
RX Packets 512-1023 Octets 0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets 0
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 0
```

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7

show chassis ethernet-switch fpc interconnect-device port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 port 2
```

Summary for switch on FC0

Link is good on GE port 2 connected to device: FWD-SWITCH-0

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 319466

RX Octets 221869

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 210295

RX Octets 151164

Link is good on XE port 28 connected to device: CB0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 373033

RX Octets 529760

Link is good on XE port 29 connected to device: CB1

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 1

RX Octets 0

Summary for switch on FC1

Link is good on GE port 2 connected to device: FWD-SWITCH-0

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 210760

RX Octets 152617

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 104587

RX Octets 77315

Link is good on XE port 28 connected to device: CB0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

```
Flow Control RX is Disabled
TX Octets          229932
RX Octets          315346
```

```
Link is good on XE port 29 connected to device: CB1
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0
```

show chassis ethernet-switch fpc interconnect-device detail port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 detail port 2
```

```
Port statistics for FC0 switch
```

```
Statistics for port 2 connected to device FWD-SWITCH-0:
```

```
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 311974
TX Packets 256-511 Octets 5552
TX Packets 512-1023 Octets 2084
TX Packets 1024-1518 Octets 2
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                 319613
TX Multicast Packets      0
TX Broadcast Packets      1
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           70091196
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 220284
RX Packets 256-511 Octets 1486
RX Packets 512-1023 Octets 198
RX Packets 1024-1518 Octets 1
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                 221969
RX Multicast Packets      0
RX Broadcast Packets      0
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0
```



```

RX Unsupported opcodes      0
RX Out of Range Length     0
RX False Carrier Errors    0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            52192002
Statistics for port 4 connected to device FWD-SWITCH-1:
TX Packets 64 Octets       0
TX Packets 65-127 Octets   1
TX Packets 128-255 Octets  205595
TX Packets 256-511 Octets  3426
TX Packets 512-1023 Octets 1366
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                  210391
TX Multicast Packets       0
TX Broadcast Packets       1
TX Single Collision frames 0
TX Mult. Collision frames  0
TX Late Collisions         0
TX Excessive Collisions    0
TX Collision frames        0
TX PAUSEMAC Ctrl Frames    0
TX MAC ctrl frames         0
TX Frame deferred Xtns     0
TX Frame excessive deferl   0
TX Oversize Packets        0
TX Jabbers                 0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            46380018
RX Packets 64 Octets       0
RX Packets 65-127 Octets   0
RX Packets 128-255 Octets  149866
RX Packets 256-511 Octets  1194
RX Packets 512-1023 Octets 173
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                  151235
RX Multicast Packets       0
RX Broadcast Packets       0
RX FCS Errors              0
RX Align Errors            0
RX Fragments               0
RX Symbol errors           0
RX Unsupported opcodes     0
RX Out of Range Length     0
RX False Carrier Errors    0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0

```

```
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            35496911
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets       0
TX Packets 65-127 Octets   0
TX Packets 128-255 Octets  370150
TX Packets 256-511 Octets  2680
TX Packets 512-1023 Octets 371
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  373204
TX Multicast Packets       0
TX Broadcast Packets       0
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets        0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            87688913
RX Packets 64 Octets       0
RX Packets 65-127 Octets   1
RX Packets 128-255 Octets  517569
RX Packets 256-511 Octets  8978
RX Packets 512-1023 Octets 3450
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  530003
RX Multicast Packets       0
RX Broadcast Packets       1
RX FCS Errors              0
RX Fragments               0
RX MAC Control Packets     0
RX Out of Range Length     0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            116471142
Statistics for port 29 connected to device CB1:
TX Packets 64 Octets       0
TX Packets 65-127 Octets   1
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  1
TX Multicast Packets       0
TX Broadcast Packets       1
```

TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

show chassis fabric connectivity

Syntax	<code>show chassis fabric connectivity (device slot)</code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the status of the data plane connections in your QFabric system.
Options	<p>none—Display the status of all data plane connections in your QFabric system.</p> <p>device—Display the status of the data plane connections for a specific device.</p> <p>slot—Display the status of the data plane connections for a specific Flexible PIC Concentrator (FPC) slot on a specific device.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request chassis fabric fpc on page 1443 • show chassis fabric device on page 1527 • Understanding Interconnect Devices on page 1183
List of Sample Output	<p>show chassis fabric connectivity on page 1521</p> <p>show chassis fabric connectivity device on page 1524</p> <p>show chassis fabric connectivity device device-name slot on page 1525</p>
Output Fields	Table 143 on page 1520 lists the output fields for the show chassis fabric connectivity command. Output fields are listed in the approximate order in which they appear.

Table 143: show chassis fabric connectivity Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.
Type	Model number of the QFabric system component. Values include qfxc08-3008 (QFX3008-I Interconnect device), qfx3600-I (QFX3600-I Interconnect device), qfx3500 (QFX3500 Node device), and qfx3600-16q (QFX3600 Node device).
Fabric: Incoming links	Displays inbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as Ok).
Fabric: Outgoing links	Displays outbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as Ok).

Sample Output

show chassis fabric connectivity

```

user@qfabric> show chassis fabric connectivity
Device ID: ED1487, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/1          -> ED1487:fte-0/1/2          Ok
A0010:fte-3/0/1          -> ED1487:fte-0/1/3          Ok
Fabric: Outgoing links:
ED1487:fte-0/1/2         -> A0010:fte-0/0/1          Ok
ED1487:fte-0/1/3         -> A0010:fte-3/0/1          Ok
Device ID: ED3683, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/2          -> ED3683:fte-0/1/2          Ok
A0010:fte-3/0/2          -> ED3683:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3683:fte-0/1/2         -> A0010:fte-0/0/2          Ok
ED3683:fte-0/1/3         -> A0010:fte-3/0/2          Ok
Device ID: ED3705, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/0          -> ED3705:fte-0/1/2          Ok
A0010:fte-3/0/0          -> ED3705:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3705:fte-0/1/2         -> A0010:fte-0/0/0          Ok
ED3705:fte-0/1/3         -> A0010:fte-3/0/0          Ok
Device ID: ED3707, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/8          -> ED3707:fte-0/1/2          Ok
A0010:fte-3/0/8          -> ED3707:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3707:fte-0/1/2         -> A0010:fte-0/0/8          Ok
ED3707:fte-0/1/3         -> A0010:fte-3/0/8          Ok
Device ID: ED3711, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/9          -> ED3711:fte-0/1/2          Ok
A0010:fte-3/0/9          -> ED3711:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3711:fte-0/1/2         -> A0010:fte-0/0/9          Ok
ED3711:fte-0/1/3         -> A0010:fte-3/0/9          Ok
Device ID: ED3702, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/10         -> ED3702:fte-0/1/2          Ok
A0010:fte-3/0/10         -> ED3702:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3702:fte-0/1/2         -> A0010:fte-0/0/10         Ok
ED3702:fte-0/1/3         -> A0010:fte-3/0/10         Ok
Device ID: BBAK8737, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/11         -> BBAK8737:fte-0/1/1        Ok
Fabric: Outgoing links:
BBAK8737:fte-0/1/1       -> A0010:fte-0/0/11        Ok
Device ID: BBAK8777, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/5          -> BBAK8777:fte-0/1/0        Ok
A0010:fte-0/0/6          -> BBAK8777:fte-0/1/1        Ok
Fabric: Outgoing links:
BBAK8777:fte-0/1/0       -> A0010:fte-0/0/5          Ok
BBAK8777:fte-0/1/1       -> A0010:fte-0/0/6          Ok
Device ID: BBAK8866, Type: qfx3500
Fabric: Incoming links:

```

```

A0010:fte-0/0/3                -> BBAK8866:fte-0/1/0                Ok
A0010:fte-0/0/4                -> BBAK8866:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8866:fte-0/1/0            -> A0010:fte-0/0/3                Ok
BBAK8866:fte-0/1/1            -> A0010:fte-0/0/4                Ok
Device ID: BBAK8810, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/5                -> BBAK8810:fte-0/1/0                Ok
A0010:fte-3/0/6                -> BBAK8810:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8810:fte-0/1/0            -> A0010:fte-3/0/5                Ok
BBAK8810:fte-0/1/1            -> A0010:fte-3/0/6                Ok
Device ID: BBAK8854, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/3                -> BBAK8854:fte-0/1/0                Ok
A0010:fte-3/0/4                -> BBAK8854:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8854:fte-0/1/0            -> A0010:fte-3/0/3                Ok
BBAK8854:fte-0/1/1            -> A0010:fte-3/0/4                Ok
Device ID: BBAK8885, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/14              -> BBAK8885:fte-0/1/0                Ok
A0010:fte-0/0/15              -> BBAK8885:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8885:fte-0/1/0            -> A0010:fte-0/0/14                Ok
BBAK8885:fte-0/1/1            -> A0010:fte-0/0/15                Ok
Device ID: BBAK8864, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/7                -> BBAK8864:fte-0/1/0                Ok
A0010:fte-3/0/11              -> BBAK8864:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8864:fte-0/1/0            -> A0010:fte-3/0/7                Ok
BBAK8864:fte-0/1/1            -> A0010:fte-3/0/11               Ok
Device ID: BBAK8759, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/12              -> BBAK8759:fte-0/1/0                Ok
A0010:fte-3/0/13              -> BBAK8759:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8759:fte-0/1/0            -> A0010:fte-3/0/12                Ok
BBAK8759:fte-0/1/1            -> A0010:fte-3/0/13                Ok
Device ID: BBAK8704, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/12              -> BBAK8704:fte-0/1/0                Ok
A0010:fte-0/0/13              -> BBAK8704:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8704:fte-0/1/0            -> A0010:fte-0/0/12                Ok
BBAK8704:fte-0/1/1            -> A0010:fte-0/0/13                Ok
Device ID: BBAK8714, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/14              -> BBAK8714:fte-0/1/0                Ok
A0010:fte-3/0/15              -> BBAK8714:fte-0/1/1                Ok
Fabric: Outgoing links:
BBAK8714:fte-0/1/0            -> A0010:fte-3/0/14                Ok
BBAK8714:fte-0/1/1            -> A0010:fte-3/0/15                Ok
Device ID: A0010, Type: qfxc08-3008
Front Card 0 : Incoming links:
ED3705:fte-0/1/2              -> A0010:fte-0/0/0                Ok
ED1487:fte-0/1/2              -> A0010:fte-0/0/1                Ok
ED3683:fte-0/1/2              -> A0010:fte-0/0/2                Ok
BBAK8866:fte-0/1/0            -> A0010:fte-0/0/3                Ok
BBAK8866:fte-0/1/1            -> A0010:fte-0/0/4                Ok

```

```

BBAK8777:fte-0/1/0      -> A0010:fte-0/0/5      Ok
BBAK8777:fte-0/1/1      -> A0010:fte-0/0/6      Ok
ED3707:fte-0/1/2        -> A0010:fte-0/0/8      Ok
ED3711:fte-0/1/2        -> A0010:fte-0/0/9      Ok
ED3702:fte-0/1/2        -> A0010:fte-0/0/10     Ok
BBAK8737:fte-0/1/1      -> A0010:fte-0/0/11     Ok
BBAK8704:fte-0/1/0      -> A0010:fte-0/0/12     Ok
BBAK8704:fte-0/1/1      -> A0010:fte-0/0/13     Ok
BBAK8885:fte-0/1/0      -> A0010:fte-0/0/14     Ok
BBAK8885:fte-0/1/1      -> A0010:fte-0/0/15     Ok
Front Card 0 : Outgoing links:
A0010:fte-0/0/8          -> ED3707:fte-0/1/2      Ok
A0010:fte-0/0/9          -> ED3711:fte-0/1/2      Ok
A0010:fte-0/0/10         -> ED3702:fte-0/1/2      Ok
A0010:fte-0/0/11         -> BBAK8737:fte-0/1/1    Ok
A0010:fte-0/0/12         -> BBAK8704:fte-0/1/0    Ok
A0010:fte-0/0/13         -> BBAK8704:fte-0/1/1    Ok
A0010:fte-0/0/14         -> BBAK8885:fte-0/1/0    Ok
A0010:fte-0/0/15         -> BBAK8885:fte-0/1/1    Ok
A0010:fte-0/0/0          -> ED3705:fte-0/1/2      Ok
A0010:fte-0/0/1          -> ED1487:fte-0/1/2      Ok
A0010:fte-0/0/2          -> ED3683:fte-0/1/2      Ok
A0010:fte-0/0/3          -> BBAK8866:fte-0/1/0    Ok
A0010:fte-0/0/4          -> BBAK8866:fte-0/1/1    Ok
A0010:fte-0/0/5          -> BBAK8777:fte-0/1/0    Ok
A0010:fte-0/0/6          -> BBAK8777:fte-0/1/1    Ok
Front Card 3 : Incoming links:
ED3705:fte-0/1/3         -> A0010:fte-3/0/0      Ok
ED1487:fte-0/1/3         -> A0010:fte-3/0/1      Ok
ED3683:fte-0/1/3         -> A0010:fte-3/0/2      Ok
BBAK8854:fte-0/1/0       -> A0010:fte-3/0/3      Ok
BBAK8854:fte-0/1/1       -> A0010:fte-3/0/4      Ok
BBAK8810:fte-0/1/0       -> A0010:fte-3/0/5      Ok
BBAK8810:fte-0/1/1       -> A0010:fte-3/0/6      Ok
BBAK8864:fte-0/1/0       -> A0010:fte-3/0/7      Ok
ED3707:fte-0/1/3         -> A0010:fte-3/0/8      Ok
ED3711:fte-0/1/3         -> A0010:fte-3/0/9      Ok
ED3702:fte-0/1/3         -> A0010:fte-3/0/10     Ok
BBAK8864:fte-0/1/1       -> A0010:fte-3/0/11     Ok
BBAK8759:fte-0/1/0       -> A0010:fte-3/0/12     Ok
BBAK8759:fte-0/1/1       -> A0010:fte-3/0/13     Ok
BBAK8714:fte-0/1/0       -> A0010:fte-3/0/14     Ok
BBAK8714:fte-0/1/1       -> A0010:fte-3/0/15     Ok
Front Card 3 : Outgoing links:
A0010:fte-3/0/8          -> ED3707:fte-0/1/3      Ok
A0010:fte-3/0/9          -> ED3711:fte-0/1/3      Ok
A0010:fte-3/0/10         -> ED3702:fte-0/1/3      Ok
A0010:fte-3/0/11         -> BBAK8864:fte-0/1/1    Ok
A0010:fte-3/0/12         -> BBAK8759:fte-0/1/0    Ok
A0010:fte-3/0/13         -> BBAK8759:fte-0/1/1    Ok
A0010:fte-3/0/14         -> BBAK8714:fte-0/1/0    Ok
A0010:fte-3/0/15         -> BBAK8714:fte-0/1/1    Ok
A0010:fte-3/0/0          -> ED3705:fte-0/1/3      Ok
A0010:fte-3/0/1          -> ED1487:fte-0/1/3      Ok
A0010:fte-3/0/2          -> ED3683:fte-0/1/3      Ok
A0010:fte-3/0/3          -> BBAK8854:fte-0/1/0    Ok
A0010:fte-3/0/4          -> BBAK8854:fte-0/1/1    Ok
A0010:fte-3/0/5          -> BBAK8810:fte-0/1/0    Ok
A0010:fte-3/0/6          -> BBAK8810:fte-0/1/1    Ok
A0010:fte-3/0/7          -> BBAK8864:fte-0/1/0    Ok

```

show chassis fabric connectivity device

```

user@qfabric> show chassis fabric connectivity device BBAK8714
Device ID: BBAK8714, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/14          -> BBAK8714:fte-0/1/0          Ok
A0010:fte-3/0/15          -> BBAK8714:fte-0/1/1          Ok
Fabric: Outgoing links:
BBAK8714:fte-0/1/0        -> A0010:fte-3/0/14          Ok
BBAK8714:fte-0/1/1        -> A0010:fte-3/0/15          Ok

user@qfabric> show chassis fabric connectivity device A0010
Device ID: A0010, Type: qfxc08-3008
Front Card 0 : Incoming links:
ED3705:fte-0/1/2          -> A0010:fte-0/0/0          Ok
ED1487:fte-0/1/2          -> A0010:fte-0/0/1          Ok
ED3683:fte-0/1/2          -> A0010:fte-0/0/2          Ok
BBAK8866:fte-0/1/0        -> A0010:fte-0/0/3          Ok
BBAK8866:fte-0/1/1        -> A0010:fte-0/0/4          Ok
BBAK8777:fte-0/1/0        -> A0010:fte-0/0/5          Ok
BBAK8777:fte-0/1/1        -> A0010:fte-0/0/6          Ok
ED3707:fte-0/1/2          -> A0010:fte-0/0/8          Ok
ED3711:fte-0/1/2          -> A0010:fte-0/0/9          Ok
ED3702:fte-0/1/2          -> A0010:fte-0/0/10         Ok
BBAK8737:fte-0/1/1        -> A0010:fte-0/0/11         Ok
BBAK8704:fte-0/1/0        -> A0010:fte-0/0/12         Ok
BBAK8704:fte-0/1/1        -> A0010:fte-0/0/13         Ok
BBAK8885:fte-0/1/0        -> A0010:fte-0/0/14         Ok
BBAK8885:fte-0/1/1        -> A0010:fte-0/0/15         Ok
Front Card 0 : Outgoing links:
A0010:fte-0/0/8          -> ED3707:fte-0/1/2          Ok
A0010:fte-0/0/9          -> ED3711:fte-0/1/2          Ok
A0010:fte-0/0/10         -> ED3702:fte-0/1/2          Ok
A0010:fte-0/0/11         -> BBAK8737:fte-0/1/1        Ok
A0010:fte-0/0/12         -> BBAK8704:fte-0/1/0        Ok
A0010:fte-0/0/13         -> BBAK8704:fte-0/1/1        Ok
A0010:fte-0/0/14         -> BBAK8885:fte-0/1/0        Ok
A0010:fte-0/0/15         -> BBAK8885:fte-0/1/1        Ok
A0010:fte-0/0/0          -> ED3705:fte-0/1/2          Ok
A0010:fte-0/0/1          -> ED1487:fte-0/1/2          Ok
A0010:fte-0/0/2          -> ED3683:fte-0/1/2          Ok
A0010:fte-0/0/3          -> BBAK8866:fte-0/1/0        Ok
A0010:fte-0/0/4          -> BBAK8866:fte-0/1/1        Ok
A0010:fte-0/0/5          -> BBAK8777:fte-0/1/0        Ok
A0010:fte-0/0/6          -> BBAK8777:fte-0/1/1        Ok
Front Card 3 : Incoming links:
ED3705:fte-0/1/3          -> A0010:fte-3/0/0          Ok
ED1487:fte-0/1/3          -> A0010:fte-3/0/1          Ok
ED3683:fte-0/1/3          -> A0010:fte-3/0/2          Ok
BBAK8854:fte-0/1/0        -> A0010:fte-3/0/3          Ok
BBAK8854:fte-0/1/1        -> A0010:fte-3/0/4          Ok
BBAK8810:fte-0/1/0        -> A0010:fte-3/0/5          Ok
BBAK8810:fte-0/1/1        -> A0010:fte-3/0/6          Ok
BBAK8864:fte-0/1/0        -> A0010:fte-3/0/7          Ok
ED3707:fte-0/1/3          -> A0010:fte-3/0/8          Ok
ED3711:fte-0/1/3          -> A0010:fte-3/0/9          Ok
ED3702:fte-0/1/3          -> A0010:fte-3/0/10         Ok
BBAK8864:fte-0/1/1        -> A0010:fte-3/0/11         Ok
BBAK8759:fte-0/1/0        -> A0010:fte-3/0/12         Ok
BBAK8759:fte-0/1/1        -> A0010:fte-3/0/13         Ok
BBAK8714:fte-0/1/0        -> A0010:fte-3/0/14         Ok

```



```

BBAK8714:fte-0/1/1          -> A0010:fte-3/0/15          Ok
Front Card 3 : Outgoing links:
A0010:fte-3/0/8             -> ED3707:fte-0/1/3          Ok
A0010:fte-3/0/9             -> ED3711:fte-0/1/3          Ok
A0010:fte-3/0/10            -> ED3702:fte-0/1/3          Ok
A0010:fte-3/0/11            -> BBAK8864:fte-0/1/1        Ok
A0010:fte-3/0/12            -> BBAK8759:fte-0/1/0        Ok
A0010:fte-3/0/13            -> BBAK8759:fte-0/1/1        Ok
A0010:fte-3/0/14            -> BBAK8714:fte-0/1/0        Ok
A0010:fte-3/0/15            -> BBAK8714:fte-0/1/1        Ok
A0010:fte-3/0/0             -> ED3705:fte-0/1/3          Ok
A0010:fte-3/0/1             -> ED1487:fte-0/1/3          Ok
A0010:fte-3/0/2             -> ED3683:fte-0/1/3          Ok
A0010:fte-3/0/3             -> BBAK8854:fte-0/1/0        Ok
A0010:fte-3/0/4             -> BBAK8854:fte-0/1/1        Ok
A0010:fte-3/0/5             -> BBAK8810:fte-0/1/0        Ok
A0010:fte-3/0/6             -> BBAK8810:fte-0/1/1        Ok
A0010:fte-3/0/7             -> BBAK8864:fte-0/1/0        Ok
Ok
A0010:fte-3/0/3             -> BBAK8854:fte-0/1/0        Ok
A0010:fte-3/0/4             -> BBAK8854:fte-0/1/1        Ok
A0010:fte-3/0/5             -> BBAK8810:fte-0/1/0        Ok
A0010:fte-3/0/6             -> BBAK8810:fte-0/1/1        Ok
A0010:fte-3/0/7             -> BBAK8864:fte-0/1/0        Ok

```

show chassis fabric connectivity device device-name slot

```

user@qfabric> show chassis fabric connectivity device A0010 slot 3
Device ID: A0010, Type: qfxc08-3008
Front Card 3 : Incoming links:
ED3705:fte-0/1/3            -> A0010:fte-3/0/0          Ok
ED1487:fte-0/1/3            -> A0010:fte-3/0/1          Ok
ED3683:fte-0/1/3            -> A0010:fte-3/0/2          Ok
BBAK8854:fte-0/1/0          -> A0010:fte-3/0/3          Ok
BBAK8854:fte-0/1/1          -> A0010:fte-3/0/4          Ok
BBAK8810:fte-0/1/0          -> A0010:fte-3/0/5          Ok
BBAK8810:fte-0/1/1          -> A0010:fte-3/0/6          Ok
BBAK8864:fte-0/1/0          -> A0010:fte-3/0/7          Ok
ED3707:fte-0/1/3            -> A0010:fte-3/0/8          Ok
ED3711:fte-0/1/3            -> A0010:fte-3/0/9          Ok
ED3702:fte-0/1/3            -> A0010:fte-3/0/10         Ok
BBAK8864:fte-0/1/1          -> A0010:fte-3/0/11         Ok
BBAK8759:fte-0/1/0          -> A0010:fte-3/0/12         Ok
BBAK8759:fte-0/1/1          -> A0010:fte-3/0/13         Ok
BBAK8714:fte-0/1/0          -> A0010:fte-3/0/14         Ok
BBAK8714:fte-0/1/1          -> A0010:fte-3/0/15         Ok
Front Card 3 : Outgoing links:
A0010:fte-3/0/8             -> ED3707:fte-0/1/3          Ok
A0010:fte-3/0/9             -> ED3711:fte-0/1/3          Ok
A0010:fte-3/0/10            -> ED3702:fte-0/1/3          Ok
A0010:fte-3/0/11            -> BBAK8864:fte-0/1/1        Ok
A0010:fte-3/0/12            -> BBAK8759:fte-0/1/0        Ok
A0010:fte-3/0/13            -> BBAK8759:fte-0/1/1        Ok
A0010:fte-3/0/14            -> BBAK8714:fte-0/1/0        Ok
A0010:fte-3/0/15            -> BBAK8714:fte-0/1/1        Ok
A0010:fte-3/0/0             -> ED3705:fte-0/1/3          Ok
A0010:fte-3/0/1             -> ED1487:fte-0/1/3          Ok
A0010:fte-3/0/2             -> ED3683:fte-0/1/3          Ok
A0010:fte-3/0/3             -> BBAK8854:fte-0/1/0        Ok
A0010:fte-3/0/4             -> BBAK8854:fte-0/1/1        Ok
A0010:fte-3/0/5             -> BBAK8810:fte-0/1/0        Ok

```

A0010:fte-3/0/6	-> BBAK8810:fte-0/1/1	Ok
A0010:fte-3/0/7	-> BBAK8864:fte-0/1/0	Ok

show chassis fabric device

Syntax	<code>show chassis fabric device <i>device-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the fabric management status of devices in your QFabric system.
Options	<p>none—Display the fabric management status for all devices in your QFabric system.</p> <p>device-name—Display the fabric management status for a specific device in your QFabric system. You can enter either the alias name or the serial number of the device.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request chassis fabric fpc on page 1443 • show chassis fabric connectivity on page 1520 • show fabric administration inventory on page 1543 • Understanding Interconnect Devices on page 1183
List of Sample Output	show chassis fabric device on page 1527
Output Fields	<p>Table 144 on page 1527 lists the output fields for the show chassis fabric device command. Output fields are listed in the approximate order in which they appear.</p>

Table 144: show chassis fabric device Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.
Type	Model number of the QFabric system component. Values include qfxc08-3008 (QFX3008-I Interconnect device), qfx3600-I (QFX3600-I Interconnect device, qfx3500 (QFX3500 Node device), and qfx3600-16q (QFX3600 Node device).
Management status	Displays keepalive status for fabric management processes on a specific device. Values include On and Off .
Hardware status	Displays operational status of the device participating in fabric management (such as Ok).

Sample Output

show chassis fabric device

```

user@qfabric> show chassis fabric device

Device ID: node2, Type: qfx3500
Management status: On, Hardware status: Ok

```

Device ID: node3, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node0, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node1, Type: qfx3500
Management status: On, Hardware status: Ok

show chassis lcd

List of Syntax	show chassis lcd (EX Series) on page 1529 show chassis lcd (QFX Series and QFabric Systems) on page 1529
show chassis lcd (EX Series)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <menu <(all-members local member <i>member-id</i>)>></pre>
show chassis lcd (QFX Series and QFabric Systems)	<pre>show chassis lcd <fpc-slot <i>fpc-slot-number</i>> <interconnect-device <i>device-id</i>> <node-device <i>device-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>menu option introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.1 for QFabric systems.</p>
Description	<p>Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.</p>
Options	<p>none—Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.</p> <p>fpc-slot <<i>fpc-slot-number</i>>—(Optional) Display the information as follows:</p> <ul style="list-style-type: none"> (EX3200, EX3300, EX4200, and EX4500 switches, or the QFX Series) Display the information that appears on the LCD panel for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by fpc-slot 0. fpc-slot refers to the switch itself and 0 is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the none option. <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no <i>fpc-slot-number</i> value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the none option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the information for the specified member. <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> (EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the <i>fpc-slot-number</i> value.

Also display the status of the currently selected port parameter of the Status LED for each network port.

interconnect-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

menu—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member *member-id*—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

node-device *device-id*—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

Required Privilege Level

view

Related Documentation

- *LCD Panel in EX3200 Switches*
- *LCD Panel in EX4200 Switches*
- *LCD Panel in EX4500 Switches*
- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- [set chassis display message on page 1469](#)

List of Sample Output

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 1531](#)
[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 1533](#)
[show chassis lcd \(EX8200 Switch\) on page 1533](#)
[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 1535](#)
[show chassis lcd menu \(EX4200 Switch\) on page 1535](#)
[show chassis lcd menu \(EX8200 Switch\) on page 1535](#)
[show chassis lcd \(QFX3500 Switches\) on page 1536](#)
[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 1536](#)
[show chassis lcd interconnect-device \(QFabric Systems\) on page 1539](#)

[show chassis lcd node-device \(QFabric Systems\) on page 1541](#)

Output Fields [Table 64 on page 787](#) lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

Table 145: show chassis lcd Output Fields

Field Name	Field Description
membernumber (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
Front panel contents for slot	FPC slot number of the switch whose content is being displayed. The number is always 0 , except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
Front panel contents (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the member number field for each member device in the EX8200 Virtual Chassis.</p>
LCD screen	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays RE and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet (EX3200 and EX4200 switches only)
LEDs status	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: The XRE200 External Routing Engine always displays the NA parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • NA—Not applicable. • POE—Power over Ethernet
fpcx	On standalone EX Series and QFX Series switches, always 0 . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

Sample Output

show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
```

Front panel contents for slot: 0

LCD screen:
 00:BK switch1
 LED:SPD ALARM 00
LEDs status:
 Alarms LED: Off
 System LED: Green
 Master LED: Off
Interface LED(ADM/SPD/DPX/POE)

ge-0/0/0 Off
ge-0/0/1 Off
ge-0/0/2 Off
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Off
ge-0/0/6 Off
ge-0/0/7 Off
ge-0/0/8 Off
ge-0/0/9 Off
ge-0/0/10 Off
ge-0/0/11 Off
ge-0/0/12 Off
ge-0/0/13 Off
ge-0/0/14 Off
ge-0/0/15 Off
ge-0/0/16 Off
ge-0/0/17 Off
ge-0/0/18 Off
ge-0/0/19 Off
ge-0/0/20 Off
ge-0/0/21 Off
ge-0/0/22 Off
ge-0/0/23 Off

Front panel contents for slot: 1

LCD screen:
 01:RE switch2
 LED:SPD ALARM 01
LEDs status:
 Alarms LED: Yellow
 System LED: Green
 Master LED: Green
Interface LED(ADM/SPD/DPX/POE)

ge-1/0/0 Off
ge-1/0/1 Off
ge-1/0/2 Off
ge-1/0/3 Off
ge-1/0/4 Off
ge-1/0/5 Off
ge-1/0/6 Off
ge-1/0/7 Off
ge-1/0/8 Off
ge-1/0/9 Off
ge-1/0/10 Off
ge-1/0/11 Off
ge-1/0/12 Off
ge-1/0/13 Off
ge-1/0/14 Off


```

ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off

```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```

user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0       Off
ge-1/0/1       Off
ge-1/0/2       Off
ge-1/0/3       Off
ge-1/0/4       Off
ge-1/0/5       Off
ge-1/0/6       Off
ge-1/0/7       Off
ge-1/0/8       Off
ge-1/0/9       Off
ge-1/0/10      Off
ge-1/0/11      Off
ge-1/0/12      Off
ge-1/0/13      Off
ge-1/0/14      Off
ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off

```

show chassis lcd (EX8200 Switch)

```

user@switch> show chassis lcd
Front panel contents:
-----
LCD screen:
  RE st-8200-r
  LED:ADM ALARM 01

```

LEDs status:

Alarms LED: Yellow

System LED: Yellow

Master LED: Green

Interface	LED(ADM/SPD/DPX)
-----------	------------------

ge-0/0/0	Off
ge-0/0/1	Off
ge-0/0/2	Off
ge-0/0/3	Off
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	Off
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	Off
ge-0/0/41	Off
ge-0/0/42	Off
ge-0/0/43	Off
ge-0/0/44	Off
ge-0/0/45	Off
ge-0/0/46	Off
ge-0/0/47	Off
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off

```

xe-2/0/7      Off
xe-3/0/0      Off
xe-3/0/1      Off
xe-3/0/2      Off
xe-3/0/3      Off
xe-3/0/4      Off
xe-3/0/5      Off
xe-3/0/6      Off
xe-3/0/7      Off
xe-5/0/0      Off
xe-5/0/1      Off
xe-5/0/2      Off
xe-5/0/3      Off
xe-5/0/4      Off
xe-5/0/5      Off
xe-5/0/6      On
xe-5/0/7      On
xe-7/0/5      Off

```

show chassis lcd fpc-slot 2 (EX8200 Switch)

```
show chassis lcd fpc-slot 2
```

Interface	LED (ADM/SPD/DPX)
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
```

```

-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default

```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu** **all-members** command is the same as the output for the **show chassis lcd menu** command.

show chassis lcd menu (EX8200 Switch)

```

user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu

```

```
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off
```

show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis lcd
member0:
-----
Front panel contents:
-----
LCD screen:
  RE ex8200-member0
  LED:ADM ALARM 04
LEDs status:
  Alarms LED: Red
  System LED: Yellow
  Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
  BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
  09:RE xre200-member9
  LED: NA ALARM 01
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      On
ge-0/0/1      On
ge-0/0/2      On
ge-0/0/3      On
```

ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On
ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off
ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off

ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On
ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On
ge-22/0/17	Off
ge-22/0/18	On
ge-22/0/19	Off
ge-22/0/20	On
ge-22/0/21	Off
ge-22/0/22	On
ge-22/0/23	Off

```

ge-22/0/24      On
ge-22/0/25      Off
ge-22/0/26      Off
ge-22/0/27      Off
ge-22/0/28      Off
ge-22/0/29      Off
ge-22/0/30      Off
ge-22/0/31      Off
ge-22/0/32      On
ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off

```

show chassis lcd interconnect-device (QFabric Systems)

```

show chassis lcd interconnect-device IC-F1012
      Front Panel Module Information
      -----
      LCD screen:
      IC-F1012          3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off
  PEM 5 LED : off

      LED info for: CB - 0
      -----

LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface          STATUS LED    LINK/ACTIVITY LED

```

```

-----
IC-F1012:pme0 :      Green      N/A
IC-F1012:pme1 :      Green      N/A
IC-F1012:pme2 :      off        N/A
IC-F1012:pme3 :      off        N/A

```

LED info for: CB - 1

```

-----
LEDs status:
  Status LED: Green
  Mastership LED: Amber

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

```

-----
LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

LED info for: FC 1 FPC - 1

```

-----
LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

LED info for: RC 0 FPC - 8

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 1 FPC - 9

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 2 FPC - 10

```

-----
LEDs status:
  Status LED: Green

```

LED info for: RC 3 FPC - 11


```

-----
LEDs status:
  Status LED: Green

      LED info for: RC 4 FPC - 12
-----
LEDs status:
  Status LED: Green

      LED info for: RC 5 FPC - 13
-----
LEDs status:
  Status LED: Green

      LED info for: RC 6 FPC - 14
-----
LEDs status:
  Status LED: Green

      LED info for: RC 7 FPC - 15
-----
LEDs status:
  Status LED: Green

```

show chassis lcd node-device (QFabric Systems)

```

show chassis lcd node-device P3774-C
  Front panel contents for: P3774-C
  -----
  LCD screen:
  P3774-C
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

show chassis nonstop-upgrade node-group

Syntax	show chassis nonstop-upgrade node-group <i>node-group-name</i>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display the status of the Node group after the most recent nonstop software upgrade (NSSU).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 105 • request system software nonstop-upgrade on page 410
List of Sample Output	show chassis nonstop-upgrade node-group on page 1542
Output Fields	Table 68 on page 817 lists the output fields for the show chassis nonstop-upgrade node-group command. Output fields are listed in the approximate order in which they appear.

Table 146: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```

user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
Item           Status           Reason
P1550-C       Online

```

show fabric administration inventory

Syntax show fabric administration inventory
 <brief | detail | summary | terse>
 <director-group (status)>
 <infrastructure (fabric-controls | fabric-managers | diagnostic-routing-engines)>
 <interconnect-devices *interconnect-device-name*>
 <node-devices *node-device-name*>
 <node-groups *node-group-name*>
 <summary>

Release Information Command introduced in Junos OS Release 11.3 for the QFX Series.

Description (QFabric systems only) Display all devices that belong to the QFabric system. You can narrow the level of output by specifying a device type.



NOTE: If your Node devices do not appear in the output of the **show fabric administration inventory** command, check the cabling of your system.

Options none—Display all devices within a QFabric system.

brief | detail | summary | terse—(Optional) Display the specified level of output.

director-group (status)—(Optional) Display the status for the Director group within a QFabric system.

infrastructure (fabric-controls | fabric-managers | diagnostic-routing-engines)—(Optional) Display information for the fabric control Routing Engine, fabric manager Routing Engine, and diagnostic Routing Engine running on the Director group for the QFabric system.

interconnect-devices *interconnect-device-name*—(Optional) Display a specific Interconnect device within a QFabric system.

node-devices *node-device-name*—(Optional) Display a specific Node device within a QFabric system.

node-groups *node-group-name*—(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level admin

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)
- [show fabric administration inventory infrastructure on page 1553](#)

- [show fabric administration inventory director-group status on page 1548](#)
- [show fabric administration inventory node-devices on page 1558](#)
- [show fabric administration inventory node-groups on page 1560](#)
- [show fabric administration inventory interconnect-devices on page 1556](#)
- [show fabric inventory on page 1563](#)

List of Sample Output [show fabric administration inventory on page 1545](#)
[show fabric administration inventory detail on page 1546](#)
[show fabric administration inventory summary on page 1547](#)

Output Fields [Table 147 on page 1544](#) lists the output fields for the **show fabric administration inventory** command. Output fields are listed in the approximate order in which they appear.

Table 147: show fabric administration inventory Output Fields

Field Name	Field Description	Level of Output
Item	Type of QFabric system component being viewed. Possible values include Node group , Interconnect device , Fabric control , Fabric manager , Diagnostic routing engine , Director group , and Ungrouped Node device .	detail none
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.	detail none
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.	detail none
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (with details about the failure), Pending (in the process of being written or retried), or Unknown .	detail none
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.	detail none
Interconnect device	Name of the Interconnect devices associated with the Node group. The device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Interconnect device.	detail none
Fabric manager	Name of the primary virtual Junos Routing Engine associated with the QFabric system. The fabric manager Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric manager Routing Engine.	detail none
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.	detail none

Table 147: show fabric administration inventory Output Fields (*continued*)

Field Name	Field Description	Level of Output
Diagnostic routing engine	Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition. The diagnostic Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.	detail none
Director group	Identifier for the Director devices that are part of the Director group in a QFabric system. Each Director device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device.	detail none
Connection uptime	Length of time the component has been operational. The time is listed in days, hours, minutes, and seconds (D+HH:MM:SS).	detail
Network domain	Indicates whether a Node group is a network Node group (Yes) or a server Node group (No).	detail
Member id	Member identification number for a Node device within a Node group.	detail
Node group master	Indicates whether or not a Node device acts as a master device within a Node group. Values for this field are Yes or No .	detail
Configuration checkout failed	Provides troubleshooting details about a failed configuration on a component. This field includes the following output: <ul style="list-style-type: none"> • Edit path—Displays the configuration hierarchy level where the problem occurs. • Statement—Displays the configuration statement that causes the problem. • Message—Provides a suggested workaround to resolve the problem. 	detail
Member	Name of a Routing Engine allocated to an Interconnect device.	detail
Master	Indicates whether or not a Routing Engine acts as a master device within an Interconnect device. Values are Yes or No .	detail
Total Node devices	Number of connected and disconnected Node devices in the QFabric system.	summary
Total connected Node devices	Number of available Node devices in the QFabric system.	summary
Total disconnected Node devices	Number of unavailable Node devices in the QFabric system.	summary
Total Interconnect devices	Number of Interconnect devices in the QFabric system.	summary

Sample Output

show fabric administration inventory

```
user@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
------	------------	------------	---------------

Ungrouped Node device			
Node6	BBAK8979	Disconnected	
Node group			
P3359-C		Connected	Configured
P3359-C		Connected	
P3865-C		Connected	Configured
P3865-C		Connected	
RSNG-1		Connected	Configured
Node-3	BBAK8276	Connected	
Node-4	BBAK8273	Connected	
NW-NG-0		Connected	Configured
Node-0	BBAK8309	Connected	
Node-1	BBAK8283	Connected	
Interconnect device			
IC-F1032		Connected	Configured
F1032/RE0		Connected	
F1032/RE1		Connected	
IC-F1092		Connected	Configured
F1092/RE0		Connected	
F1092/RE1		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured
Director group			
0281112011000023		Connected	
0281112011000082		Connected	

show fabric administration inventory detail

```

user@qfabric> show fabric administration inventory detail
Node group: NW-NG-0, Connected, Configured
  Connection uptime: 2+22:40:46
  Network domain: Yes

Node device: node01, Connected,
  Member id: 0

Node group: RSNG, Connected, Configured
  Connection uptime: 1:20:22

Node device: node02, Connected,
  Member id: 0

Node device: node03, Connected,
  Member id: 1
  Node group master: Yes

Node group: BBAK0423, Connected, Failed (invalid configuration)
  Connection uptime: 0:01:06

Configuration checkout failed:
  Edit path: edit ethernet-switching-options
  Statement: analyzer
  Message: Vlan vlan1 which is configured as output vlan for analyzer session
an1 contains untagged interface.
  The analyzer output vlan should only contain tagged interface.

```

```
Node device: node0, Connected,  
  Member id: 0  
  Node group master: Yes  
  
Interconnect device: IC-F4912, Connected, Configured  
  Connection uptime: 0:08:05  
  
Member: F4912/RE0, Connected,  
  Connection uptime: 0:08:05  
  Master: Yes  
  
Fabric manager: FM-0, Connected, Configured  
  Connection uptime: 2+22:40:47  
  
Fabric control: FC-0, Connected, Configured  
  Connection uptime: 2+22:17:38  
  
Fabric control: FC-1, Connected, Configured  
  Connection uptime: 2+22:17:57  
  
Diagnostic routing engine: DRE-0, Connected, Configured  
  Connection uptime: 2+22:39:56
```

show fabric administration inventory summary

```
user@qfabric> show fabric administration inventory summary  
Total Node devices: 3  
Total connected Node devices: 3  
Total disconnected Node devices: 0  
Total Interconnect devices: 1
```

show fabric administration inventory director-group status

Syntax	show fabric administration inventory director-group status (target <i>director-device-name</i> all)
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the status of Director devices that belong to a QFabric system Director group.
Options	<p>none—Display the status of all Director devices within a QFabric system.</p> <p>all—Display the status of all Director devices within a QFabric system.</p> <p>target <i>director-device-name</i>—Display the status of a specific Director device within a QFabric system.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request fabric administration director-group change-master on page 1446 • Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335 • Understanding the Director Group on page 1180
List of Sample Output	<p>show fabric administration inventory director-group status on page 1549</p> <p>show fabric administration inventory director-group status target on page 1551</p>
Output Fields	Table 148 on page 1548 lists the output fields for the show fabric administration inventory director-group status command. Output fields are listed in the approximate order in which they appear.

Table 148: show fabric administration inventory director-group status Output Fields

Field Name	Field Description
Director Group Status	Timestamp for the Director group status report.
Member	Name of the Director device.
Status	Current operational mode of the Director device: online or offline .
Role	High availability operational role of the Director device: master or standby .
Mgmt Address	Management IP address of the Director device.
CPU	Percentage of CPU processing memory being used by the Director device.
Free Memory	Available storage memory on the Director device.

Table 148: show fabric administration inventory director-group status Output Fields (*continued*)

Field Name	Field Description
VMs	Number of virtual machines operating on the Director device. A Routing Engine issue exists on a particular Director device when the system displays a star (*) in the VMs column along with the following message: Error in retrieving VM Count in dgX .
Up Time	Length of time the Director device has been operating.
Device Id/Alias	Name or identifier of the Director device.
Master Services	Operational status of the database server, load balancer, and QFabric partition address.
Director Group Managed Services	Operational status of the shared file system, network file system, virtual machine server, and DHCP load balancer.
Hard Drive Status	Operational status of the Director device hard drive, including information about the volume identifier, physical identifiers, and SCSI identifiers. There is also status information for the drive partitions, including directory size, available and used drive space, utilization, and directory locations. A hard drive issue exists when the system displays one of the following messages: Error in retrieving Hard disk status or Error in retrieving Hard disk storage status .
Director Group Processes	Operational status of the Director group processes, such as device managers, SSH, NFS, FTP, system log messages, and SNMP.
Interface Link Status	Operational status of the Director device interfaces, such as the management interface, the control plane bridge interface, the control plane LAG, the control plane links (where the first number represents the Ethernet module [0 or 1] and the second number represents the port [0 - 3]), and the inter-Director crossover LAG. The state of the interfaces can be up or down . The Active designation indicates that a control plane link is active.

Sample Output

show fabric administration inventory director-group status

```

user@qfabric> show fabric administration inventory director-group status
Director Group Status Thu Aug  2 17:36:34 UTC 2012

Member Status Role      Mgmt Address      CPU Free Memory VMs Up Time
-----
dg0    online master  10.94.215.38      8% 15191684k    4 6 days, 06:24 hrs
dg1    online backup  10.94.215.39      7% 17733160k    3 6 days, 06:24 hrs

Member Device Id/Alias Status Role
-----
dg0    0281052011000001 online master

Master Services
-----
Database Server      online
Load Balancer Director online
QFabric Partition Address online

Director Group Managed Services
-----

```

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Physical ID:0	online
Physical ID:1	online

Size	Used	Avail	Used%	Mounted on
------	------	-------	-------	------------

423G	9.4G	391G	3%	/
99M	16M	79M	17%	/boot
93G	11G	83G	12%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up	
Control Plane Bridge	up	
Control Plane LAG	up	
CP Link [0/2]	up	Active
CP Link [0/1]	up	Active
CP Link [0/0]	up	Active
CP Link [1/2]	down	
CP Link [1/1]	down	
CP Link [1/0]	down	
Crossover LAG	up	
CP Link [0/3]	up	
CP Link [1/3]	up	

Member	Device Id/Alias	Status	Role
--------	-----------------	--------	------

dg1	0281052011000032	online	backup
-----	------------------	--------	--------

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

```
Physical ID:0          online
Physical ID:1          online
```

```
Size  Used Avail Used% Mounted on
----  -
423G  9.8G 391G   3%  /
99M   16M  79M  17% /boot
93G   11G  83G  12% /pbdata
```

Director Group Processes

```
-----
Director Group Manager      online
Partition Manager           online
Software Mirroring           online
Shared File System master   online
Secure Shell Process         online
Network File System          online
DHCP Server master           online    backup
FTP Server                   online
Syslog                       online
Distributed Management        online
SNMP Trap Forwarder          online
SNMP Process                 online
Platform Management          online
```

Interface Link Status

```
-----
Management Interface        up
Control Plane Bridge         up
Control Plane LAG            up
CP Link [0/2]                up      Active
CP Link [0/1]                up      Active
CP Link [0/0]                up      Active
CP Link [1/2]                down
CP Link [1/1]                down
CP Link [1/0]                down
Crossover LAG               up
CP Link [0/3]                up
CP Link [1/3]                up
```

show fabric administration inventory director-group status target

```
user@qfabric> show fabric administration inventory director-group status target
0281052011000004
```

```
Director Group Status Thu Aug 16 02:25:37 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.94.195.109	7%	10009364k	3	23:44 hrs
dg1	online	master	10.94.195.110	9%	6120712k	4	1 day, 33:17 mins

Member	Device Id/Alias	Status	Role
dg1	0281052011000004	online	master

Master Services

```
-----
Database Server              online
Load Balancer Director        online
QFabric Partition Address     online
```

Director Group Managed Services

```

-----
Shared File System          online
Network File System         online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Physical ID:0               online
Physical ID:1               online

```

Size Used Avail Used% Mounted on

```

-----
423G 7.2G 394G 2% /
99M 20M 75M 21% /boot
93G 8.8G 85G 10% /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System          online
DHCP Server master          online    master
FTP Server                  online
Syslog                      online
Distributed Management       online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge         up
Control Plane LAG           up
CP Link [0/2]               up    Active
CP Link [0/1]               up    Active
CP Link [0/0]               up    Active
CP Link [1/2]               up
CP Link [1/1]               up
CP Link [1/0]               up
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```

show fabric administration inventory infrastructure

Syntax	show administrator inventory infrastructure <brief detail> (fabric-controls fabric-managers diagnostic-routing-engines)
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the services running on the Director group for the QFabric system. These services can include external Routing Engines that are used to support QFabric system operations, such as partitioning and fabric control.
Options	<p>none—Display all services running on the Director group, which are used to support the QFabric system.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>fabric-managers—Display information for the fabric manager Routing Engine running on the Director group, which is used to support all partitions in the QFabric system.</p> <p>fabric-controls—Display information for the fabric control Routing Engine running on the Director group, which is used to support route information in the QFabric system.</p> <p>diagnostic-routing-engines—Display information for the diagnostic Routing Engine running on the Director group, which is responsible for troubleshooting and diagnostic utilities within a QFabric system partition.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • show fabric administration inventory on page 1543
List of Sample Output	show fabric administration inventory infrastructure on page 1554 show fabric administration inventory infrastructure fabric-controls on page 1555 show fabric administration inventory infrastructure fabric-managers on page 1555 show fabric administration inventory infrastructure diagnostic-routing-engines on page 1555
Output Fields	Table 149 on page 1553 lists the output fields for the show fabric administration inventory infrastructure command. Output fields are listed in the approximate order in which they appear.

Table 149: show fabric administration inventory infrastructure Output Fields

Field Name	Field Description
Routing Engine Type	Type of virtual Junos Routing Engine being viewed. Examples include the network Node group, fabric control, fabric manager, and diagnostic Routing Engines.
Hostname	Name of the QFabric system component.

Table 149: show fabric administration inventory infrastructure Output Fields (*continued*)

Field Name	Field Description
PID	Process identifier for the component.
CPU-Use (%)	Percentage of CPU processing memory being used by the component.
Fabric control	<p>Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition.</p> <p>With the fabric-controls option, the fabric control Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric control Routing Engine.</p>
Fabric manager	<p>Name of the virtual Junos Routing Engine that manages the QFabric system.</p> <p>With the fabric-managers option, the fabric manager Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric manager Routing Engine.</p>
Network Node group	Name of the virtual Junos Routing Engine instance that handles routing processes for a network Node group.
Diagnostic	<p>Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition.</p> <p>With the diagnostic-routing-engines option, the diagnostic Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.</p>
Item	Type of QFabric system component being viewed.
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory infrastructure

```

user@qfabric> show fabric administration inventory infrastructure
dg0:
Routing Engine Type      Hostname                      PID      CPU-Use(%)
-----
Fabric manager           FM-0                          9832      1.0
Network Node group       QFabric_default_NW-NG-1_RE1  24633     4.2
Fabric control           QFabric_default_FC-1_RE0     25374     1.8

```

Diagnostic	QFabric_DRE	6789	1.3
dgl:			
Routing Engine Type	Hostname	PID	CPU-Use(%)

Fabric manager	FM-1	572	1.6
Network Node group	QFabric_default_NW-NG-0_RE0	19217	7.8
Fabric control	QFabric_default_FC-0_RE0	20071	1.9

show fabric administration inventory infrastructure fabric-controls

```
user@qfabric> show fabric administration inventory infrastructure fabric-controls fabric-controls
```

Item	Identifier	Connection	Configuration
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric administration inventory infrastructure fabric-managers

```
user@qfabric> show fabric administration inventory infrastructure fabric-managers
```

Item	Identifier	Connection	Configuration
Fabric manager			
FM-0		Connected	Configured

show fabric administration inventory infrastructure diagnostic-routing-engines

```
user@qfabric> show fabric administration inventory infrastructure diagnostic-routing-engines
```

Item	Identifier	Connection	Configuration
Diagnostic routing engine			
DRE-0		Connected	Configured

show fabric administration inventory interconnect-devices

Syntax	<code>show fabric administration inventory interconnect-devices <i>interconnect-device-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the Interconnect devices that belong to a QFabric system.
Options	<p>none—Display all Interconnect devices within a QFabric system.</p> <p><i>interconnect-device-name</i>—Display a specific Interconnect device within a QFabric system.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request chassis fabric fpc on page 1443 • show chassis fabric connectivity on page 1520 • show chassis fabric device on page 1527 • show fabric administration inventory on page 1543 • Understanding Interconnect Devices on page 1183
List of Sample Output	<p>show fabric administration inventory interconnect-devices on page 1557</p> <p>show fabric administration inventory interconnect-devices device-name on page 1557</p>
Output Fields	Table 150 on page 1556 lists the output fields for the show fabric administration inventory interconnect-devices command. Output fields are listed in the approximate order in which they appear.

Table 150: show fabric administration inventory interconnect-devices Output Fields

Field Name	Field Description
Interconnect device	Name of the Interconnect devices associated with the partition.
Item	Type of QFabric system component being viewed. Interconnect devices either display the alias name (if configured) or the hardware serial identifier and Control Board Routing Engine numbers.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory interconnect-devices

```
user@qfabric> show fabric administration inventory interconnect-devices
Item                               Identifier           Connection           Configuration
Interconnect device
  IC-YW3781                        YW3781/RE0           Connected             Configured
  YW3781/RE1                       YW3781/RE1           Connected
  IC-YW3798                        YW3798/RE0           Connected             Configured
  YW3798/RE1                       YW3798/RE1           Connected
```

show fabric administration inventory interconnect-devices device-name

```
user@qfabric> show fabric administration inventory interconnect-devices IC-YW3781
Item                               Identifier           Connection           Configuration
Interconnect device
  IC-YW3781                        YW3781/RE0           Connected             Configured
  YW3781/RE1                       YW3781/RE1           Connected
```

show fabric administration inventory node-devices

Syntax `show fabric administration inventory node-devices node-device-name`

Release Information Command introduced in Junos OS Release 11.3 for the QFX Series.

Description (QFabric systems only) Display the Node devices that belong to the QFabric system.



NOTE: If your Node devices do not appear in the output of the `show fabric administration inventory node-devices` command, check the cabling of your system.

Options **none**—Display all Node devices within the QFabric system.

node-device-name—Display a specific Node device within the QFabric system.

Required Privilege Level admin

- Related Documentation**
- [Configuring Aliases for the QFabric System on page 1353](#)
 - [Configuring Node Groups for the QFabric System on page 1363](#)
 - [show fabric administration inventory node-groups on page 1560](#)
 - [show fabric administration inventory on page 1543](#)
 - [Understanding Node Devices on page 1186](#)

List of Sample Output [show fabric administration inventory node-devices on page 1559](#)
[show fabric administration inventory node-devices device-name on page 1559](#)

Output Fields [Table 151 on page 1558](#) lists the output fields for the `show fabric administration inventory node-devices` command. Output fields are listed in the approximate order in which they appear.

Table 151: show fabric administration inventory node-devices Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed.
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either Connected or Not Connected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Node device	Name of the Node devices associated with the Node group. The device can be either Connected or Not Connected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.

Sample Output

show fabric administration inventory node-devices

```
user@qfabric> show fabric administration inventory node-devices
Item                Identifier      Connection
Node device
node1               P3749-C       Connected
node2               P3767-C       Connected
node3               P3850-C       Connected
node4               P3947-C       Connected
```

show fabric administration inventory node-devices device-name

```
user@qfabric> show fabric administration inventory node-devices node0
Item                Identifier      Connection
Node device
node1               P3749-C       Connected
```

show fabric administration inventory node-groups

Syntax	<code>show fabric administration inventory node-groups <i>node-group-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the Node groups and the corresponding Node devices that belong to the QFabric system.
Options	<p>none—Display all Node groups within the QFabric system.</p> <p><i>node-group-name</i>—Display information for a specific Node group within the QFabric system.</p>
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • Configuring Node Groups for the QFabric System on page 1363 • show fabric administration inventory node-devices on page 1558 • show fabric administration inventory on page 1543
List of Sample Output	<p>show fabric administration inventory node-groups on page 1561</p> <p>show fabric administration inventory node-groups node-group-name on page 1561</p>
Output Fields	Table 152 on page 1560 lists the output fields for the show fabric administration inventory node-groups command. Output fields are listed in the approximate order in which they appear.

Table 152: show fabric administration inventory node-groups Output Fields

Field Name	Field Description
Node group	Name of the Node groups associated with the partition.
Item	<p>Type of QFabric system component being viewed.</p> <ul style="list-style-type: none"> • Autogenerated Node groups display the hardware serial identifier for both the name of the Node group and the name of the included Node device. • User-configured Node groups either display the alias name (if configured) or the hardware serial identifier for each Node device contained in the Node group.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory node-groups

```

user@qfabric> show fabric administration inventory node-groups
Item                               Identifier      Connection      Configuration
Node group
  BBAK8891                         Connected      Configured
  BBAK8891                         Connected
  BBAK8868                         Connected      Configured
  BBAK8868                         Connected
  RSNG-1                           Connected      Configured
  Node-3                           BBAK8276      Connected
  Node-4                           BBAK8273      Connected
  NW-NG-0                           Connected      Configured
  Node-0                           BBAK8309      Connected
  Node-1                           BBAK8283      Connected

```

show fabric administration inventory node-groups node-group-name

```

user@qfabric> show fabric administration inventory node-groups RSNG-1
Item                               Identifier      Connection      Configuration
Node group
  RSNG-1                           Connected      Configured
  Node-3                           BBAK8276      Connected
  Node-4                           BBAK8273      Connected

```

show fabric administration system mac-pool

Syntax	show fabric administration system mac-pool
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display the MAC addresses that belong to a QFabric Director group.
Options	There are no options for this command.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none"> • request fabric administration system mac-pool add on page 1449 • request fabric administration system mac-pool delete on page 1450
List of Sample Output	show fabric administration system mac-pool on page 1562
Output Fields	Table 153 on page 1562 lists the output fields for the show fabric administration system mac-pool command. Output fields are listed in the approximate order in which they appear.

Table 153: show fabric administration system mac-pool Output Fields

Field Name	Field Description
MAC Block Base	Starting MAC address for the pool assigned to the QFabric system.
Total MACs	Total number of MAC addresses assigned to the QFabric system.
Available MACs	Number of available MAC addresses from the total.

Sample Output

show fabric administration system mac-pool

```

user@qfabric> show fabric administration system mac-pool
  Mac Block Base      Total MACs      Available MACs
  00:11:00:00:00:00    4096           4084
  02:00:00:11:22:00     10             10

```

show fabric inventory

Syntax `show fabric inventory`
`<brief | detail | terse>`
`<infrastructure fabric-controls <FC-0 | FC-1>>`
`<node-devices node-device-name>`
`<node-groups node-group-name>`

Release Information Command introduced in Junos OS Release 12.2 for the QFX Series.

Description (QFabric systems only) Display Node devices, Node groups, and fabric control Routing Engines that belong to the QFabric system. You can narrow the level of output by specifying a device type.



NOTE:

- If you have administrator privileges, issue the `show fabric administration inventory` command to view all devices in your QFabric system (including Interconnect devices and Director devices).
- If your Node devices do not appear in the output of the `show fabric inventory` command, check the cabling of your system.

Options `none`—Display all devices within a QFabric system.

`brief | detail | terse`—(Optional) Display the specified level of output.

`infrastructure fabric-controls <FC-0 | FC-1>`—(Optional) Display information for all fabric control Routing Engines running on the Director group within the QFabric system, or the individual fabric control Routing Engine you specify (either FC-0 or FC-1).

`node-devices node-device-name`—(Optional) Display a specific Node device within a QFabric system.

`node-groups node-group-name`—(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level admin

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Configuring Aliases for the QFabric System on page 1353](#)
- [Configuring Node Groups for the QFabric System on page 1363](#)
- [show fabric administration inventory on page 1543](#)

List of Sample Output [show fabric inventory on page 1564](#)
[show fabric inventory infrastructure fabric-controls on page 1565](#)

[show fabric inventory node-devices on page 1565](#)

[show fabric inventory node-groups on page 1565](#)

Output Fields Table 154 on page 1564 lists the output fields for the **show fabric inventory** command. Output fields are listed in the approximate order in which they appear.

Table 154: show fabric inventory Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed. Possible values include Node device , Node group , Fabric control , and Ungrouped Node device .
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.
Node device	Name of the Node devices associated with the Node group. The device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.

Sample Output

show fabric inventory

```
user@qfabric> show fabric inventory
```

Item	Identifier	Connection	Configuration
Ungrouped Node device			
P3747-C		Disconnected	
Node group			
NW-NG-0		Connected	Configured
Node-1	P4093-C	Connected	
RSNG-1		Connected	Configured
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric inventory infrastructure fabric-controls

```
user@qfabric> show fabric inventory infrastructure fabric-controls
```

Item	Identifier	Connection	Configuration
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric inventory node-devices

```
user@qfabric> show fabric inventory node-devices
```

Item	Identifier	Connection	Configuration
Node device			
Node-1	P4093-C	Connected	
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	

show fabric inventory node-groups

```
user@qfabric> show fabric inventory node-groups
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
Node-1	P4093-C	Connected	
RSNG-1		Connected	Configured
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	

show fabric session-host

Syntax	show fabric session-host
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	(QFabric systems only) Display the Director device within the Director group that hosts the QFabric CLI session.
Options	none —Display the Director device hosting the QFabric CLI session.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none">• Understanding the Director Group on page 1180• show fabric administration inventory director-group status on page 1548
List of Sample Output	show fabric session-host on page 1566
Output Fields	Table 155 on page 1566 lists the output fields for the show fabric session-host command. Output fields are listed in the approximate order in which they appear.

Table 155: show fabric session-host Output Fields

Field Name	Field Description
Identifier	Hardware serial identifier of the Director device that hosts the SSH QFabric CLI session.

Sample Output

show fabric session-host

```
user@qfabric> show fabric session-host
Identifier: 0281052011000032
```

show log

List of Syntax	Syntax on page 1567 Syntax (QFabric System) on page 1567 Syntax (TX Matrix Routers) on page 1567
Syntax	<pre>show log <filename user <username>></pre>
Syntax (QFabric System)	<pre>show log filename <device-type (device-id device-alias)></pre>
Syntax (TX Matrix Routers)	<pre>show log <all-lcc lcc number scc> <filename user <username>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p>
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.
Options	<p>none—List all log files.</p> <p><all-lcc lcc number scc>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p>device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p> <ul style="list-style-type: none"> • director-device—Display logs for Director devices. • infrastructure-device—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup). • interconnect-device—Display logs for Interconnect devices. • node-device—Display logs for Node devices.



NOTE: If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

Required Privilege Level trace

List of Sample Output [show log on page 1568](#)
[show log filename on page 1568](#)
[show log filename \(QFabric System\) on page 1569](#)
[show log user on page 1569](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

show log user

```

user@host> show log user
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

show system software upgrade status

Syntax	show system software upgrade status
Release Information	Command introduced in Junos OS Release 13.1 for the QFX Series.
Description	(QFabric systems only) Display the status of a software upgrade, including details for both nonstop software upgrades and <i>component all</i> style upgrades.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 105 • Verifying Nonstop Software Upgrade for QFabric Systems on page 1410 • Upgrading Software on a QFabric System on page 134 • request system software nonstop-upgrade on page 410
List of Sample Output	show system software upgrade status (Nonstop Software Upgrade) on page 1570 show system software upgrade status (Component All Upgrade) on page 1571
Output Fields	Table 156 on page 1570 lists the output fields for the show system software upgrade status command. Output fields are listed in the approximate order in which they appear.

Table 156: show system software upgrade status Output Fields

Field Name	Field Description
Timestamp	Displays the day of the week, month, date, hour, minute, second, and year when you issue the show system software upgrade status command. An example of the timestamp format is as follows: Wed Jan 16 22:06:02 2013.
Software nonstop upgrade on:	Status of the upgrade: <ul style="list-style-type: none"> • FM-0 in progress—A Director group nonstop software upgrade is in process for the fabric manager Routing Engine. • NW-NG-0 in progress—A Node group nonstop software upgrade is in process for the network Node group. • RSNG in progress—A Node group nonstop software upgrade is in process for the redundant server Node group. • all in progress—A <i>component all</i> style upgrade is in process for the entire QFabric system.

Sample Output

show system software upgrade status (Nonstop Software Upgrade)

```

user@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
                        NW-NG-0 in progress
                        RSNG in progress

```

show system software upgrade status (Component All Upgrade)

```
user@qfabric> show system software upgrade status
Wed Jan 16 22:37:48 2013 Software component upgrade on:
all in progress
```


CHAPTER 12

Troubleshooting

- [QFabric System Troubleshooting on page 1573](#)

QFabric System Troubleshooting

- [Performing System Backup and Recovery for a QFabric System on page 1573](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 1574](#)
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 1581](#)
- [Creating an Emergency Boot Device for a QFX Series Device on page 1583](#)

Performing System Backup and Recovery for a QFabric System

Many routers and switches require an administrator to recover the software package and the configuration file for the device separately. In the case of a device failure, this means the administrator might need to perform two separate tasks (if neither the software package nor the configuration file can be recovered).

In contrast, the QFabric system uses a unique mechanism that saves the backup and recovery files for both the Junos OS software and the system configuration into a single collection. The following QFabric system backup and recovery mechanism simplifies and streamlines the recovery process so you can return to normal operations as quickly as possible.

To backup and recover your QFabric system:

1. (First time only) Implement the following one-time procedure to prepare your QFabric system to use the system backup and recovery feature:
 - Insert a Juniper Networks software installation USB flash drive into the master Director device. (This drive was provided to you as one of the components of your QFabric system shipment.)
 - Issue the **request system software format-qfabric-backup** command. The contents and format of the USB flash drive are copied to the Director group shared directory and are used as the basis for all future backup and recovery operations.

```
user@qfabric> request system software format-qfabric-backup
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```

- Remove the Juniper Networks software installation USB drive from the master Director device.
- 2. Issue the **request system software system-backup** command to backup the software package and configuration file. This command saves the current files necessary to recover the QFabric system. The files are saved to a shared memory directory in the Director group.



NOTE: As you upgrade your system with new software and change the system configuration over time, remember to reissue this command periodically to save the newest files for recovery purposes.

```
user@qfabric> request system software system-backup
```

```
user@qfabric>
```

- 3. Insert a 4 GB or larger USB flash drive into the master Director device for your Director group, and issue the **request system software system-backup usb-create** command. This command copies the recovery files that have been backed up in the Director group and transfers them to the USB flash drive to create a recovery USB drive.



NOTE: Issuing this command overwrites the contents of the USB flash drive with the QFabric system recovery files.

```
user@qfabric> request system software system-backup usb-create /dev/sdb
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```

- 4. Remove the recovery USB drive from the Director device, and store it securely in a known location that you will remember when you need to use the recovery USB drive.
- 5. If the QFabric system fails, power off the Director group, insert the recovery USB drive into the master Director device of your Director group, turn on power to the Director device, and follow the prompts to recover your system. This step restores the software package and the configuration file for your QFabric system.

Related Documentation

- [request system software format-qfabric-backup on page 1460](#)
- [request system software system-backup on page 1467](#)

Performing a QFabric System Recovery Installation on the Director Group

If the software on your QFabric system is damaged in some way that prevents the software from loading correctly, or you need to upgrade the software on your QFabric system, you may need to perform a recovery installation on the Director group.

If possible, perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device (for example, an external USB flash drive) for each of your Director devices to use during the recovery installation.

You can either use the external USB flash drive containing the software supplied by Juniper Networks, or you can use an external USB flash drive supplied by Juniper Networks on which you install the QFabric system install media.

2. Because the recovery installation process completely overwrites the entire contents of the Director device, make sure you back up any configuration files and initial setup information on a different external USB flash drive before you begin a recovery installation. You will need to restore this information as part of recovery process.

Use the **request system software configuration-backup** command to back up your configuration files and initial setup information:

```
user@switch> request system software configuration-backup path
```



NOTE: To recover the Director group, you must upgrade both Director devices in parallel. If you are recovering only one Director device in a Director group, and the software version will remain the same between the two Director devices, make sure that the other Director device is powered on and operational. If the software version of the Director device you are recovering will be different, make sure that the other Director device is powered off and is not operational.

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive on page 1575
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software on page 1577

(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive

If you do not have an external USB flash drive preloaded with the software from Juniper Networks to use as an emergency boot device, you can create your own, using a blank external USB flash drive provided by Juniper Networks. Download the install media from the Juniper Networks Support website onto your UNIX workstation, uncompress and untar the software, and then burn the software image onto your Juniper Networks external USB (4-gigabyte) flash drive. Make sure you create two emergency boot devices, one for each Director device, so you can perform a recovery installation in parallel.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.
4. In the **QFX Series** section, click the name of the platform for which you want to download software.

5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. Select the complete install media you want to download in the **QFabric System Install Media** section.
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Log in and save the install media file to your UNIX workstation.
10. Use FTP to access the UNIX workstation where the install media resides.

```
ftp ftp://hostname/pathname install-media-qfabric-<version>.img.tgz
```
11. When prompted, enter your username and password.
12. Make sure you are in binary mode by entering **binary** at the prompt.

```
binary
```
13. Use the **get** command to transfer the installation package from the FTP host to your UNIX workstation.

```
get install-media-qfabric-<version>.img.tgz
```
14. Close the FTP session:

```
bye
```
15. Untar the *install-media-qfabric-<version>.img.tgz* file on your UNIX workstation.

```
tar -xvzf install-media-qfabric-11.3X30.6.img.tgz
```
16. Insert a blank external USB (4-gigabyte) flash drive supplied by Juniper Networks into your UNIX workstation.
17. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive using the **dd** command:

```
dd if=install-media-qfabric-11.3X30.6.img of=/dev/sdb bs=16k
250880+0 records in
250880+0 records out
4110417920 bytes (4.1 GB) copied, 5.10768 seconds, 805 MB/s
```
18. Perform the steps in [“Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software” on page 115](#) to continue with the recovery installation.

Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software

This procedure describes how to perform a recovery installation using an external USB flash drive that contains Junos OS software.



NOTE: Since the recovery installation process completely overwrites the entire contents of the Director device, you will need to restore the required configuration files and initial setup information. The following procedure assumes you previously saved these backup files with the **request system software configuration-backup** command. Ensure that you have these backup files available on an external USB flash drive before you perform the following steps.

1. Insert the external USB flash drive into the Director device.
2. Perform one of the following tasks:
 - If you have access to the default partition, reboot the Director device by issuing the **request system reboot director-group** command.
 - If you do not have access to the default partition, power cycle the Director device.

The following menu appears on the Director device console when the Director device boots up:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

3. Type **install** and then press **Enter** to install the software on the Director device.

Once the installation process is complete, the Director device reboots, and the following menu appears on the Director device console:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

4. Press **Enter**.

The Director device reboots from the local disk on which the software was just installed.

5. Log in as root on the Director device.

The following menu appears on the Director device console:

```
Before you can access the QFabric system, you must complete the initial setup
of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of the
Director device and then log back in to restart
this setup menu.
```

```
Continue?[y/n]
```

6. Enter **n** to bypass the initial setup script and enter the Director device root directory, where you can mount the external USB flash drive containing the configuration files and initial setup information.

7. Issue the **ls /mnt** command to list the **mount** directory.

```
root@dg0 ~]# ls /mnt
```

8. Issue the **mkdir** command to create a directory within the mount directory.

```
root@dg0 ~]# mkdir /mnt/myusb
```

9. Issue the **mount /dev/sdb2 /mnt/myusb/** command to mount the external USB flash drive to the local drive of the Director device.

```
root@dg0 ~]# mount /dev/sdb2 /mnt/myusb/
```

10. Issue the **ls -la /mnt/myusb/** command to verify the contents of your mounted external USB flashdrive.

```
root@dg0 ~]# ls -la /mnt/myusb/
total 1770884
drwxr-xr-x 2 root root      4096 Sep  7 05:16 .
drwxr-xr-x 3 root root      4096 Sep  7 10:15 ..
-rw-r--r-- 1 root root    4249 Sep  7 03:52 mybackup-20110907
```

11. Exit the Director device and log back in as root on the Director device.

The following menu appears:

Before you can access the QFabric system, you must complete the initial setup of the Director group by using the steps that follow.

If the initial setup procedure does not complete successfully, log out of the Director device and then log back in to restart this setup menu.

```
Continue?[y/n] y
Initial Configuration
```

You may enter the configuration manually or restore from a backup.

```
Specify a backup file? [y/n] : y
Please specify the full path of the configuration backup file. :
/mnt/myusb/mybackup-20110907
```

12. Enter **y** to continue.

13. Enter **y** and specify the path to the backup configuration file located on the external USB flash drive.

```
/mnt/myusb/mybackup-20110907
```

The following messages appear:

```
Saving temporary configuration...
Configuring peer...
connect error for 1.1.1.2:9001
Configuring local interfaces...
Configuring interface eth0 with [10.49.213.163/24:10.49.213.254]
Configured interface eth0 with [10.49.213.163/24:10.49.213.254]
Configuring QFabric software with initial pool of 4000 MAC addresses
[00:10:00:00:00:00 - 00:10:00:00:0f:3b]
Configuring QFabric address [10.49.213.50]
Reconfiguring QFabric software static configuration
Applying the new Director Device password
Applying the QFabric component password
```

First install initial configuration, generating and sharing SSH keys.
 First install initial configuration, generating SSH keys.
 connect error for 1.1.1.2:9001
 Shared SSH keys.
 Configuration complete. Director Group services will auto start within 30 seconds.

The Director device reboots from the local disk on which the software was just installed.
 Exit the Director device session and log in to the QFabric default partition CLI.

14. Issue the **request system software configuration-restore** command and specify the path to the backup configuration file located on the external USB flash drive to load the previously saved QFabric system configuration.

15. From the default partition, issue the **request system reboot node-group all** command to reboot all of the Node groups in the QFabric system to ensure that all Node devices are running the same version of software as the Director-group.

```
user@switch> request system reboot node-group all
```

16. From the default partition, issue the **request system reboot fabric** command to reboot the Interconnect devices and the other components in the fabric in the QFabric system to ensure that Interconnect devices are running the same version of software as the Director group.

```
user@switch> request system reboot fabric
```

17. Log in to the default partition and issue the **show version component all** command to verify that all components are running the same version of software.

```
user@switch> show version component all
```

```
dg1:
```

```
-
```

```
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]
```

```
dg0:
```

```
-
```

```
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]
```

```
NW-NG-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

```
FC-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
```

JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-1:

Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

DRE-0:

-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FM-0:

-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

nodedevice1:

-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

interconnectdevice1:

-
Hostname: qfabric


```

Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
warning: from interconnectdevice0: Disconnected

```

Related Documentation

- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
- [Upgrading Software on a QFabric System on page 134](#)
- [request system software configuration-backup on page 402](#)
- [request system software configuration-restore on page 403](#)

Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device for a QFX Series Device” on page 165](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



WARNING: The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the QFX Series device.
2. Reboot the QFX Series device.



NOTE: Do not power off the device if it is already on.

[edit system]
user@device> **request system reboot**

If you do not have access to the CLI, power cycle the QFX Series device.

The emergency boot device (external USB install media) is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following options:

```
External USB install media detected.  
You can load Junos from this media onto an internal drive.  
Press 'y' to proceed, 'f' to format and install, or 'n' to abort.  
Do you wish to continue ([y]/f/n)? f
```

4. Type **f** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.

If you do not want to format the internal flash storage, type **y**.

The following messages are displayed:

```
Installing packages from external USB drive da1  
Packages will be installed to da0, media size: 8G
```

```
Processing format options  
Fri September 4 01:18:44 UTC 2012
```

```
-- IMPORTANT INFORMATION --  
Installer has detected settings to format system boot media.  
This operation will erase all data from your system.
```

```
Formatting installation disk .. this will take a while, please wait  
Disabling platform watchdog - threshold 12 mins
```

```
Determining installation slice  
Fri September 4 01:27:07 UTC 2012
```

5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the device is finished copying the software, you are presented with the following prompt:

```
*** Fri September 4 01:19:00 UTC 2012***  
Installation successful..  
Please select one of the following options:  
Reboot to installed Junos after removing install media (default) ... 1  
Reboot to installed Junos by disabling install media ..... 2  
Exit to installer debug shell ..... 3  
Install Junos to alternate slice ..... 4  
Your choice: 4  
NOTE: System installer will now install Junos to alternate slice  
Do not power off or remove the external installer media or  
interrupt the installation mechanism.
```

6. Select **4** to install Junos OS to the alternate slice of the partition, and then press Enter.

7. Remove the emergency boot device when prompted and then press Enter. The device then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the login prompt.
8. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.

**Related
Documentation**

- [Creating an Emergency Boot Device for a QFX Series Device on page 165](#)

Creating an Emergency Boot Device for a QFX Series Device

If Junos OS on the QFX Series is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



NOTE: In the following procedure, we assume that you are creating the emergency boot device on a QFX3500 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device from a QFX3500 device:

1. Use FTP to copy the installation media image into the `/var/tmp` directory on the QFX3500 device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```



NOTE: The password is the root password for the QFX3500 device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
```

```
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

**Related
Documentation**

- *USB Port Specifications for the QFX Series*
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)

PART 6

Configuration and File Management

- [Overview on page 1587](#)
- [Configuration on page 1593](#)
- [Administration on page 1629](#)
- [Troubleshooting on page 1657](#)

CHAPTER 13

Overview

- [Configuration Files Overview on page 1587](#)
- [Software Overview on page 1588](#)

Configuration Files Overview

- [Configuration File Terms on page 1587](#)

Configuration File Terms

[Table 33 on page 49](#) lists the various configuration file terms used for the QFX Series and their definitions.

Table 157: Configuration File Terms

Term	Definition
active configuration	Current committed configuration of a switch.
candidate configuration	Working copy of the configuration that allows users to make configurational changes without causing any operational changes until this copy is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Check configuration for proper syntax, activate and mark as the current configuration file running on the switching platform.
configuration hierarchy	Junos OS configuration consists of a hierarchy of statements. There are two types of statements: container statements, which contain other statements, and leaf statements, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
default configuration	Default configuration contains the initial values set for each configuration parameter when a switch is shipped.
rescue configuration	Well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the CLI.
roll back a configuration	Return to a previously committed configuration.

- Related Documentation**
- [Loading a Previous Configuration File on page 1600](#)
 - [Reverting to the Rescue Configuration on page 175](#)
 - [Understanding Configuration Files on page 1590](#)

Software Overview

- [Forms of the configure Command on page 1588](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 1589](#)
- [Understanding Configuration Files on page 1590](#)
- [Understanding How the Junos Configuration Is Stored on page 1591](#)

Forms of the configure Command

The Junos OS supports three forms of the **configure** command: **configure**, **configure private**, and **configure exclusive**. These forms control how users edit and commit configurations and can be useful when multiple users configure the software. See [Table 158 on page 1588](#).

Table 158: Forms of the configure Command

Command	Edit Access	Commit Access
configure	<ul style="list-style-type: none"> • No one can lock the configuration. All users can make configuration changes. <p>When you enter configuration mode, the CLI displays the following information:</p> <ul style="list-style-type: none"> • A list of other users editing the configuration. • Hierarchy levels the users are viewing or editing. • Whether the configuration has been changed, but not committed. • When multiple users enter conflicting configurations, the most recent change to be entered takes precedence. 	<ul style="list-style-type: none"> • No one can lock the configuration. All users can commit all changes to the configuration. • If you and another user make changes and the other user commits changes, your changes are committed as well.

Table 158: Forms of the configure Command (*continued*)

Command	Edit Access	Commit Access
configure exclusive	<ul style="list-style-type: none"> One user locks the configuration and makes changes without interference from other users. Other users can enter and exit configuration mode, but they cannot commit the configuration. If you enter configuration mode while another user has locked the configuration (with the configure exclusive command), the CLI displays the user and the hierarchy level the user is viewing or editing. If you enter configuration mode while another user has locked the configuration, you can forcibly log out that user with the request system logout operational mode command. For details, see the CLI Explorer. 	
configure private	<ul style="list-style-type: none"> Multiple users can edit the configuration at the same time. Each user has a private candidate configuration to edit independently of other users. When multiple users enter conflicting configurations, the first commit operation takes precedence over subsequent commit operations. 	<ul style="list-style-type: none"> When you commit the configuration, the router verifies that the operational (running) configuration has not been modified by another user before accepting your private candidate configuration as the new operational configuration. If the configuration has been modified by another user, you can merge the modifications into your private candidate configuration and attempt to commit again.

Related Documentation

- *Committing a Junos OS Configuration*
- *Example: Using the configure Command*
- *Displaying Users Currently Editing the Configuration*
- *Using the configure exclusive Command*
- *Updating the configure private Configuration*
- *Displaying set Commands from the Junos OS Configuration*

Junos OS Commit Model for Router or Switch Configuration

The router or switch configuration is saved using a commit model: that is, a candidate configuration is modified as desired and then committed to the system. Once a configuration has been committed, the router or switch checks the configuration for syntax errors, and if no errors are found, the configuration is saved as **juniper.conf.gz** and activated. The former active configuration file is saved as the first rollback configuration file (**juniper.conf.1.gz**), and all other rollback configuration files are incremented by 1. For example, **juniper.conf.1.gz** is incremented to **juniper.conf.2.gz**, making it the second rollback configuration file. The router or switch can have a maximum of 49 rollback configurations (1–49) saved on the system.

On the router or switch, the active configuration file and the first three rollback files (**juniper.conf.gz.1**, **juniper.conf.gz.2**, **juniper.conf.gz.3**) are located in the **/config** directory. If the file **rescue.conf.gz** is saved on the system, this file should also be saved in the **/config** directory. The factory default files are located in the **/etc/config** directory.

There are two mechanisms used to propagate the configurations between Routing Engines within a router or switch:

- Synchronization—Propagates a configuration from one Routing Engine to a second Routing Engine within the same router or switch chassis.



NOTE: The QFX3500 switch has only one Routing Engine.

To synchronize configurations, use the **commit synchronize** CLI command. If one of the Routing Engines is locked, the synchronization fails. If synchronization fails because of a locked configuration file, you can use the **commit synchronize force** command. This command overrides the lock and synchronizes the configuration files.

- Distribution—Propagates a configuration across the routing plane on a multichassis router or switch. Distribution occurs automatically. There is no user command available to control the distribution process. If a configuration is locked during a distribution of a configuration, the locked configuration does not receive the distributed configuration file, so the synchronization fails. You need to clear the lock before the configuration and resynchronize the routing planes.



NOTE: When you use the **commit synchronize force** CLI command on a multichassis platform, the forced synchronization of the configuration files does not affect the distribution of the configuration file across the routing plane. If a configuration file is locked on a router or switch remote from the router or switch where the command was issued, the synchronization fails on the remote router or switch. You need to clear the lock and reissue the **synchronization** command.

Related Documentation

- *Configuring Junos OS for the First Time on a Router or Switch with a Single Routing Engine*

Understanding Configuration Files

A configuration file stores the complete configuration of a switch. The current configuration of a switch is called the active configuration. You can alter this current configuration and you can also return to a previous configuration or to a rescue configuration.

Juniper Networks Junos OS saves the 50 most recently committed configuration files on a switch so that you can return to a previous configuration. The configuration files are named:

- **juniper.conf.gz**—The current active configuration.

- `juniper.conf.1.gz` to `juniper.conf.49.gz`—Rollback configurations.

To make changes to the configuration file, you have to work in the configuration mode in the CLI. When making changes to a configuration file, you are viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the active configuration or causing potential damage to your current network operations. Once you commit the changes made to the candidate configuration, the system updates the active configuration.

Related Documentation

- [Uploading a Configuration File on page 1609](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Reverting to the Rescue Configuration on page 175](#)
- [Configuration File Terms on page 48](#)

Understanding How the Junos Configuration Is Stored

When you edit a configuration, you work in a copy of the current configuration to create a candidate configuration. The changes you make to the candidate configuration are visible in the CLI immediately, so if multiple users are editing the configuration at the same time, all users can see all changes.

To have a candidate configuration take effect, you *commit* the changes. At this point, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration, when you commit the candidate configuration, all changes made by all the users take effect.

In addition to saving the current configuration, the CLI saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0, which is the current operational version and the default configuration that the system returns to if you roll back to a previous configuration. The oldest saved configuration is version 49.

The currently operational Junos OS configuration is stored in the file `juniper.conf` and the last three committed configurations are stored in the files `juniper.conf.1`, `juniper.conf.2`, and `juniper.conf.3`. These four files are located in the directory `/config`, which is on the switch's hard disk. The remaining 46 previous versions of committed configurations, the files `juniper.conf.4` through `juniper.conf.49`, are stored in the directory `/var/db/config` on the hard disk.

Related Documentation

- [Returning to the Most Recently Committed Junos Configuration on page 1600](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1601](#)
- [Loading a Configuration from a File on page 1597](#)

CHAPTER 14

Configuration

- [Configuration Tasks on page 1593](#)
- [Configuration Statements on page 1613](#)
- [Default Configurations on page 1619](#)
- [Configuration Examples on page 1625](#)

Configuration Tasks

- [Comparing Configuration Changes with a Prior Version on page 1593](#)
- [Compressing the Current Configuration File on page 1595](#)
- [Creating and Returning to a Rescue Configuration on page 1596](#)
- [Loading a Configuration from a File on page 1597](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Returning to the Most Recently Committed Junos Configuration on page 1600](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1601](#)
- [Reverting to the Default Factory Configuration on page 1606](#)
- [Reverting to the Rescue Configuration on page 1606](#)
- [Rolling Back Junos OS Configuration Changes on page 1607](#)
- [Saving a Configuration to a File on page 1608](#)
- [Setting or Deleting the Rescue Configuration on page 1609](#)
- [Uploading a Configuration File on page 1609](#)
- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611](#)

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

[edit]

```
user@host# show | compare (filename| rollback n)
```

filename is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.

n is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (`/config/juniper.conf`).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ().

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 60;
  advertise-inactive;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  peer-as 33333;
  allow 2.2.2.2/32;
}
group test-peers {
  type external;
  allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
```

```
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
  -type external;
  -allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 90;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  advertise-inactive;
  peer-as 3333;
  allow 2.2.2.2/32;
}
```

**Related
Documentation**

- [Creating and Returning to a Rescue Configuration on page 1596](#)

Compressing the Current Configuration File

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

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



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

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

```
[edit system]
compress-configuration-files;
```

Commit the current configuration file to include the **compression-configuration-files** statement. Commit the configuration again to compress the current configuration file:

```
[edit system]
user@host# set compress-configuration-files
user@host# commit
```

```
commit complete
user@host# commit
commit complete
```

- If you do not want to compress the current operational configuration file, include the **no-compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-compression-configuration-files;
```

Commit the current configuration file to include the **no-compress-configuration-files** statement. Commit the configuration again to uncompress the current configuration file:

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

Related Documentation

- [Junos OS Commit Model for Router or Switch Configuration on page 50](#)
- [compress-configuration-files on page 250](#)

Creating and Returning to a Rescue Configuration

A *rescue configuration* allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]
user@host# rollback rescue
load complete
```



NOTE: If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, the **rollback** command fails, an error message appears, and the current configuration remains active.

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]
user@host# rollback rescue
```



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

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

- Related Documentation**
- [Comparing Configuration Changes with a Prior Version on page 1593](#)
 - [Saving a Configuration to a File on page 1605](#)

Loading a Configuration from a File

You can create a file, copy the file to the local router, and then load the file into the CLI. After you have loaded the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively using the CLI and commit it at a later time.

You can also create a configuration while typing at the terminal and then load it. Loading a configuration from the terminal is generally useful when you are cutting existing portions of the configuration and pasting them elsewhere in the configuration.

To load an existing configuration file that is located on the router, use the **load** configuration mode command:

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
filename <relative>
```

For information about specifying the filename, see *Viewing Files and Directories on a Device Running Junos OS*.

To load a configuration from the terminal, use the following version of the **load** configuration mode command. Press Ctrl-d to end input.

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
terminal <relative>
```

To replace an entire configuration, specify the **override** option at any level of the hierarchy. A **load override** operation completely replaces the current candidate configuration with the file you are loading. Thus, if you saved a complete configuration, use this option.

An **override** operation discards the current candidate configuration and loads the configuration in **filename** or the configuration that you type at the terminal. When you use the **override** option and commit the configuration, all system processes reparse the configuration. For an example, see [Figure 36 on page 1625](#).

To replace portions of a configuration, specify the **replace** option. The **load replace** operation looks for **replace:** tags that you added to the loaded file, and replaces the parts of the candidate configuration with whatever is specified after the tag. This is useful when you want more control over exactly what is being changed. For this operation to work, you must include **replace:** tags in the file or configuration you type at the terminal. The software searches for the **replace:** tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the **replace** operation adds to the configuration the statements marked with the **replace:** tag. For an example, see [Figure 37 on page 1626](#).

If, in an **override** or **merge** operation, you specify a file or type text that contains **replace:** tags, the **replace:** tags are ignored and the **override** or **merge** operation is performed.

If you are performing a **replace** operation and the file you specify or text you type does not contain any **replace:** tags, the **replace** operation is effectively equivalent to a **merge** operation. This might be useful if you are running automated scripts and cannot know in advance whether the scripts need to perform a **replace** or a **merge** operation. The scripts can use the **replace** operation to cover either case.

The **load merge** operation adds the saved file to the existing candidate configuration. This is useful if you are adding new configuration sections. For example, suppose that you are adding a BGP configuration to the **[edit protocols]** hierarchy level, where there was no BGP configuration before, you can use the **load merge** operation to combine the saved file configuration to the existing candidate configuration. If the existing configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the existing configuration.

To replace only the configuration that has changed, specify the **update** option at any level of the hierarchy. The **load update** operation compares the candidate configuration and the file you are loading, and only changes the parts of the candidate configuration that are different from the new configuration. You would use this, for example, if there is an existing BGP configuration and the file you are loading changes it in some way.

To change part of the configuration with a patch file, specify the **patch** option. The **load patch** operation loads a file or terminal input that contains configuration changes. First, on a device that already has the configuration changes, you type the **show | compare** command to output the differences between two configurations. Then you can load the differences on another router. The advantage of the **load patch** command is that it saves you from having to copy snippets from different hierarchy levels into a text file prior to loading them into the target device. This might be a useful time saver if you are configuring several devices with the same options. For example, suppose that you configure a routing policy on Device router1 and you want to replicate the policy configuration on Device router2, router3, and router4, you can use the **load patch** operation.

First, run the **show | compare** command.

```
user@router1# show | compare rollback 3
[edit protocols ospf]
+ export default-static;
- export static-default
[edit policy-options]
+ policy-statement default-static {
```

```
+      from protocol static;
+      then accept;
+  }
```

Copy the output of the **show | compare** command to the clipboard, making sure to include the hierarchy levels. On Device router2, router3, and router4, type **load patch terminal** and paste the output. Press Enter and then press Ctrl-d to end the operation. If the patch input specifies different values for an existing statement, the patch input overrides the existing statement.

To use the **merge**, **replace**, **set**, or **update** option without specifying the full hierarchy level, specify the **relative** option. For example:

```
[edit system]
user@host# show static-host-mapping
bob sysid 987.654.321ab
[edit system]
user@host# load replace terminal relative
[Type ^D at a new line to end input]
replace: static-host-mapping {
  bob sysid 0123.456.789bc;
}
load complete
[edit system]
user@host# show static-host-mapping
bob sysid 0123.456.789bc;
```

To load a configuration that contains the **set** configuration mode command, specify the **set** option. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as **set**, **edit**, **exit**, and **top**. For an example, see [Figure 40 on page 1627](#).

To copy a configuration file from another network system to the local router, you can use the SSH and Telnet utilities, as described in the [CLI Explorer](#).



NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

Related Documentation

- [Examples: Loading a Configuration from a File on page 1625](#)

Loading a Previous Configuration File

You can use the **rollback** <*number*> command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

Syntax

rollback <*number*>

Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
 - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
 - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

Related Documentation

- [Configuration File Terms on page 48](#)

Returning to the Most Recently Committed Junos Configuration

To return to the most recently committed configuration and load it into configuration mode without activating it, use the **rollback** configuration mode command:

```
[edit]
user@host# rollback

load complete
```

To activate the configuration to which you rolled back, use the **commit** command:

```
[edit]
user@host# rollback
load complete
[edit]
user@host# commit
```

- Related Documentation**
- [Rolling Back Junos OS Configuration Changes on page 1607](#)
 - [Returning to a Previously Committed Junos OS Configuration on page 1601](#)
 - [Understanding How the Junos Configuration Is Stored on page 1591](#)

Returning to a Previously Committed Junos OS Configuration

This topic explains how you can return to a configuration prior to the most recently committed one, and contains the following sections:

- [Returning to a Configuration Prior to the One Most Recently Committed on page 1601](#)
- [Displaying Previous Configurations on page 1601](#)
- [Comparing Configuration Changes with a Prior Version on page 1602](#)
- [Creating and Returning to a Rescue Configuration on page 1604](#)
- [Saving a Configuration to a File on page 1605](#)

Returning to a Configuration Prior to the One Most Recently Committed

To return to a configuration prior to the most recently committed one, include the configuration number, 0 through 49, in the **rollback** command. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

```
[edit]
user@host# rollback number
load complete
```

Displaying Previous Configurations

To display previous configurations, including the rollback number, date, time, the name of the user who committed changes, and the method of commit, use the **rollback ?** command.

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0      2005-02-27 12:52:10 PST by abc via cli
1      2005-02-26 14:47:42 PST by def via cli
2      2005-02-14 21:55:45 PST by ghi via cli
3      2005-02-10 16:11:30 PST by jkl via cli
4      2005-02-10 16:02:35 PST by mno via cli
5      2005-03-16 15:10:41 PST by pqr via cli
6      2005-03-16 14:54:21 PST by stu via cli
7      2005-03-16 14:51:38 PST by vwx via cli
8      2005-03-16 14:43:29 PST by yzz via cli
9      2005-03-16 14:15:37 PST by abc via cli
10     2005-03-16 14:13:57 PST by def via cli
11     2005-03-16 12:57:19 PST by root via other
12     2005-03-16 10:45:23 PST by root via other
13     2005-03-16 10:08:13 PST by root via other
```

```
14      2005-03-16 01:20:56 PST by root via other
15      2005-03-16 00:40:37 PST by ghi via cli
16      2005-03-16 00:39:29 PST by jkl via cli
17      2005-03-16 00:32:36 PST by mno via cli
18      2005-03-16 00:31:17 PST by pqr via cli
19      2005-03-15 19:59:00 PST by stu via cli
20      2005-03-15 19:53:39 PST by vwx via cli
21      2005-03-15 18:07:19 PST by yzz via cli
22      2005-03-15 17:59:03 PST by abc via cli
23      2005-03-15 15:05:14 PST by def via cli
24      2005-03-15 15:04:51 PST by ghi via cli
25      2005-03-15 15:03:42 PST by jkl via cli
26      2005-03-15 15:01:52 PST by mno via cli
27      2005-03-15 14:58:34 PST by pqr via cli
28      2005-03-15 13:09:37 PST by root via other
29      2005-03-12 11:01:20 PST by stu via cli
30      2005-03-12 10:57:35 PST by vwx via cli
31      2005-03-11 10:25:07 PST by yzz via cli
32      2005-03-10 23:40:58 PST by abc via cli
33      2005-03-10 23:40:38 PST by def via cli
34      2005-03-10 23:14:27 PST by ghi via cli
35      2005-03-10 23:10:16 PST by jkl via cli
36      2005-03-10 23:01:51 PST by mno via cli
37      2005-03-10 22:49:57 PST by pqr via cli
38      2005-03-10 22:24:07 PST by stu via cli
39      2005-03-10 22:20:14 PST by vwx via cli
40      2005-03-10 22:16:56 PST by yzz via cli
41      2005-03-10 22:16:41 PST by abc via cli
42      2005-03-10 20:44:00 PST by def via cli
43      2005-03-10 20:43:29 PST by ghi via cli
44      2005-03-10 20:39:14 PST by jkl via cli
45      2005-03-10 20:31:30 PST by root via other
46      2005-03-10 18:57:01 PST by mno via cli
47      2005-03-10 18:56:18 PST by pqr via cli
48      2005-03-10 18:47:49 PST by stu via cli
49      2005-03-10 18:47:34 PST by vw via cli
|Pipe through a command
[edit]
```

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]
user@host# show | compare (filename) rollback n)
```

filename is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.

n is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (`/config/juniper.conf`).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ().

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 1.1.1.1/32;
}
group fred {
    type external;
    peer-as 33333;
    allow 2.2.2.2/32;
}
group test-peers {
    type external;
    allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
```

```
-type external;  
-allow 3.3.3.3/32;  
}  
[edit protocols bgp]  
user@host# show  
group my-group {  
  type internal;  
  hold-time 90;  
  allow 1.1.1.1/32;  
}  
group fred {  
  type external;  
  advertise-inactive;  
  peer-as 3333;  
  allow 2.2.2.2/32;  
}
```

Creating and Returning to a Rescue Configuration

A *rescue configuration* allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]  
user@host# rollback rescue  
load complete
```



NOTE: If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, the **rollback** command fails, an error message appears, and the current configuration remains active.

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]  
user@host# rollback rescue  
load complete  
[edit]  
user@host# commit
```

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:


```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

Saving a Configuration to a File

Save the Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    ...
  }
}
```

- Related Documentation**
- [Returning to the Most Recently Committed Junos Configuration on page 1600](#)
 - [Loading a Configuration from a File on page 1597](#)
 - [Viewing Files and Directories on a Device Running Junos OS](#)

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1.

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

- Related Documentation**
- [Understanding Configuration Files on page 1590](#)
 - [Loading a Previous Configuration File on page 1600](#)
 - [Reverting to the Rescue Configuration on page 175](#)

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```
2. Commit your changes.

```
[edit]
user@switch# commit filename
```

- Related Documentation**
- [Setting or Deleting the Rescue Configuration on page 1609](#)
 - [Reverting to the Default Factory Configuration on page 173](#)

- [Configuration File Terms on page 48](#)

Rolling Back Junos OS Configuration Changes

This topic shows how to use the **rollback** command to return to the most recently committed Junos OS configuration. The **rollback** command is useful if you make configuration changes and then decide not to keep the changes.

The following procedure shows how to configure an SNMP health monitor on a device running Junos OS and then return to the most recently committed configuration that does not include the health monitor. When configured, the SNMP health monitor provides the network management system (NMS) with predefined monitoring for file system usage, CPU usage, and memory usage on the device.

1. Enter configuration mode:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

2. Show the current configuration (if any) for SNMP:

```
[edit]
user@host# show snmp
```

No **snmp** statements appear because SNMP has not been configured on the device.

3. Configure the health monitor:

```
[edit]
user@host# set snmp health-monitor
```

4. Show the new configuration:

```
[edit]
user@host# show snmp
health-monitor;
```

The **health-monitor** statement indicates that SNMP health monitoring is configured on the device.

5. Enter the **rollback** configuration mode command to return to the most recently committed configuration:

```
[edit]
user@host# rollback
load complete
```

6. Show the configuration again to make sure your change is no longer present:

```
[edit]
user@host# show snmp
```

No **snmp** configuration statements appear. The health monitor is no longer configured.

7. Enter the **commit** command to activate the configuration to which you rolled back:

```
[edit]
```

```
user@host# commit
```

8. Exit configuration mode:

```
[edit]
user@host# exit
Exiting configuration mode
```

You can also use the **rollback** command to return to earlier configurations.

Related Documentation

- [Returning to the Most Recently Committed Junos Configuration on page 1600](#)

Saving a Configuration to a File

Save the Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

```

    }
    ospf {
        traffic-engineering;
        reference-bandwidth 4g;
        ...
    }
}

```

Setting or Deleting the Rescue Configuration

A rescue configuration is user-defined configuration that restores connectivity to switch. You set a current committed configuration to be the rescue configuration through the CLI. If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To set the current active configuration as the rescue configuration:

```
user@switch> request system configuration rescue save
```

To delete an existing rescue configuration:

```
user@switch> request system configuration rescue delete
```

Related Documentation

- [Reverting to the Default Factory Configuration on page 173](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Configuration File Terms on page 48](#)
- [CLI Explorer](#)

Uploading a Configuration File

You can create a configuration file on your local system, copy the file to the switch, and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Create the configuration file using a text editor such as Notepad, making sure that the syntax of the configuration file is correct. For more information about testing the syntax of a configuration file see the *Junos OS System Basics and Services Command Reference* at <http://www.juniper.net/techpubs/software/junos/index.html>.
2. In the configuration text file, use an option to perform the required action when the file is loaded. [Table 159 on page 1610](#) lists and describes some options for the **load** command.

Table 159: Options for the load Command

Options	Description
merge	Combines the current active configuration and the configuration in the filename you specify or the one that you type at the terminal. A merge operation is useful when you are adding a new section to an existing configuration. If the active configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the active configuration.
override	Discards the current candidate configuration and loads the configuration in the filename you specify or the one that you type at the terminal. When you use the override option and commit the configuration, all system processes reparse the configuration. You can use the override option at any level of the hierarchy.
replace	Searches for the replace tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the replace operation adds the statements marked with the replace tag to the active configuration. NOTE: For this operation to work, you must include replace tags in the text file or in the configuration you type at the terminal.

- Press Ctrl+a to select all the text in the configuration file.
- Press Ctrl+c to copy the contents of the configuration text file to the Clipboard.
- Log in to the switch using your username and password.
- To enter configuration mode:
user@switch> **configure**

You will see this output, with the hash or pound mark indicating configuration mode.
Entering configuration mode
[edit]
user@switch#
- Load the configuration file:
[edit]
user@switch# **load merge terminal**
- At the cursor, paste the contents of the Clipboard using the mouse and the Paste icon:
[edit]
user@switch# **load merge terminal**
[Type ^D at a new line to end input]
>Cursor is here. Paste the contents of the clipboard here<
- Press Enter.
- Press Ctrl+d to set the end-of-file marker.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt. You can also edit the configuration interactively using the CLI and commit it at a later time.

Related Documentation

- [Understanding Configuration Files on page 1590](#)

Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site

You can configure a router or switch to transfer its configuration to an archive file periodically. The following tasks describe how to transfer the configuration to an archive site:

1. [Configuring the Router or Switch to Transfer Its Currently Active Configuration to an Archive on page 1611](#)
2. [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1611](#)
3. [Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1612](#)
4. [Configuring Archive Sites for Transfer of Active Configuration Files on page 1612](#)

Configuring the Router or Switch to Transfer Its Currently Active Configuration to an Archive

If you want to back up your device's current configuration to an archive site, you can configure the router or switch to transfer its currently active configuration by FTP or secure copy (SCP) periodically or after each commit.

To configure the router or switch to transfer its currently active configuration to an archive site, include statements at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username<:password>@host-address<:port>/url-path;
  scp://username<:password>@host-address<:port>/url-path;
}
transfer-interval interval;
transfer-on-commit;
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path"

Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site

To configure the router or switch to periodically transfer its currently active configuration to an archive site, include the **transfer-interval** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-interval interval;
```

The *interval* is a period of time ranging from 15 through 2880 minutes.

Configuring Transfer of the Current Active Configuration When a Configuration Is Committed

To configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration, include the **transfer-on-commit** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-on-commit;
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([]). For example,
"scp://username<:password>@[ipv6-host-address]<:port>/url-path"

Configuring Archive Sites for Transfer of Active Configuration Files

When you configure the router or switch to transfer its configuration files, you specify an archive site to which the files are transferred. If you specify more than one archive site, the router or switch attempts to transfer files to the first archive site in the list, moving to the next site only if the transfer fails.

When you use the **archive-sites** statement, you can specify a destination as an FTP URL, or SCP-style remote file specification. The URL type **file://** is also supported.

To configure the archive site, include the **archive-sites** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username@host:<port>url-path password password;
  scp://username@host:<port>url-path password password;
  file://<path>/<filename>;
}
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([]). For example,
"scp://username<:password>@[ipv6-host-address]<:port>/url-path"

When you specify the archive site, do not add a forward slash (/) to the end of the URL.

The destination filename is saved in the following format, where *n* corresponds to the number of the compressed configuration rollback file that has been archived:

```
<router-name>_juniper.conf.n.gz_YYYYMMDD_HHMMSS
```




.....



NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.

.....



Configuration Statements

- [archival on page 1614](#)
- [archive-sites \(Configuration File\) on page 1615](#)
- [configuration on page 1617](#)
- [transfer-interval \(Configuration\) on page 1618](#)
- [transfer-on-commit on page 1619](#)

archival

Syntax	<pre> archival { configuration { archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } transfer-interval interval; transfer-on-commit; } } </pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.
<div>  NOTE: The <code>edit system archival</code> hierarchy is not available on QFabric systems. </div>	
Options	The remaining statements are explained separately.
<div>  NOTE: The <code>[edit system archival]</code> hierarchy is not available on QFabric systems. </div>	
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611

archive-sites (Configuration File)


Syntax	<pre>archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; }</pre>
Hierarchy Level	[edit system archival configuration]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "scp://username<:password>@[ipv6-host-address]<:port>/url-path"</p> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <p><i>router-name_juniper.conf.n.gz_YYYYMMDD_HHMMSS.</i></p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
Options	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p>file:// —transfer on a path to a named file</p> <p>ftp:// —transfer using active FTP server</p> <p>pasvftp:// —transfer to a device that only accepts passive FTP services</p>

scp:// —transfer to a known host using background SCP file transfers

Required Privilege	system—To view this statement in the configuration.
Level	system-control—To add this statement to the configuration.

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Configuring Archive Sites for Transfer of Active Configuration Files on page 1612• Junos OS Commit Model for Router or Switch Configuration on page 50• configuration on page 1617• transfer-on-commit on page 1619 |
|------------------------------|--|

configuration

Syntax	<pre>configuration { transfer-interval interval; transfer-on-commit; archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } }</pre>
Hierarchy Level	[edit system archival]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the router or switch to periodically transfer its currently active configuration (or after each commit).
<div>  <p>NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>	
Options	The remaining statements are explained separately.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611 • archive on page 6366 • archive-sites on page 1615 • transfer-interval on page 1618 • transfer-on-commit on page 1619

transfer-interval (Configuration)

Syntax	<code>transfer-interval <i>interval</i>;</code>
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to periodically transfer its currently active configuration to an archive site.



NOTE: The `edit system archival` hierarchy is not available on QFabric systems.

Options *interval*—Interval at which to transfer the current configuration to an archive site.
Range: 15 through 2880 minutes



NOTE: The `[edit system archival]` hierarchy is not available on QFabric systems.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1611](#)
- [archive on page 6366](#)
- [configuration on page 1617](#)
- [transfer-on-commit on page 1619](#)

transfer-on-commit

Syntax	transfer-on-commit;
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path" .



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1612 • archive on page 6366 • configuration on page 1617 • transfer-interval on page 1618

Default Configurations

- [QFX3500 Switch Default Configuration on page 1619](#)

QFX3500 Switch Default Configuration

Each QFX Series product is programmed with a factory default configuration that contains the values set for each configuration parameter when a switch is shipped. The default configuration file sets values for system parameters such as **syslog** and **commit**, configures storm control and Ethernet switching on all interfaces, and enables IGMP snooping, RSTP, and LLDP protocols.

When you commit changes to the configuration, a new configuration file is created, which becomes the active configuration. You can always revert to the factory default configuration if you need to.

The following factory default configuration file is for a QFX3500 switch with 48 ports:



NOTE: In this example, xe-0/0/0 through xe-0/0/47 are the network interface ports.

```
protocols {
  igmp-snooping {
    vlan all;
  }
  rstp;
  lldp {
    interface all;
  }
}
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/3 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/4 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/5 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/6 {
    unit 0 {
      family ethernet-switching;
    }
  }
```



```
}
xe-0/0/7 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/8 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/9 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/10 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/12 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/13 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/14 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/15 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/16 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/17 {
  unit 0 {
    family ethernet-switching;
  }
}
```

```
xe-0/0/18 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/19 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/20 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/21 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/22 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/23 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/24 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/25 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/26 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/27 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/28 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/29 {
```

```
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/30 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/31 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/32 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/33 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/34 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/35 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/36 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/37 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/38 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/39 {  
    unit 0 {  
        family ethernet-switching;  
    }  
}  
xe-0/0/40 {  
    unit 0 {
```

```
        family ethernet-switching;
    }
}
xe-0/0/41 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/42 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/43 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/44 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/45 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/46 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/47 {
    unit 0 {
        family ethernet-switching;
    }
}
}
ethernet-switching-options {
    storm-control {
        interface all;
    }
}
system {
    syslog {
        archive size 256k;
        file default-log-messages {
            structured-data;
        }
    }
    user * {
        any emergency;
    }
    file messages {
        any notice;
        authorization info;
    }
}
```

```
    }
    file interactive-commands {
        interactive-commands any;
    }
}
ports {
    console type vt100;
}
compress-configuration-files;
login {
    password {
        minimum-length 6;
        minimum-changes 1;
        change-type set transitions;
        format md5;
    }
}
commit {
    factory-settings {
        reset-chassis-lcd-menu;
    }
}
}
```

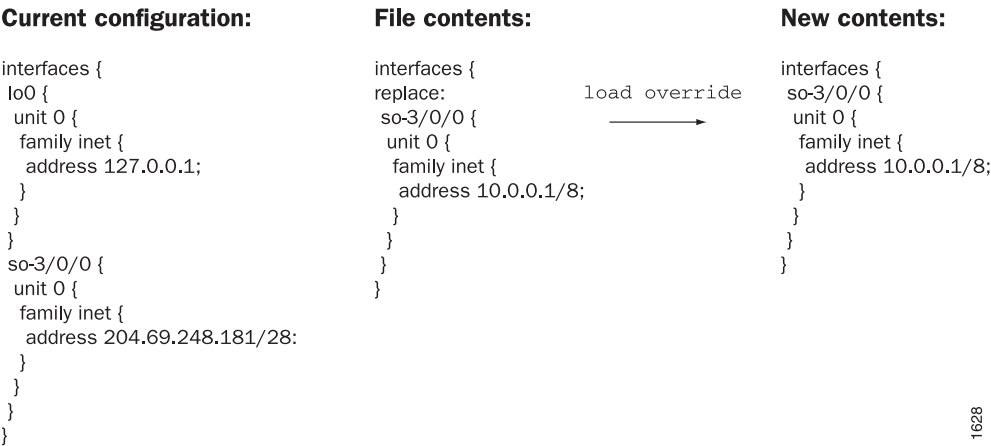
- Related Documentation
- [Reverting to the Default Factory Configuration on page 173](#)
 - [Configuring a QFX3500 Device as a Standalone Switch on page 163](#)
 - [Understanding Configuration Files on page 1590](#)
 - [Interfaces Overview on page 2415](#)

Configuration Examples

- [Examples: Loading a Configuration from a File on page 1625](#)

Examples: Loading a Configuration from a File

Figure 36: Overriding the Current Configuration



1628

Figure 37: Using the replace Option

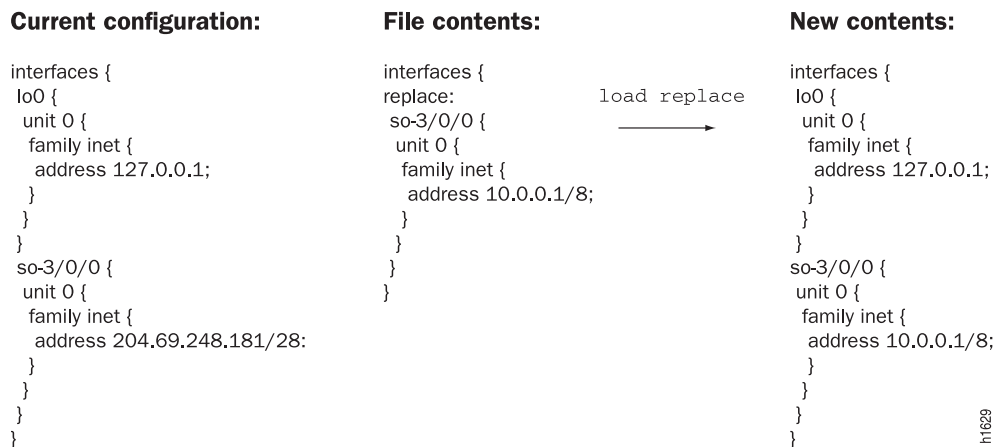


Figure 38: Using the merge Option

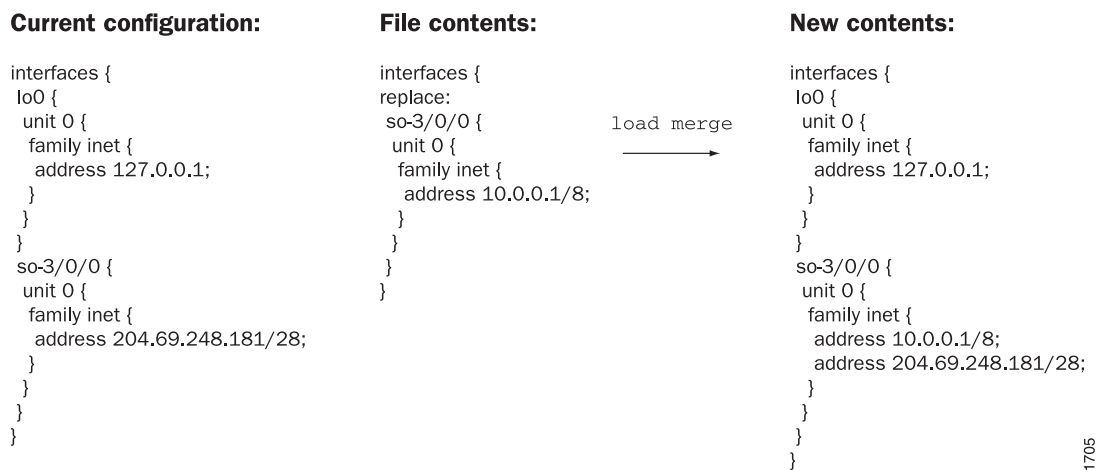


Figure 39: Using a Patch File

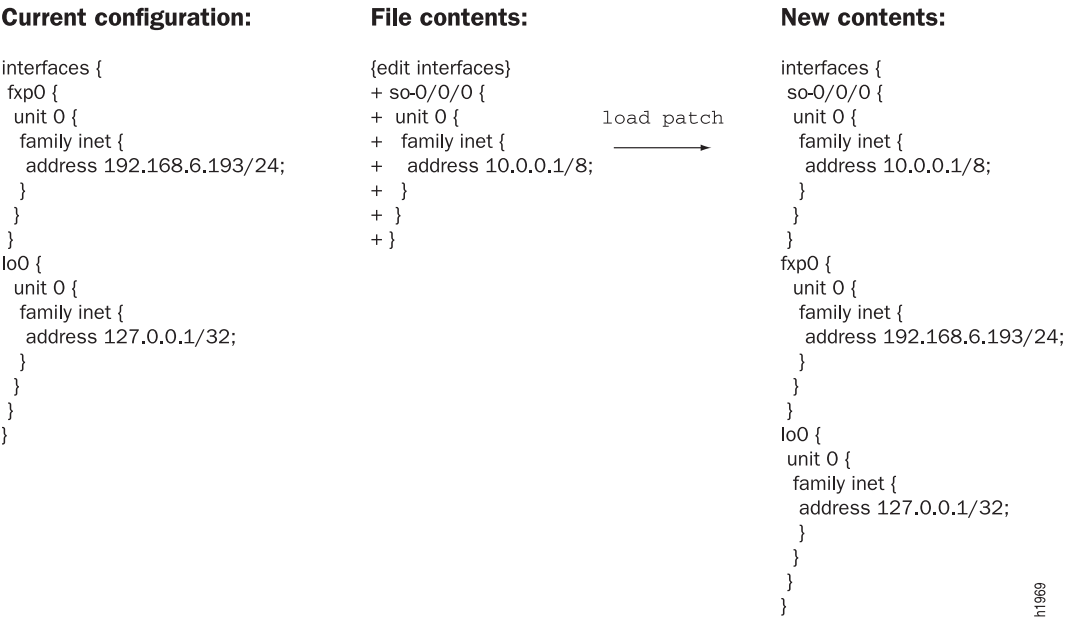
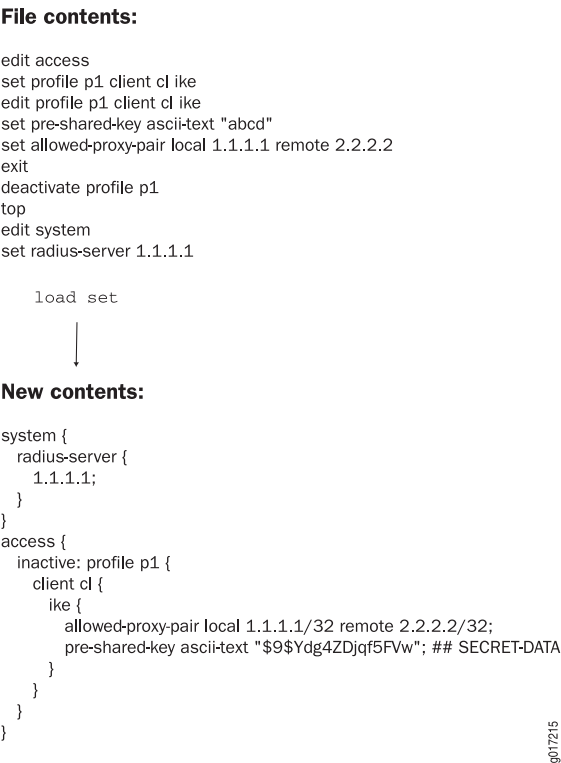


Figure 40: Using the set Option



Related Documentation

- [Loading a Configuration from a File on page 1597](#)

CHAPTER 15

Administration

- [Operational Commands on page 1629](#)

Operational Commands

- [clear log](#)
- [clear system commit](#)
- [file archive](#)
- [file checksum md5](#)
- [file checksum sha1](#)
- [file checksum sha-256](#)
- [file compare](#)
- [file delete](#)
- [file list](#)
- [file rename](#)
- [file show](#)
- [request system configuration rescue delete](#)
- [request system configuration rescue save](#)
- [show system commit](#)
- [show system configuration archival](#)
- [show system configuration rescue](#)
- [show system rollback](#)
- [test configuration](#)

clear log

Syntax	<code>clear log <i>filename</i></code> <code><all></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Remove contents of a log file.
Options	<i>filename</i> —Name of the specific log file to delete. all —(Optional) Delete the specified log file and all archived versions of it.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show log on page 889
List of Sample Output	clear log on page 1630
Output Fields	See file list for an explanation of output fields.

Sample Output

clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```

clear system commit

Syntax	clear system commit
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear any pending commit operation.
Options	This command has no options.
Required Privilege Level	maintenance (or the actual user who scheduled the commit)
Related Documentation	<ul style="list-style-type: none"> • show system commit on page 939
List of Sample Output	clear system commit on page 1631 clear system commit (None Pending) on page 1631 clear system commit (User Does Not Have Required Privilege Level) on page 1631
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

file archive

Syntax	<code>file archive destination <i>destination</i> source <i>source</i> <compress></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
Options	<p>destination <i>destination</i>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none">• For archived files—The suffix .tar• For archived and compressed files—The suffix .tgz <p>source <i>source</i>—Source of the original file or files. Specify the source as a URL or filename.</p> <p>compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix .tgz.</p>
Required Privilege Level	maintenance
List of Sample Output	file archive (Multiple Files) on page 1632 file archive (Single File) on page 1632 file archive (with Compression) on page 1633
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory `/var/log/messages` as the single file `messages-archive.tgz`.

```
user@host> file archive compress source /var/log/messages* destination  
/var/log/messages-archive.tgz  
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

file checksum md5

Syntax	<code>file checksum md5 <pathname> filename</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Message Digest 5 (MD5) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the MD5 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum sha-256 on page 341• file checksum sha1 on page 340• <i>op</i>
List of Sample Output	file checksum md5 on page 1634
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

file checksum sha1

Syntax	<code>file checksum sha1 <pathname> filename</code>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
Options	<p>pathname—(Optional) Path to a filename.</p> <p>filename—Name of a local file for which to calculate the SHA-1 checksum.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i> • file checksum md5 on page 339 • file checksum sha-256 on page 341 • <i>op</i>
List of Sample Output	file checksum sha1 on page 1635
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

file checksum sha-256

Syntax	<code>file checksum sha-256 <pathname> filename</code>
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
Options	pathname —(Optional) Path to a filename. filename —Name of a local file for which to calculate the SHA-256 checksum.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i>• file checksum md5 on page 339• file checksum sha1 on page 340• <i>op</i>
List of Sample Output	file checksum sha-256 on page 1636
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```


file compare

Syntax	<pre>file compare (files <i>filename filename</i>) <context unified> <ignore-white-space></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> • Default—In the first line of output, c means lines were changed between the two files, d means lines were deleted between the two files, and a means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (<) in front of output lines refers to the first file. A right angle bracket (>) in front of output lines refers to the second file. • Context—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-). • Unified—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.
Options	<p>files <i>filename</i>—Names of two local files to compare.</p> <p>context—(Optional) Display output in context format.</p> <p>ignore-white-space—(Optional) Ignore changes in the amount of white space.</p> <p>unified—(Optional) Display output in unified format.</p>
Required Privilege Level	none
List of Sample Output	<p>file compare files on page 1638</p> <p>file compare files context on page 1638</p> <p>file compare files unified on page 1638</p> <p>file compare files unified ignore-white-space on page 1638</p>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

file delete

Syntax	<code>file delete <i>filename</i></code> <code><purge></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete a file on the local router or switch.
Options	<i>filename</i> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued. <i>purge</i> —(Optional) Overwrite regular files before deleting them.
Required Privilege Level	maintenance
List of Sample Output	file delete on page 1640 file delete (Routing Matrix) on page 1640
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax	file list <detail recursive> <filename>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of files on the local router or switch.
Options	<p>none—Display a list of all files for the current directory.</p> <p>detail recursive—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.</p> <p>filename—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.</p>
Additional Information	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the file list command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the file list command.
Required Privilege Level	maintenance
List of Sample Output	file list on page 1641 file list (Routing Matrix) on page 1641
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

file rename

Syntax	<code>file rename <i>source destination</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Rename a file on the local router or switch.
Options	<i>destination</i> —New name for the file. <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
Required Privilege Level	maintenance
List of Sample Output	file rename on page 1643 file rename (Routing Matrix) on page 1643
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
```

```
-----
/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```


file show

Syntax	<code>file show <i>filename</i></code> <code><encoding (base64 raw)></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the contents of a file.
Options	<i>filename</i> —Name of a file. For a routing matrix, the filename must include the chassis information. <code>encoding (base64 raw)</code> —(Optional) Encode file contents with base64 encoding or show raw text.
Required Privilege Level	maintenance
List of Sample Output	file show on page 1645 file show (Routing Matrix) on page 1645
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/gdbinit
lcc0-re0:
-----
#####
# Settings
#####


set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

request system configuration rescue delete


Syntax	request system configuration rescue delete
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete an existing rescue configuration.
<div>  NOTE: The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • request system configuration rescue save on page 374 • request system software rollback on page 416 • show system commit on page 939
List of Sample Output	request system configuration rescue delete on page 1647
Output Fields	This command produces no output.

Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax	request system configuration rescue save
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the rollback command.
<div> NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.</div>	
Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• request system software delete on page 404• request system software rollback on page 416• show system commit on page 939
List of Sample Output	request system configuration rescue save on page 1648
Output Fields	This command produces no output.

Sample Output

request system configuration rescue save

```
user@host> request system configuration rescue save
```

show system commit

Syntax	show system commit
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the system commit history and any pending commit operation.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear system commit on page 331
List of Sample Output	show system commit on page 1650 show system commit (At a Particular Time) on page 1650 show system commit (At the Next Reboot) on page 1650 show system commit (Rollback Pending) on page 1650 show system commit (QFX Series) on page 1650
Output Fields	Table 83 on page 939 describes the output fields for the show system commit command. Output fields are listed in the approximate order in which they appear.

Table 160: show system commit Output Fields

Field Name	Field Description
Junos XML protocol	Displays the last 50 commit operations listed, most recent to first. The identifier Junos XML protocol designates a configuration created for recovery using the request system configuration rescue save command.
Junos XML protocol	Date and time of the commit operation.
Junos XML protocol	User who executed the commit operation.
Junos XML protocol	Method used to execute the commit operation: <ul style="list-style-type: none"> • Junos XML protocol—CLI interactive user performed the commit operation. • Junos XML protocol—Junos XML protocol client performed the commit operation. • synchronize—The commit synchronize command was performed on the other Routing Engine. • snmp—An SNMP set request caused the commit operation. • button—A button on the router or switch was pressed to commit a rescue configuration for recovery. • autoinstall—A configuration obtained through autoinstallation was committed. • other—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.

Sample Output

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

show system configuration archival

Syntax show system configuration archival

Release Information Introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display directory and number of files queued for archival transfer.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options This command has no options.

Required Privilege Level maintenance

List of Sample Output [show system configuration archival on page 1651](#)

Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

show system configuration rescue

Syntax	show system configuration rescue
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a rescue configuration, if one exists.



NOTE: The [edit system configuration] hierarchy is not available on QFabric systems.

Options	This command has no options.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• show system configuration archival on page 941
List of Sample Output	show system configuration rescue on page 1652

Sample Output

show system configuration rescue

```
user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
  services {
```



```
        ftp;  
        rlogin;  
        rsh;  
        telnet;  
    }  
}  
.....
```

show system rollback

Syntax `show system rollback number`
`<compare number>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the contents of a previously committed configuration, or the differences between two previously committed configurations.



NOTE: The `show system rollback` command is a purely operational mode command and cannot be issued with `run` from the configuration mode.

Options *number*—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.

compare number—(Optional) Number of another previously committed (rollback) configuration to compare to rollback *number*. The output displays the differences between the two configurations. The range of values is 0 through 49.

Required Privilege Level view

List of Sample Output [show system rollback compare on page 1654](#)

Sample Output

show system rollback compare

```
user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
+       }
+     }
+   }
+ }
```

```
+      }
+    }
+    ge-1/3/0 {
+      unit 0 {
+        family inet {
+          filter {
+            input mf_plp;
+          }
+          address 12.1.1.1/30;
+        }
+      }
+    }
+  }
+}
```

test configuration

Syntax	<code>test configuration filename</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found.
Options	filename —Name of the configuration file. syntax-only —Check the syntax of a partial configuration file, without checking for commit errors. This option introduced in Junos OS Release 12.1.
Required Privilege Level	view
List of Sample Output	test configuration on page 1656
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
    syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
    statement must contain additional statements
configuration syntax failed
```

CHAPTER 16

Troubleshooting

- [Troubleshooting Procedures on page 1657](#)

Troubleshooting Procedures

- [Loading a Previous Configuration File on page 1657](#)
- [Reverting to the Default Factory Configuration on page 1658](#)
- [Reverting to the Rescue Configuration on page 1658](#)

Loading a Previous Configuration File

You can use the **rollback** <*number*> command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

Syntax

rollback <*number*>

Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
 - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
 - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
```

```
user@switch# commit
```

Related Documentation

- [Configuration File Terms on page 48](#)

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1.

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

Related Documentation

- [Understanding Configuration Files on page 1590](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Reverting to the Rescue Configuration on page 175](#)

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```
2. Commit your changes.

```
[edit]
user@switch# commit filename
```

Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1609](#)
- [Reverting to the Default Factory Configuration on page 173](#)
- [Configuration File Terms on page 48](#)

PART 7

User and Access Management

- [Overview on page 1661](#)
- [Configuration on page 1691](#)
- [Administration on page 1831](#)

CHAPTER 17

Overview

- [Software Overview on page 1661](#)
- [Access Control Overview on page 1664](#)

Software Overview

- [Understanding Software Infrastructure and Processes on page 1661](#)
- [Understanding User and Access Management Features on the QFabric System on page 1663](#)

Understanding Software Infrastructure and Processes

The QFX Series products run the Juniper Networks Junos OS. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the switch.

Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

Using the Junos OS command-line interface (CLI), you configure switching features and set the properties of network interfaces. After activating a software configuration, use either the Junos Space or CLI user interface to monitor, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 1661](#)
- [Junos OS Processes on page 1662](#)

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
 - Creates the packet forwarding switch, which provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network.

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of the Junos OS for added flexibility.

Table 35 on page 68 describes the primary Junos OS processes.

Table 161: Junos OS Processes

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
DNS Server process	named-service	Resolves hostnames into addresses.
Dynamic Host Configuration Protocol (DHCP) process	dhcp-service	Enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree Protocol, and access port security.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Firewall management process	firewall	Manages the firewall configuration and helps accept or reject packets that are transiting an interface on a switch.
Forwarding process	pfem	Defines how routing protocols operate on the partition. The overall performance of the partition is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Integrated Local Management Interface (ILMI) process	ilmi	Provides bidirectional exchange of management information between two ATM interfaces across a physical connection.

Table 161: Junos OS Processes (*continued*)

Process	Name	Description
Link Management Protocol (LMP) process	link-management	Establishes and maintains LMP control channels.
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the partition.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Multicast snooping process	multicast-snooping	Makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
Secure Neighbor Discovery (SEND) Protocol process	send	Protects Neighbor Discovery Protocol (NDP) messages.
Simple Network Management Protocol (SNMP) process	snmp	Enables the monitoring of network devices from a central location and provides the switch's SNMP master agent.
Tunnel OAM process	tunnel-oamd	Enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
Virtual Router Redundancy Protocol (VRRP) process	vrrp	Enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

- Related Documentation**
- *Junos OS Baseline Network Operations Guide*
 - *Junos OS System Basics Configuration Guide*

Understanding User and Access Management Features on the QFabric System

The QFabric system supports the following user and access management features:

- User authentication
- RADIUS
- Link Layer Discovery Protocol (LLDP)

- SSH
- TACACS+
- Access privilege management

The specific functionality, features, options, syntax, and hierarchy levels of some of the user and access management commands and configuration statements implemented on the QFabric system may differ somewhat from the same commands and configuration statements on standard Junos OS. See the configuration statement or command topic in the documentation set for additional information, and use the help (?) command-line function to display specific information as needed.

Some user and access management features are not yet fully supported in the full QFabric architecture, although full support is planned for future releases. The user and access management features currently unsupported on the QFabric system include:

- Full RADIUS server support, including RADIUS accounting
- **accounting-options** configuration statement hierarchy
- **tacplus-options** configuration statement

Access Control Overview

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [Understanding Login Authentication on page 1665](#)
- [Understanding LLDP on page 1666](#)
- [Understanding RADIUS Accounting on page 1667](#)
- [Understanding VSAs on the QFX Series on page 1668](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1668](#)
- [Juniper Networks Vendor-Specific TACACS+ Attributes on page 1670](#)
- [Understanding Junos OS Access Privilege Levels on page 1672](#)
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1676](#)
- [Junos OS User Authentication Methods on page 1681](#)
- [Junos OS User Accounts Overview on page 1681](#)
- [Junos OS Login Classes Overview on page 1683](#)
- [Understanding QFabric System Login Classes on page 1684](#)
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1685](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1686](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1687](#)

Overview of Template Accounts for RADIUS and TACACS+ Authentication

When you use local password authentication, you must create a local user account for every user who wants to access the system. However, when you are using RADIUS or TACACS+ authentication, you can create single accounts (for authorization purposes) that are shared by a set of users. You create these accounts using the remote and local user template accounts. When a user is using a template account, the command-line interface (CLI) username is the login name; however, the privileges, file ownership, and effective user ID are inherited from the template account.

Related Documentation

- [Configuring Remote Template Accounts for User Authentication on page 1701](#)
- [Configuring Local User Template Accounts for User Authentication on page 1695](#)

Understanding Login Authentication

You can control access to your network through the Juniper Networks QFX Series using several different authentication methods—media access control (MAC) RADIUS, for example. Authentication prevents unauthorized devices and users from gaining access to your LAN. For MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server.

You can enable end devices to access the network without authenticating on the RADIUS server by configuring the MAC address of the end device in the static MAC bypass list by configuring the MAC address using the **authentication-whitelist** statement.

You can configure one or more authentication methods on a single interface and thereby enable fallback to the next method if the first or second method is unsuccessful.

On a single interface you can configure one or a combination of several authentication methods.

This topic covers:

- [MAC RADIUS Authentication on page 1665](#)

MAC RADIUS Authentication

You can configure MAC RADIUS authentication on interfaces that are connected to end devices.

The EAP method supported for MAC RADIUS authentication on the QFX Series is EAP-MD5.

When you configure the **mac-radius restrict** option, the switch immediately attempts a MAC- RADIUS authentication by sending a request to the RADIUS server for authentication of the MAC address of the end device. If MAC address of the end device is configured for RADIUS authentication, LAN access between the two switches is created.

Related Documentation

- [Configuring RADIUS Authentication on page 1699](#)

Understanding LLDP

The QFX Series product uses Link Layer Discovery Protocol (LLDP) to learn and distribute device information on network links. The information enables the switch to identify a variety of devices quickly. This quick identification results in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include specifics, such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in Junos OS.

The QFX Series products support the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.
- **Port Identifier**—The port identification for the specified port in the local system.
- **Port Description**—The user-configured port description. The port description can be a maximum of 256 characters.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.
- **System Description**—The system description containing information about the software and current image running on the system. This information cannot be configured, but is taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports are defined; for example, bridge or router. This information cannot be configured, but is based on the model of the product.
- **Management Address**—The IP management address of the local system.

The QFX Series products support the following 802.3 TLVs:

- **Power via MDI**—A TLV that advertises media dependent interface (MDI) power support, power source equipment (PSE) power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information cannot be configured, but is based on the physical interface structure.
- **Link Aggregation**—A TLV that advertises whether the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

Related Documentation

- [Configuring LLDP on page 1693](#)

Understanding RADIUS Accounting

Juniper Networks QFX Series products support IETF RFC 2866, *RADIUS Accounting*. Configuring RADIUS accounting on the QFX Series supports collecting statistical data about users logging in to or out from a LAN and sending the data to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, analyzing and tracking usage patterns, or billing a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the switch, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. If the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Junos OS.

The RADIUS accounting process between the switch and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.
2. The switch forwards an *accounting-request* packet containing an event record to the accounting server. The event record associated with this supplicant contains an *Acct-Status-Type* attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request contains an *Acct-Status-Type* attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.
3. The RADIUS accounting server logs these events in a file as start-accounting or stop-accounting records. On FreeRADIUS, the filename is the server's address; for example, 122.69.1.250.
4. The accounting server sends an *accounting-response* packet back to the switch confirming it has received the accounting request.
5. If the switch does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

Related Documentation

- [Configuring RADIUS System Accounting on page 1697](#)

Understanding VSAs on the QFX Series

The Juniper Networks QFX Series products support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs) and are described in RFC 2138, *Remote Authentication Dial In User Service* (RADIUS).

Through VSAs, you can configure port-filtering attributes on the RADIUS server. VSAs are cleartext fields sent from the RADIUS server to the switch as a result of authentication success or failure. Authentication prevents unauthorized user access by blocking a supplicant at the port until the device is authenticated by the RADIUS server. The VSA attributes are interpreted by the switch during authentication, and the switch takes appropriate actions. Implementing port-filtering attributes with authentication on the RADIUS server provides a central location for controlling LAN access for supplicants.

These port-filtering attributes specific to Juniper Networks are encapsulated in a RADIUS server VSA with the vendor ID set to the Juniper Networks ID number, 2636.

As well as configuring port-filtering attributes through VSAs, you can apply a port firewall filter that has already been configured on the switch directly to the RADIUS server. Like port-filtering attributes, the filter is applied during the authentication process, and its actions are applied at the switch port. Adding a port firewall filter to a RADIUS server eliminates the need to add the filter to multiple ports and switches.

- Related Documentation**
- [Configuring Firewall Filters on page 4747](#)
 - [Configuring RADIUS Authentication on page 1699](#)
 - [VSA Match Conditions and Actions on page 1724](#)

Juniper Networks Vendor-Specific RADIUS Attributes

Junos OS supports the configuration of Juniper Networks RADIUS vendor-specific attributes (VSAs). These VSAs are encapsulated in a RADIUS vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. [Table 162 on page 1668](#) lists the Juniper Networks VSAs you can configure.

Table 162: Juniper Networks Vendor-Specific RADIUS Attributes

Name	Description	Type	Length	String
Juniper-Local-User-Name	Indicates the name of the user template used by this user when logging in to a device. This attribute is used only in Access-Accept packets.	1	≥3	One or more octets containing printable ASCII characters.

Table 162: Juniper Networks Vendor-Specific RADIUS Attributes (*continued*)

Name	Description	Type	Length	String
Juniper-Allow-Commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	2	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands" on page 1686.
Juniper-Deny-Commands	Contains an extended regular expression that denies the user permission to run operation mode commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	3	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands" on page 1686.
Juniper-Allow-Configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	4	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1685.
Juniper-Deny-Configuration	Contains an extended regular expression that denies the user permission to run configuration commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	5	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See "Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1685.
Juniper-Interactive-Command	Indicates the interactive command entered by the user. This attribute is used only in Accounting-Request packets.	8	≥3	One or more octets containing printable ASCII characters.
Juniper-Configuration-Change	Indicates the interactive command that results in a configuration (database) change. This attribute is used only in Accounting-Request packets.	9	≥3	One or more octets containing printable ASCII characters.

Table 162: Juniper Networks Vendor-Specific RADIUS Attributes (*continued*)

Name	Description	Type	Length	String
Juniper-User-Permissions	<p>Contains information the server uses to specify user permissions. This attribute is used only in Access-Accept packets.</p> <p>NOTE: When the Juniper-User-Permissions attribute is configured to grant the Junos OS maintenance or all permissions on a RADIUS server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the su root command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions maintenance or all, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	10	≥3	<p>One or more octets containing printable ASCII characters.</p> <p>The string is a list of permission flags separated by a space. The exact name of each flag must be specified in its entirety. See Table 164 on page 1672.</p>

For more information about the VSAs, see RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

- Related Documentation**
- [Configuring RADIUS Authentication](#)
 - [Configuring RADIUS Authentication on page 1699](#)

Juniper Networks Vendor-Specific TACACS+ Attributes

Junos OS supports the configuration of Juniper Networks TACACS+ vendor-specific attributes (VSAs). These VSAs are encapsulated in a TACACS+ vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. [Table 163 on page 1670](#) lists the Juniper Networks VSAs you can configure.

Table 163: Juniper Networks Vendor-Specific TACACS+ Attributes

Name	Description	Length	String
local-user-name	Indicates the name of the user template used by this user when logging in to a device.	≥3	One or more octets containing printable ASCII characters.

Table 163: Juniper Networks Vendor-Specific TACACS+ Attributes (*continued*)

Name	Description	Length	String
allow-commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 168 on page 1686 .
allow-configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See “Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies” on page 1685 .
deny-commands	Contains an extended regular expression that denies the user permission to run operational mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 168 on page 1686 .
deny-configuration	Contains an extended regular expression that denies the user permission to run configuration mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See Table 167 on page 1686 .
user-permissions	<p>Contains information the server uses to specify user permissions.</p> <p>NOTE: When the user-permissions attribute is configured to grant the Junos OS maintenance or all permissions on a TACACS+ server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the su root command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions maintenance or all, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	≥3	One or more octets containing printable ASCII characters. See Table 164 on page 1672 .

Related Documentation

- [Configuring TACACS+ Authentication](#)
- [Configuring TACACS+ Authentication on page 1711](#)

Understanding Junos OS Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with them. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more *permission flags*.

For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

The following sections provide additional information about permissions:

- [Junos OS Login Class Permission Flags on page 1672](#)
- [Allowing or Denying Individual Commands for Junos OS Login Classes on page 1675](#)

Junos OS Login Class Permission Flags

The **permissions** statement specifies one or more of the permission flags listed in [Table 164 on page 1672](#). Permission flags are not cumulative, so for each class you must list all the permission flags needed, including **view** to display information and **configure** to enter configuration mode. Two forms of permissions control for individual parts of the configuration are:

- "Plain" form—Provides read-only capability for that permission type. An example is **interface**.
- Form that ends in **-control**—Provides read and write capability for that permission type. An example is **interface-control**.

[Table 164 on page 1672](#) lists the Junos® operating system (Junos OS) login class permission flags that you can configure by including the **permissions** statement at the **[edit system login class class-name]** hierarchy level.

Table 164: Login Class Permission Flags

Permission Flag	Description
access	Can view the access configuration in configuration mode and with the show configuration operational mode command.
access-control	Can view and configure access information at the [edit access] hierarchy level.
admin	Can view user account information in configuration mode and with the show configuration operational mode command.
admin-control	Can view user accounts and configure them at the [edit system login] hierarchy level.

Table 164: Login Class Permission Flags (*continued*)

Permission Flag	Description
all-control	Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels.
clear	Can clear (delete) information learned from the network that is stored in various network databases by using the clear commands.
configure	Can enter configuration mode by using the configure command.
control	Can perform all control-level operations—all operations configured with the -control permission flags.
field	Can view field debug commands. Reserved for debugging support.
firewall	Can view the firewall filter configuration in configuration mode.
firewall-control	Can view and configure firewall filter information at the [edit firewall] hierarchy level.
floppy	Can read from and write to the removable media.
flow-tap	Can view the flow-tap configuration in configuration mode.
flow-tap-control	Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the [edit services flow-tap] hierarchy level.
flow-tap-operation	<p>Can make flow-tap requests to the router or switch. For example, a Dynamic Tasking Control Protocol (DTCP) client must authenticate itself to the Junos OS as an administrative user. That account must have flow-tap-operation permission.</p> <p>NOTE: The flow-tap-operation option is not included in the all-control permissions flag.</p>
idp-profiler-operation	Can view profiler data.
interface	Can view the interface configuration in configuration mode and with the show configuration operational mode command.

Table 164: Login Class Permission Flags (*continued*)

Permission Flag	Description
interface-control	Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the following hierarchy levels: <ul style="list-style-type: none"> • [edit chassis] • [edit class-of-service] • [edit groups] • [edit forwarding-options] • [edit interfaces]
maintenance	Can perform system maintenance, including starting a local shell on the router and becoming the superuser in the shell by using the su root command, and can halt and reboot the router by using the request system commands.
network	Can access the network by using the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view the pgcp session mirroring configuration.
pgcp-session-mirroring-control	Can modify the pgcp session mirroring configuration.
reset	Can restart software processes by using the restart command and can configure whether software processes are enabled or disabled at the [edit system processes] hierarchy level.
rollback	Can use the rollback command to return to a previously committed configuration other than the most recently committed one.
routing	Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.
routing-control	Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the [edit routing-options] hierarchy level, routing protocols at the [edit protocols] hierarchy level, and routing policy at the [edit policy-options] hierarchy level.
secret	Can view passwords and other authentication keys in the configuration.
secret-control	Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.
security	Can view security configuration in configuration mode and with the show configuration operational mode command.

Table 164: Login Class Permission Flags (*continued*)

Permission Flag	Description
security-control	Can view and configure security information at the [edit security] hierarchy level.
shell	Can start a local shell on the router or switch by using the start shell command.
snmp	Can view Simple Network Management Protocol (SNMP) configuration information in configuration and operational modes.
snmp-control	Can view SNMP configuration information and can modify SNMP configuration at the [edit snmp] hierarchy level.
system	Can view system-level information in configuration and operational modes.
system-control	Can view system-level configuration information and configure it at the [edit system] hierarchy level.
trace	Can view trace file settings and configure trace file properties.
trace-control	Can modify trace file settings and configure trace file properties.
view	Can use various commands to display current system-wide, routing table, and protocol-specific values and statistics. Cannot view the secret configuration.
view-configuration	Can view all of the configuration excluding secrets, system scripts, and event options. NOTE: Only users with the maintenance permission can view commit script, op script, or event script configuration.

Allowing or Denying Individual Commands for Junos OS Login Classes

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

- The **all** login class permission bits take precedence over extended regular expressions when a user with **rollback** permission issues the **rollback** command.
- Expressions used to allow and deny commands for users on RADIUS and TACACS+ servers have been simplified. Instead of a single, long expression with multiple commands (**allow-commands=cmd1 cmd2 ... cmdn**), you can specify each command as a separate expression. This new syntax is valid for **allow-configuration** and

deny-configuration, **allow-commands** and **deny-commands**, and all user permission bits.

- Users cannot issue the **load override** command when specifying an extended regular expression. Users can only issue the **merge**, **replace**, and **patch** configuration commands.
- If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by the **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.
- Regular expressions for **allow-commands** and **deny-commands** can also include the **commit**, **load**, **rollback**, **save**, **status**, and **update** commands.
- If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

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

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

Related Documentation

- [Configuring Access Privilege Levels on page 1692](#)

Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

If the **authentication-order** is remote-server then local, Junos OS will retry the local server if the remote-server is unreachable or has timed out. However, if the remote-server rejects the authentication, Junos OS will not retry the authentication.

If none of the configured authentication methods accept the login credentials and if a reject response is received, the login attempt fails. If no response is received from any configured authentication method, the Junos OS consults local password authentication as a last resort.

Using RADIUS or TACACS+ Authentication

You can configure the Junos OS to be both a RADIUS and TACACS+ authentication client.

If an authentication method included in the **[authentication-order]** statement is not available, or if the authentication is available but returns a reject response, the Junos OS tries the next authentication method included in the **authentication-order** statement.

The RADIUS or TACACS+ server authentication might fail because of the following reasons:

- The authentication method is configured, but the corresponding authentication servers are not configured. For instance, the RADIUS and TACACS+ authentication methods are included in the **authentication-order** statement, but the corresponding RADIUS or TACACS+ servers are not configured at the respective **[edit system radius-server]** and **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server does not respond within the timeout period configured at the **[edit system radius-server]** or **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server is not reachable because of a network problem.

The RADIUS or TACACS+ server authentication might return a reject response because of the following reasons:

- The user profiles of users accessing a router or switch might not be configured on the RADIUS or TACACS+ server.
- The user enters incorrect logon credentials.

Using Local Password Authentication

You can explicitly configure the password authentication method or use this method as a fallback mechanism when remote authentication servers fail. The password authentication method consults the local user profiles configured at the **[edit system login]** hierarchy level. Users can log in to a router or switch using their local username and password in the following scenarios:

- The password authentication method (password) is explicitly configured as one of the authentication methods in the **[authentication-order authentication-methods]** statement. In this case, the password authentication method is tried if no previous authentication accepts the logon credentials. This is true whether the previous authentication method fails to respond or returns a reject response because of an incorrect username or password.
- The password authentication method is not explicitly configured as one of the authentication methods in the **authentication-order authentication-methods** statement. In this case, the password authentication method is tried only if all configured authentication methods fail to respond. It is not consulted if any configured authentication method returns a reject response because of an incorrect username or password.

Order of Authentication Attempts

[Table 165 on page 1678](#) describes how the **authentication-order** statement at the **[edit system]** hierarchy level determines the procedure that the Junos OS uses to authenticate users for access to a router or switch.

Table 165: Order of Authentication Attempts

Syntax	Order of Authentication Attempts
authentication-order radius;	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS server is available but authentication is rejected, deny access. 4. If RADIUS servers are not available, try password authentication. <p>NOTE: If a RADIUS server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [radius tacplus];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ server is available but authentication is rejected, deny access. 6. If both RADIUS and TACACS+ servers are not available, try password authentication. <p>NOTE: If either RADIUS or TACACS+ servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [radius tacplus password];	<ol style="list-style-type: none"> 1. Try configured RADIUS authentication servers. 2. If RADIUS server is available and authentication is accepted, grant access. 3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers. 4. If TACACS+ server is available and authentication is accepted, grant access. 5. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.

Table 165: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
authentication-order tacplus;	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ server is available but authentication is rejected, deny access. 4. If TACACS+ servers are not available, try password authentication. <p>NOTE: If a TACACS+ server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.
authentication-order [tacplus radius];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS server is available but authentication is rejected, deny access. 6. If both TACACS+ and RADIUS servers are not available, try password authentication. <p>NOTE: If either TACACS+ or RADIUS servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
authentication-order [tacplus radius password];	<ol style="list-style-type: none"> 1. Try configured TACACS+ authentication servers. 2. If TACACS+ server is available and authentication is accepted, grant access. 3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers. 4. If RADIUS server is available and authentication is accepted, grant access. 5. If RADIUS servers fail to respond or return a reject response try password authentication, because it is explicitly configured in the authentication order.

Table 165: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<code>authentication-order password;</code>	<ol style="list-style-type: none"> 1. Try to authenticate the user, using the password configured at the <code>[edit system login]</code> hierarchy level. 2. If the authentication is accepted, grant access. 3. If the authentication is rejected, deny access.



NOTE: If SSH public keys are configured, SSH user authentication first tries to perform public key authentication before using the authentication methods configured in the `authentication-order` statement. If you want SSH logins to use the authentication methods configured in the `authentication-order` statement without first trying to perform public key authentication, do not configure SSH public keys.

In a routing matrix based on a TX Matrix router, the authentication order must be configured only at the configuration groups `re0` and `re1`. The authentication order must not be configured at the `[edit system]` hierarchy. This is because the authentication order for the routing matrix is controlled on the switch-card chassis (or TX Matrix router) or switch-fabric chassis (for TX Matrix Plus router) only.

In Junos OS Release 10.0 and later, the superuser (belonging to the super-user login class) is also authenticated based on the authentication order that is configured for TACACS+, RADIUS, or password authentication using the `authentication-order` statement. For example, if the only configured authentication order is TACACS+, the superuser can only be authenticated by the TACACS+ server and password authentication cannot be used as an alternative. However, in Junos OS Release 9.6 and earlier, the superuser can use password authentication to login, even if password authentication is not configured explicitly using the `authentication-order` statement.

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1694](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1716](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1734](#)

Junos OS User Authentication Methods

The Junos OS supports three methods of user authentication: local password authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+).

With local password authentication, you configure a password for each user allowed to log in to the router or switch.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router or switch using telnet. They are both distributed client-server systems—the RADIUS and TACACS+ clients run on the router or switch, and the server runs on a remote network system.

You can configure the router or switch to be both a RADIUS and TACACS+ client, and you can also configure authentication passwords in the Junos OS configuration file. You can prioritize the methods to configure the order in which the software tries the different authentication methods when verifying user access.

Related Documentation

- *Configuring RADIUS Authentication*
- *Configuring TACACS+ Authentication*
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1676](#)
- [Configuring RADIUS Authentication on page 1699](#)
- [Configuring TACACS+ Authentication on page 1711](#)

Junos OS User Accounts Overview

User accounts provide one way for users to access the switch. (Users can access the switch without accounts if you configured RADIUS or TACACS+ servers, as described in [“Junos OS User Authentication Methods” on page 1681](#).) For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

For each user account, you can define the following:

- Username—(Optional) Name that identifies the user. It must be unique within the switch. Do not include spaces, colons, or commas in the username. The username can be up to 64 characters long.
- User's full name—(Optional) If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
- User identifier (UID)—(Optional) Numeric identifier that is associated with the user account name. The identifier must be in the range from 100 through 64,000 and must

be unique within the switch. If you do not assign a UID to a username, the software assigns one when you commit the configuration, preferring the lowest available number.

- You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning when you commit the configuration and then assigns the duplicate UID.
- User's access privilege—(Required) One of the login classes you defined in the **class** statement at the **[edit system login]** hierarchy level, or one of the default classes listed in [“Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies” on page 1685](#).
- Authentication method or methods and passwords that the user can use to access the switch—(Optional) You can use SSH or a Message Digest 5 (MD5) password, or you can enter a plain-text password that Junos OS encrypts using MD5-style encryption before entering it in the password database. For each method, you can specify the user's password. If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system login user user-name]
user@switch# set authentication plain-text-password
New password: type password here
Retype new password: retype password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
 - You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Valid passwords must contain at least one change of case or character class.

For each user account and for root logins, you can configure more than one public RSA or DSA key for user authentication. When a user logs in using a user account or as root, the configured public keys are referenced to determine whether the private key matches any of them.

For SSH authentication, you can also copy the contents of an SSH key file into the configuration.

To load an SSH key file, use the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH key entries, use the configuration mode **show** command. For example:

```
[edit system login user boojum]
user@switch# set authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@switch# show
root-authentication {
```

```
ssh-rsa "1024 35 9727638204084251055468226757249864241630322
207404962528390382038690141584534964170019610608358722961563
475784918273603361276441874265946893207739108344813125957722
625461667999278316123500438660915866283822489746732605661192
181489539813862940327687806538169602027491641637359132693963
44008443 boojum@juniper.net"; # SECRET-DATA
}
```

An account for the user **root** is always present in the configuration. You configure the password for **root** using the **root-authentication** statement, as described in “[Configuring the Root Password](#)” on page 1702.

Junos-FIPS and Common Criteria have special password requirements. FIPS and Common Criteria passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the switch, you cannot configure passwords unless they meet this standard.

Related Documentation

- [Configuring Junos OS User Accounts on page 1692](#)
- [Junos OS Login Classes Overview on page 1683](#)

Junos OS Login Classes Overview

All users who can log in to the router or switch must be in a login class. With login classes, you define the following:

- Access privileges that users have when they are logged in to the router or switch
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out

You can define any number of login classes and then apply one login class to an individual user account.

The Junos operating system (Junos OS) contains a few predefined login classes, which are listed in [Table 166 on page 1683](#). The predefined login classes cannot be modified.

Table 166: Predefined System Login Classes

Login Class	Permission Flag Set
operator	clear, network, reset, trace, and view
read-only	view
superuser or super-user	all
unauthorized	None

**NOTE:**

- You cannot modify a predefined login class name. If you issue the `set` command on a predefined class name, the Junos OS appends `-local` to the login class name. The following message also appears:

warning: '<class-name>' is a predefined class name; changing to '<class-name>-local'

- You cannot issue the `rename` or `copy` command on a predefined login class. Doing so results in the following error message:

error: target '<class-name>' is a predefined class

Related Documentation

- [Defining Junos OS Login Classes](#)
- [Defining Junos OS Login Classes on page 1716](#)
- [Understanding QFabric System Login Classes on page 1230](#)

Understanding QFabric System Login Classes

In some cases (such as device-level troubleshooting), it is useful to log in to individual QFabric system components so you can view and manage issues on a per-device basis. This topic explains the login classes that provide individual component access within a QFabric system.



NOTE: Under normal operating conditions, you should manage the QFabric system as a single entity by using the QFabric system default partition command-line interface (CLI). The default partition CLI provides you with the ability to configure and monitor your entire QFabric system from a central location and should be used as the primary way to manage the system.

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system:

- **qfabric-admin**—Provides the ability to log in to individual QFabric system components and manage them. This class is equivalent to setting the following permissions: **access**, **admin**, **clear**, **firewall**, **interface**, **maintenance**, **network**, **reset**, **routing**, **secret**, **security**, **snmp**, **system**, **trace**, and **view**. The *qfabric-admin* class also enables you issue all operational mode commands except **configure**. To provide QFabric system component-level login and management privileges, include the **qfabric-admin** statement at the `[edit system login user username authentication remote-debug-permission]` hierarchy level.
- **qfabric-operator**—Provides the privilege to log in to individual QFabric system components and view component operations and configurations. This class is equivalent to setting the following permissions: **trace** and **view**. The *qfabric-operator* class also enables you issue the **monitor** and **show log messages** operational mode commands. To provide limited QFabric system component-level access, include the

qfabric-operator statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.

- **qfabric-user**—Prevents access to individual QFabric system components. This class is the default setting for all QFabric system users and is equivalent to the preset Junos OS class of **unauthorized**. To prevent a user from accessing individual QFabric system components, include the **qfabric-user** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.

When you perform the initial setup for the Director group, you must specify a username and password for QFabric components. Once configured, this information is stored in the QFabric system and mapped to the QFabric system login classes. Such mapping allows users with the proper login class (**qfabric-admin** or **qfabric-operator**) to log in automatically to a component without being prompted for the username and password.

After you assign the **qfabric-admin** or **qfabric-operator** class to a user, the user can log in to an individual QFabric system component by issuing the **request component login component-name** command. You can access Node devices, Interconnect devices, and virtual Junos Routing Engines (diagnostics, fabric control, and fabric manager) one at a time when you issue this command. To leave the CLI prompt of a component and return to the QFabric system default partition CLI, issue the **exit** command from the component's operational mode CLI prompt.

Related Documentation

- [Example: Configuring QFabric System Login Classes on page 1345](#)
- [remote-debug-permission on page 1401](#)
- [request component login on page 1444](#)
- [Junos OS Login Classes Overview on page 1683](#)

Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies

Use extended regular expressions to specify which configuration mode hierarchies are denied or allowed. You specify these regular expressions in the **allow/deny-configuration-regexps** and **allow/deny-configuration** statements at the **[edit system login class]** hierarchy level, or by specifying Juniper Networks vendor-specific TACACS+ or RADIUS attributes in your authentication server's configuration. If regular expressions are received during TACACS+ or RADIUS authentication, they merge with any regular expressions configured on the local router or switch.

[Table 167 on page 1686](#) lists common regular expression operators that you can use for allowing or denying configuration mode .

Command regular expressions implement the extended (modern) regular expressions, as defined in POSIX 1003.2.

Table 167: Configuration Mode Hierarchies—Common Regular Expression Operators

Operator	Match
	One of two or more terms separated by the pipe. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive .
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated the result is then evaluated as part of the overall expression. Parentheses must be used in conjunction with pipe operators as explained .
*	Zero or more terms.
+	One or more terms.
.	Any character except for a space " ".

Related Documentation

- *Specifying Access Privileges for Junos OS Configuration Mode Hierarchies*
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1719](#)

Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands

Use extended regular expressions to specify which operational mode commands are denied or allowed. [Table 168 on page 1686](#) lists common regular expression operators that can be used in the operational mode commands. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2.

Table 168: Common Regular Expression Operators to Allow or Deny Operational Mode Commands

Operator	Match
	One of two or more terms separated by the pipe () symbol. Each term must be a complete standalone expression enclosed in parentheses (), with no spaces between the pipe and the adjacent parentheses. For example, (show system alarms) (show system software).

Table 168: Common Regular Expression Operators to Allow or Deny Operational Mode Commands *(continued)*

Operator	Match
^	At the beginning of an expression, used to denote where the command begins, and where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces\$" means that the user can issue the show interfaces command but cannot issue the show interfaces detail or show interfaces extensive command.
[]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
()	A group of commands, indicating a complete, standalone expression to be evaluated; the result is then evaluated as part of the overall expression. Parentheses must always be used in conjunction with pipe operators as explained above.

If a regular expression contains a syntax error, it becomes invalid, and although the user can log in, the permission granted or denied by the regular expression does not take effect. When regular expressions configured on TACACS+ or RADIUS servers merge with regular expressions configured on the router or switch, if the final expression has a syntax error, the overall result is an invalid regular expression. If a regular expression does not contain any operators, all varieties of the command are allowed. For example, if the following statement is included in the configuration, the user can issue the commands **show interfaces detail** and **show interfaces extensive** in addition to showing an individual interface:

allow-commands "show interfaces";

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 178](#)

Special Requirements for Junos OS Plain-Text Passwords

Junos OS has special requirements when you create plain-text passwords on a router or switch. [Table 169 on page 1687](#) shows the default requirements.

Table 169: Special Requirements for Plain-Text Passwords

Junos OS	Junos-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters long
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.

Table 169: Special Requirements for Plain-Text Passwords (*continued*)

Junos OS	Junos-FIPS
Valid passwords must contain at least one change of case or character class.	Passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

You can change the requirements for plain-text passwords.

Junos OS supports the following five character classes for plain-text passwords:

- Lowercase letters
- Uppercase letters
- Numbers
- Punctuation
- Special characters: ! @ # \$ % ^ & *, + < > ; :

Control characters are not recommended.

You can include the **plain-text-password** statement at the following hierarchy levels:

- [edit system diag-port-authentication]
- [edit system pic-console-authentication]
- [edit system root-authentication]
- [edit system login user *username* authentication]

The **change-type** statement specifies whether the password is checked for the following:

- The total number of character sets used (**character-set**)
- The total number of character set changes (**set-transitions**)

For example, the following password:

MyPassWd@2

has four character sets (uppercase letters, lowercase letters, special characters, and numbers) and seven character set changes (**M–y**, **y–P**, **P–a**, **a–W**, **W–d**, **d–@**, and **@–2**).

The **change-type** statement is optional. If you omit the **change-type** option, Junos-FIPS plain-text passwords are checked for character sets, and Junos OS plain-text passwords are checked for character set changes.

The **minimum-changes** statement specifies how many character sets or character set changes are required for the password. This statement is optional. If you do not use the **minimum-changes** statement, character sets are not checked for Junos OS. If the **change-type** statement is configured for the **character-set** option, then the **minimum-changes** value must be 5 or less, because Junos OS only supports five character sets.

The **format** statement specifies the hash algorithm (**md5**, **sha1** or **des**) for authenticating plain-text passwords. This statement is optional. For Junos OS, the default format is **md5**. For Junos-FIPS, only **sha1** is supported.

The **maximum-length** statement specifies the maximum number of characters allowed in a password. This statement is optional. By default, Junos OS passwords have no maximum; however, only the first 128 characters are significant. Junos-FIPS passwords must be 20 characters or less. The range for Junos OS maximum-length passwords is from 20 to 128 characters.

The **minimum-length** statement specifies the minimum number of characters required for a password. This statement is optional. By default, Junos OS passwords must be at least 6 characters long, and Junos-FIPS passwords must be at least 10 characters long. The range is from 6 to 20 characters.

Changes to password requirements do not take effect until the configuration is committed. When requirements change, only newly created, plain-text passwords are checked; existing passwords are not checked against the new requirements.

The default configuration for Junos OS plain-text passwords is:

```
[edit system login]
passwords {
  change-type character-sets;
  format md5;
  minimum-changes 1;
  minimum-length 6;
}
```

The default configuration for Junos-FIPS plain-text passwords is:

```
[edit system login]
passwords {
  change-type set-transitions;
  format sha1;
  maximum-length 20;
  minimum-changes 3;
  minimum-length 10;
}
```

Related Documentation

- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
- [Configuring the Root Password on page 158](#)
- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
- [Configuring the Root Password on page 1702](#)

CHAPTER 18

Configuration

- [Configuration Tasks on page 1691](#)
- [Configuration Examples on page 1726](#)
- [Configuration Statements on page 1747](#)

Configuration Tasks

- [Configuring Access Privilege Levels on page 1692](#)
- [Configuring CLI Tips on page 1692](#)
- [Configuring Junos OS User Accounts on page 1692](#)
- [Configuring LLDP on page 1693](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1694](#)
- [Configuring Local User Template Accounts for User Authentication on page 1695](#)
- [Configuring Management Access on page 1697](#)
- [Configuring RADIUS System Accounting on page 1697](#)
- [Configuring RADIUS Authentication on page 1699](#)
- [Configuring Remote Template Accounts for User Authentication on page 1701](#)
- [Configuring the Root Password on page 1702](#)
- [Configuring SNMP on page 1703](#)
- [Configuring SSH Host Keys for Secure Copying of Data on page 1707](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 1709](#)
- [Configuring TACACS+ Authentication on page 1711](#)
- [Configuring TACACS+ System Accounting on page 1714](#)
- [Defining Junos OS Login Classes on page 1716](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1716](#)
- [Recovering the Root Password on page 1717](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1719](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 1720](#)
- [Using Junos OS to Configure Logical System Administrators on page 1721](#)

- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1722](#)
- [VSA Match Conditions and Actions on page 1724](#)

Configuring Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with it. Users can execute only those commands and configure and view only those statements for which they have access privileges.

To configure access privilege levels, include the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level:

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

Related Documentation

- [Example: Configuring Access Privilege Levels on page 1729](#)
- [Understanding Junos OS Access Privilege Levels on page 1672](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 178](#)
- *permissions*

Configuring CLI Tips

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

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

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

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

Related Documentation

- [CLI User Interface Overview on page 70](#)
- [Defining Junos OS Login Classes](#)
- [login-tip on page 265](#)

Configuring Junos OS User Accounts

User accounts provide one way for users to access the router or switch. For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

To create user accounts, include the **user** statement at the **[edit system login]** hierarchy level:


```
[edit system login]
user username {
  class class-name;
  class {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
  }
  full-name complete-name;
  uid uid-value;
  class class-name;
}
```

Related Documentation

- [Example: Configuring User Accounts on page 1733](#)
- [Example: Configuring User Login Accounts on page 1744](#)
- [Junos OS User Accounts Overview on page 1681](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1716](#)

Configuring LLDP

QFX Series products use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to identify a variety of devices quickly. The result is a LAN that interoperates smoothly and efficiently.

The LLDP protocol cannot be enabled by issuing the **set protocols lldp** statement at the **[edit]** hierarchy level. Enable the LLDP protocol by configuring it on all interfaces or on specific interfaces.

To configure basic LLDP options using the CLI:

1. Configure the advertisement interval in seconds:

```
[edit protocols lldp]
user@switch# set advertisement-interval 45
```

2. Specify the multiplier used in combination with the **advertisement-interval** value to determine the length of time LLDP information is held before it is discarded:

```
[edit protocols lldp]
user@switch# set hold-multiplier 5
```

3. Configure LLDP on all interfaces or on a specific interface:

```
[edit protocols lldp]
user@switch# set interface (LLDP) all
```

4. Configure tracing operations for the LLDP protocol:

```
[edit protocols lldp]
user@switch# set traceoptions file lldptrace
```

Related Documentation

Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

To configure the authentication order, include the **authentication-order** statement at the **[edit system]** hierarchy level:

```
[edit system]
authentication-order [ authentication-methods ];
```

Specify one or more of the following authentication methods in the preferred order, from first tried to last tried:

- **radius**—Verify the user using RADIUS authentication services
- **tacplus**—Verify the user using TACACS+ authentication services.
- **password**—Verify the user using the username and password configured locally by including the authentication statement at the **[edit system login user]** hierarchy level.

The CHAP authentication sequence cannot take more than 30 seconds. If it takes longer to authenticate a client, the authentication is abandoned and a new sequence is initiated.

For example, if you configure three RADIUS servers so that the router or switch attempts to contact each server three times, and with each retry the server times out after 3 seconds, then the maximum time given to the RADIUS authentication method before CHAP considers it a failure is 27 seconds. If you add more RADIUS servers to this configuration, they might not be contacted because the authentication process might be abandoned before these servers are tried.

The Junos OS enforces a limit on the number of standing authentication server requests that the CHAP authentication can have at one time. Thus, an authentication server method—RADIUS, for example—might fail to authenticate a client when this limit is exceeded. If it fails, the authentication sequence is reinitiated by the router or switch until authentication succeeds and the link is brought up. However, if the RADIUS servers are not available and if additional authentication methods such as **tacplus** or **password** are configured along with **radius**, the next authentication method is tried.

The following example shows how to configure **radius** and **password** authentication:

```
[edit system]
user@switch# authentication-order [ radius password ];
```

The following example shows how to delete the **radius** statement from the authentication order:

```
[edit system]
user@switch# delete authentication-order radius
```

The following example shows how to insert the **tacplus** statement after the **radius** statement:

```
[edit system]
user@switch# insert authentication-order tacplus after radius
```

Related Documentation

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1676](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1722](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1734](#)
- *authentication-order*

Configuring Local User Template Accounts for User Authentication

You use local user template accounts when you need different types of templates for authentication. Each template can define a different set of permissions appropriate for the group of users who use that template. These templates are defined locally on the router and referenced by the TACACS+ and RADIUS authentication servers.

When you configure local user templates and a user logs in, the Junos OS issues a request to the authentication server to authenticate the user's login name. If a user is authenticated, the server returns the local username to the Junos OS, which then determines whether a local username is specified for that login name (**local-username** for TACACS+, **Juniper-Local-User** for RADIUS). If so, the Junos OS selects the appropriate local user template locally configured on the router. If a local user template does not exist for the authenticated user, the router defaults to the **remote** template.

To configure different access privileges for users who share the local user template account, include the **allow-commands** and **deny-commands** commands in the authentication server configuration file.

To configure a local user template, include the **user local-username** statement at the **[edit system login]** hierarchy level and specify the privileges you want to grant to the local users to whom the template applies:

```
[edit system login]
user local-username {
  full-name "Local user account";
  uid uid-value;
  class class-name;
}
```

This example configures the **sales** and **engineering** local user templates:

```
[edit]
system {
  login {
    user sales {
      uid uid-value;
```

```
        class class-name;
    }
    user engineering {
        uid uid-value;
        class class-name;
    }
}

user = simon {
    ...
    service = junos-exec {
        local-user-name = sales
        allow-commands = "configure"
        deny-commands = "shutdown"
    }
}

user = rob {
    ...
    service = junos-exec {
        local-user-name = sales
        allow-commands = "(request system) | (show rip neighbor)"
        deny-commands = "clear"
    }
}

user = harold {
    ...
    service = junos-exec {
        local-user-name = engineering
        allow-commands = "monitor | help | show | ping | traceroute"
        deny-commands = "configure"
    }
}

user = jim {
    ...
    service = junos-exec {
        local-user-name = engineering
        allow-commands = "show bgp neighbor"
        deny-commands = "telnet | ssh"
    }
}
```

When the login users Simon and Rob are authenticated, the switch applies the sales local user template. When login users Harold and Jim are authenticated, the switch applies the engineering local user template.

**Related
Documentation**

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [user \(Access\)](#)
- [user \(Access\) on page 314](#)

Configuring Management Access

To define the management access settings for the routing platform:

1. Next to Allow Telnet Access, select the check box to allow remote Telnet access to the routing platform.
2. Next to Allow SSH Access, selected the check box to allow remote SSH access to the routing platform.
3. Click **Apply** to apply the configuration.

Configuring RADIUS System Accounting

With RADIUS accounting enabled, Juniper Networks routers or switches, acting as RADIUS clients, can notify the RADIUS server about user activities such as software logins, configuration changes, and interactive commands. The framework for RADIUS accounting is described in RFC 2866.

Tasks for configuring RADIUS system accounting are:

1. [Configuring Auditing of User Events on a RADIUS Server on page 1697](#)
2. [Specifying RADIUS Server Accounting and Auditing Events on page 1697](#)
3. [Configuring RADIUS Server Accounting on page 1698](#)

Configuring Auditing of User Events on a RADIUS Server

To audit user events, include the following statements at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
destination {
  radius {
    server {
      server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number;
        timeout seconds;
      }
    }
  }
}
```

Specifying RADIUS Server Accounting and Auditing Events

To specify the events you want to audit when using a RADIUS server for authentication, include the **events** statement at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
```

events is one or more of the following:

- **login**—Audit logins
- **change-log**—Audit configuration changes
- **interactive-commands**—Audit interactive commands (any command-line input)

Configuring RADIUS Server Accounting

To configure RADIUS server accounting, include the **server** statement at the [edit system accounting destination radius] hierarchy level:

```
server {  
  server-address {  
    accounting-port port-number;  
    secret password;  
    source-address address;  
    retry number;  
    timeout seconds;  
  }  
}
```

server-address specifies the address of the RADIUS server. To configure multiple RADIUS servers, include multiple **server** statements.



NOTE: If no RADIUS servers are configured at the [edit system accounting destination radius] statement hierarchy level, the Junos OS uses the RADIUS servers configured at the [edit system radius-server] hierarchy level.

accounting-port *port-number* specifies the RADIUS server accounting port number.

The default port number is 1813.



NOTE: If you enable RADIUS accounting at the [edit access profile *profile-name* accounting-order] hierarchy level, accounting is triggered on the default port of 1813 even if you do not specify a value for the **accounting-port** statement.

You must specify a secret (password) that the local router or switch passes to the RADIUS client by including the **secret** statement. If the password contains spaces, enclose the entire password in quotation marks (" ").

In the **source-address** statement, specify a source address for the RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address configured on one of the router or switch interfaces.

Optionally, you can specify the number of times that the router or switch attempts to contact a RADIUS authentication server by including the **retry** statement. By default, the router or switch retries three times. You can configure the router or switch to retry from 1 through 10 times.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a RADIUS server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

Configuring RADIUS Authentication

RADIUS authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure RADIUS authentication are:



NOTE: The **source-address** statement is not supported at the **[edit system radius-options]** or **[edit system-radius-server *name*]** hierarchies on the QFabric system.

- [Configuring RADIUS Server Details on page 1699](#)
- [Configuring MS-CHAPv2 for Password-Change Support on page 1700](#)
- [Specifying a Source Address for the Junos OS to Access External RADIUS Servers on page 1701](#)

Configuring RADIUS Server Details

To use RADIUS authentication on the router or switch, configure information about one or more RADIUS servers on the network by including one **radius-server** statement at the **[edit system]** hierarchy level for each RADIUS server:

```
[edit system]
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

server-address is the address of the RADIUS server.

You can specify a port on which to contact the RADIUS server. By default, port number **1812** is used (as specified in RFC 2865). You can also specify an accounting port to send accounting packets. The default is **1813** (as specified in RFC 2866).

You must specify a password in the **secret password** statement. If the password contains spaces, enclose it in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the amount of time that the local router or switch waits to receive a response from a RADIUS server (in the **timeout** statement) and the number of times that the router or switch attempts to contact a RADIUS authentication server (in the **retry** statement). By default, the router or switch waits 3 seconds. You can configure this to be a value from 1 through 90 seconds. By default, the router or switch retries

connecting to the server three times. You can configure this to be a value from 1 through 10 times.

You can use the **source-address** statement to specify a logical address for individual or multiple RADIUS servers.

To configure multiple RADIUS servers, include multiple **radius-server** statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the **user** statement at the **[edit system login]** hierarchy level, as described in [“Overview of Template Accounts for RADIUS and TACACS+ Authentication” on page 1665](#).

You can also configure RADIUS authentication at the **[edit access]** and **[edit access profile]** hierarchy level. Junos OS uses the following search order to determine which set of servers are used for authentication:

1. **[edit access profile *profile-name* radius-server *server-address*]**
2. **[edit access radius-server *server-address*]**
3. **[edit system radius-server *server-address*]**

Configuring MS-CHAPv2 for Password-Change Support

You can configure the Microsoft implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the router or switch to support changing of passwords. This feature provides users accessing a router or switch the option of changing the password when the password expires, is reset, or is configured to be changed at the next login.

Before you configure MS-CHAPv2 for password-change support, ensure that you:

- Configure the RADIUS server authentication parameters
- Set the **authentication-order** to use the RADIUS server for the initial password attempt

To configure MS-CHAP-v2, include the following statements at the **[edit system radius-options]** hierarchy level:

```
[edit system radius-options]
password-protocol mschap-v2;
```

The following example shows statements for configuring the MS-CHAPv2 password protocol, password authentication order, and user accounts:

```
[edit]
system {
  authentication-order [ radius password ];
  radius-server {
    192.168.69.149 secret "$9$G-j.5Qz6tpBk.1hrlXxUjiq5Qn/C"; ## SECRET-DATA
  }
  radius-options {
    password-protocol mschap-v2;
  }
}
```



```
login {
  user bob {
    class operator;
  }
}
```

Specifying a Source Address for the Junos OS to Access External RADIUS Servers

You can specify which source address Junos OS uses when accessing your network to contact an external RADIUS server for authentication. You can also specify which source address Junos OS uses when contacting a RADIUS server for sending accounting information.

To specify a source address for a RADIUS server, include the **source-address** statement at the **[edit system radius-server server-address]** hierarchy level:

```
[edit system radius-server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.

Related Documentation

- [Example: Configuring RADIUS Authentication on page 1730](#)
- [Example: Configuring RADIUS Authentication on a QFabric System on page 1731](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1734](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1668](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [Example: Configuring RADIUS Template Accounts on page 1745](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1722](#)
- [Junos OS User Authentication Methods on page 1681](#)

Configuring Remote Template Accounts for User Authentication

By default, the Junos OS uses remote template accounts for user authentication when:

- The authenticated user does not exist locally on the router or switch.
- The authenticated user's record in the authentication server specifies local user, or the specified local user does not exist locally on the router or switch.

To configure the remote template account, include the **user remote** statement at the **[edit system login]** hierarchy level and specify the privileges you want to grant to remote users:

```
[edit system login]
user remote {
  full-name "All remote users";
  uid uid-value;
```

```
class class-name;  
}
```

To configure different access privileges for users who share the remote template account, include the **allow-commands** and **deny-commands** statements in the authentication server configuration file.

**Related
Documentation**

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [user \(Access\)](#)
- [user \(Access\) on page 314](#)

Configuring the Root Password

Junos OS is preinstalled on the router or switch. When the router or switch is powered on, it is ready to be configured. Initially, you log in as the user “root” with no password.



NOTE: If you configure a blank password using the **encrypted-password** statement at the **[edit system root-authentication]** hierarchy level for root authentication, you can commit a configuration, but you are *not* able to log in as superuser and gain root level access to the router or switch.

After you log in, you should configure the root (superuser) password by including the **root-authentication** statement at the **[edit system]** hierarchy level:

```
[edit system]  
root-authentication {  
  (encrypted-password "password" | load-key-password URL | plain-text-password);  
  ssh-dsa "public-key";  
  ssh-rsa "public-key";  
}
```

If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system]  
user@switch# set root-authentication plain-text-password  
New password: type password here  
Retype new password: retry password here
```

To load an SSH key file, enter the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

You can also configure SSH RSA keys and SSH DSA keys to authenticate root logins. You can configure more than one public RSA or DSA key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH keys entries, use the configuration mode **show** command. For example:

```
[edit system]
user@switch# set root-authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@switch# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
  SECRET-DATA
}
```

Junos-FIPS software has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the router or switch, you cannot configure passwords unless they meet this standard. If you use the **encrypted-password** option, then a null-password (empty) is not permitted.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

Related Documentation

- [Recovering the Root Password on page 1159](#)
- [Example: Configuring the Root Password on page 1733](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 188](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1733](#)

Configuring SNMP

SNMP is implemented in the Junos OS Software running on the QFX Series products. By default, SNMP is not enabled. To enable SNMP, you must include the SNMP configuration statements at the **[edit]** hierarchy level.

To configure the minimum requirements for SNMP, include the following statements at the **[edit]** hierarchy level of the configuration:

```
[edit]
snmp {
  community public;
}
```

To configure complete SNMP features, include the following statements at the **[edit]** hierarchy level of the configuration:

```
snmp {
  client-list client-list-name {
    ip-addresses;
  }
}
```

```
community community-name {
  authorization authorization;
  client-list-name client-list-name;
  clients {
    address restrict;
  }
  logical-system logical-system-name {
    routing-instance routing-instance-name {
      clients {
        addresses;
      }
    }
  }
  routing-instance routing-instance-name {
    clients {
      addresses;
    }
  }
  view view-name;
}
contact contact;
description description;
filter-duplicates;
filter-interfaces;
health-monitor {
  falling-threshold integer;
  interval seconds;
  rising-threshold integer;
}
interface [ interface-names ];
location location;
name name;
nonvolatile {
  commit-delay seconds;
}
rmon {
  alarm index {
    description description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    request-type;
    rising-event-index index;
    rising-threshold integer;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
    syslog-subtag syslog-subtag;
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type type;
  }
  history history-index {
```

```

    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
    }
  }
}

```

```
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
      privacy-aes128 {
        privacy-password privacy-password;
      }
      privacy-des {
        privacy-password privacy-password;
      }
      privacy-3des {
        privacy-password privacy-password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
      privacy-aes128 {
        privacy-password privacy-password;
      }
      privacy-des {
        privacy-password privacy-password;
      }
      privacy-3des {
        privacy-password privacy-password;
      }
      privacy-none {
        privacy-password privacy-password;
      }
    }
  }
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
          }
        }
      }
    }
  }
}
```

```

        read-view view-name;
        write-view view-name;
    }
}
}
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
view view-name {
    oid object-identifier (include | exclude);
}
}

```

**Related
Documentation**

- [Understanding the Implementation of SNMP on page 6107](#)
- [snmp on page 1810](#)

Configuring SSH Host Keys for Secure Copying of Data

Secure Shell (SSH) uses encryption algorithms to generate a host, server, and session key system that ensures secure data transfer. You can configure SSH host keys to support secure copy (SCP) as an alternative to FTP for the background transfer of data such as configuration archives and event logs. To configure SSH support for SCP, you must complete the following tasks:

- Specify SSH known hosts by including hostnames and host key information in the Routing Engine configuration hierarchy.
- Set an SCP URL to specify the host from which to receive data. Setting this attribute automatically retrieves SSH host key information from the SCP server.
- Verify that the host key is authentic.
- Accept the secure connection. Accepting this connection automatically stores host key information in the local host key database. Storing host key information in the configuration hierarchy automates the secure handshake and allows background data transfer using SCP.

Tasks to configure SSH host keys for secure copying of data are:

1. [Configuring SSH Known Hosts on page 1708](#)
2. [Configuring Support for SCP File Transfer on page 1708](#)
3. [Updating SSH Host Key Information on page 1709](#)

Configuring SSH Known Hosts

To configure SSH known hosts, include the **host** statement, and specify hostname and host key options for trusted servers at the **[edit security ssh-known-hosts]** hierarchy level:

```
[edit security ssh-known-hosts]
host corporate-archive-server, ip-address {
    dsa-key key;
}
host archive-server-url {
    rsa-key key;
}
host server-with-ssh-version-1, ip-address {
    rsa1-key key;
}
```

Host keys are one of the following:

- **dsa-key**—Base64 encoded Digital Signature Algorithm (DSA) key.
- **rsa-key**—Base64 encoded RSA public key algorithm, which supports encryption and digital signatures.
- **rsa1-key**—Base64 encoded RSA public key algorithm, which supports encryption and digital signatures for SSH version 1 and SSH version 2.

Configuring Support for SCP File Transfer

To configure a known host to support background SCP file transfers, include the **archive-sites** statement at the **[edit system archival configuration]** hierarchy level.

```
[edit system archival configuration]
archive-sites {
    scp://username<:password>@host<:port>/url-path;
}
```



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "scp://username<:password>@[host]<:port>/url-path";

Setting the **archive-sites** statement to point to an SCP URL triggers automatic host key retrieval. At this point, Junos OS connects to the SCP host to fetch the SSH public key, displays the host key message digest or fingerprint as output to the console, and terminates the connection to the server.

```
user@switch# set system archival configuration archive-sites "<scp-url-path>"
The authenticity of host <my-archive-server (<server-ip-address>)> can't be established.
RSA key fingerprint is <ascii-text key>. Are you sure you want to continue connecting
(yes/no)?
```

To verify that the host key is authentic, compare this fingerprint with a fingerprint that you obtain from the same host using a trusted source. If the fingerprints are identical,

accept the host key by entering **yes** at the prompt. The host key information is then stored in the Routing Engine configuration and supports background data transfers using SCP.

Updating SSH Host Key Information

Typically, SSH host key information is automatically retrieved when you set a URL attribute for SCP using the **archival configuration archive-sites** statement at the **[edit system]** hierarchy level. However, if you need to manually update the host key database, use one of the following methods.

1. [Retrieving Host Key Information Manually on page 1709](#)
2. [Importing Host Key Information from a File on page 1709](#)

Retrieving Host Key Information Manually

To manually retrieve SSH public host key information, use the **fetch-from-server** option with the **set security ssh-known-hosts** command. You must include a hostname attribute with the **set security ssh-known-hosts fetch-from-server** command to specify the host from which to retrieve the SSH public key.

```
user@switch# set security ssh-known-hosts fetch-from-server <hostname>
```

Importing Host Key Information from a File

To manually import SSH host key information from the known-hosts file located at **/var/tmp/known-hosts** on the server, include the **load-key-file** option with the **set security ssh-known-hosts** command. You must include the path to the **known-hosts** file with the **set security ssh-known-hosts load-key-file** command to specify the location from which to import host key information.

```
user@switch# set security ssh-known-hosts load-key-file /var/tmp/known-hosts
```

Configuring SSH Service for Remote Access to the Router or Switch

To configure the router or switch to accept SSH as an access service, include the **ssh** statement at the **[edit system services]** hierarchy level:

```
[edit system services]
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ...]
  client-alive-count-max number;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm <algorithm | no-algorithm>;
  key-exchange algorithm;
  macs algorithm;
  max-sessions-per-connection number;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login <allow | deny | deny-password>;
}
```

By default, the router or switch supports a limited number of simultaneous SSH sessions and connection attempts per minute. Use the following statements to change the defaults:

- **connection-limit *limit***—Maximum number of simultaneous connections per protocol (IPv4 and IPv6). The range is a value from 1 through 250. The default is 75. When you configure a connection limit, the limit is applicable to the number of SSH sessions per protocol (IPv4 and IPv6). For example, a connection limit of 10 allows 10 IPv6 SSH sessions and 10 IPv4 SSH sessions.
- **max-sessions-per-connection *number***—Include this statement to specify the maximum number of SSH sessions allowed per single SSH connection. This allows you to limit the number of cloned sessions tunneled within a single SSH connection. The default value is 10.
- **rate-limit *limit***—Maximum number of connection attempts accepted per minute (a value from 1 through 250). The default is 150. When you configure a rate limit, the limit is applicable to the number of connection attempts per protocol (IPv4 and IPv6). For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

For information about other configuration settings, see the following topics:

- [Configuring the Root Login Through SSH on page 1710](#)
- [Configuring the SSH Protocol Version on page 1710](#)
- [Configuring the Client Alive Mechanism on page 1711](#)

Configuring the Root Login Through SSH

By default, users are allowed to log in to the router or switch as **root** through SSH. To control user access through SSH, include the **root-login** statement at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
  root-login (allow | deny | deny-password);
```

allow—Allows users to log in to the router or switch as root through SSH. The default is **allow**.

deny—Disables users from logging in to the router or switch as root through SSH.

deny-password—Allows users to log in to the router or switch as root through SSH when the authentication method (for example, RSA) does not require a password.

Configuring the SSH Protocol Version

By default, both version 1 and version 2 of the SSH protocol are enabled. To configure the router or switch to use only version 1 of the SSH protocol, include the **protocol-version** statement and specify **v1** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
  protocol-version [ v1 ];
```

To configure the router or switch to use only version 2 of the SSH protocol, include the **protocol-version** statement and specify **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v2 ];
```

To explicitly configure the router or switch to use version 1 and 2 of the SSH protocol, include the **protocol-version** statement and specify **v1** and **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 v2 ];
```

For J Series Services Routers, the export license software supports SSH version 1 only.

Configuring the Client Alive Mechanism

The client alive mechanism is valuable when the client or server depends on knowing when a connection has become inactive. It differs from the standard keepalive mechanism because the client alive messages are sent through the encrypted channel. The client alive mechanism is not enabled at default. To enable it, configure the **client-alive-count-max** and the **client-alive-interval**. This option applies to SSH protocol version 2 only.

In the following example, unresponsive SSH clients will be disconnected after approximately 100 seconds (20 x 5).

```
[edit system services ssh]
client-alive-count-max 5;
client-alive-interval 20;
```

Configuring TACACS+ Authentication

TACACS+ authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure TACACS+ configuration are:

- [Configuring TACACS+ Server Details on page 1711](#)
- [Specifying a Source Address for the Junos OS to Access External TACACS+ Servers on page 1712](#)
- [Configuring the Same Authentication Service for Multiple TACACS+ Servers on page 1713](#)
- [Configuring Juniper Networks Vendor-Specific TACACS+ Attributes on page 1713](#)

Configuring TACACS+ Server Details

To use TACACS+ authentication on the router or switch, configure information about one or more TACACS+ servers on the network by including the **tacplus-server** statement at the **[edit system]** hierarchy level:

```
[edit system]
tacplus-server server-address {
  port port-number;
  secret password;
  single-connection;
```

```
    timeout seconds;  
}
```

server-address is the address of the TACACS+ server.

port-number is the TACACS+ server port number.

You must specify a secret (password) by using the **secret** statement. The local router or switch passes the **secret** to the TACACS+ client. If the password included spaces, enclose the password in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a TACACS+ server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

Optionally, you can use the **single-connection** statement to have the software maintain one open Transmission Control Protocol (TCP) connection to the server for multiple requests, rather than opening a connection for each connection attempt.



NOTE: Early versions of the TACACS+ server do not support the **single-connection** option. If you specify this option and the server does not support it, Junos OS will be unable to communicate with that TACACS+ server.

To configure multiple TACACS+ servers, include multiple **tacplus-server** statements.

On a TX Matrix router, TACACS+ accounting should be configured only under the groups **re0** and **re1**.



NOTE: Accounting should not be configured at the **[edit system]** hierarchy level; on a TX Matrix router, control is done under the switch-card chassis only.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the **user** statement at the **[edit system login]** hierarchy level.

Specifying a Source Address for the Junos OS to Access External TACACS+ Servers

You can specify which source address Junos OS uses when accessing your network to contact an external TACACS+ server for authentication. You can also specify which source address Junos OS uses when contacting a TACACS+ server for sending accounting information.

To specify a source address for a TACACS+ server for authentication, include the **source-address** statement at the **[edit system tacplus-server server-address]** hierarchy level:

```
[edit system tacplus-server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.

To specify a source address for a TACACS+ server for system accounting, include the **source-address** statement at the **[edit system accounting destination tacplus server server-address]** hierarchy level:

```
[edit system accounting destination tacplus server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router or switch interfaces.

Configuring the Same Authentication Service for Multiple TACACS+ Servers

To configure the same authentication service for multiple TACACS+ servers, include statements at the **[edit system tacplus-server]** and **[edit system tacplus-options]** hierarchy levels. For information about how to configure a TACACS+ server at the **[edit system tacplus-server]** hierarchy level.

To assign the same authentication service to multiple TACACS+ servers, include the **service-name** statement at the **[edit system tacplus-options]** hierarchy level:

```
[edit system tacplus-options]
service-name service-name;
```

service-name is the name of the authentication service. By default, the service name is set to **junos-exec**.

The following example shows how to configure the same authentication service for multiple TACACS+ servers:

```
[edit system]
tacplus-server {
  10.2.2.2 secret "$9$2dgoJGDiqP5ZG9A"; ## SECRET-DATA
  10.3.3.3 secret "$9$2dgoJGDiqP5ZG9A"; ## SECRET-DATA
}
tacplus-options {
  service-name bob;
}
```

Configuring Juniper Networks Vendor-Specific TACACS+ Attributes

The Juniper Networks vendor-specific TACACS+ attributes enable you to configure access privileges for users on a TACACS+ server. They are specified in the TACACS+ server configuration file on a per-user basis. Junos OS retrieves these attributes through an authorization request of the TACACS+ server after authenticating a user. You do not need to configure these attributes to run Junos OS with TACACS+.

To specify these attributes, include a **service** statement of the following form in the TACACS+ server configuration file:

```
service = junos-exec {
  local-user-name = <username-local-to-router>
  allow-commands = "<allow-commands-regex>"
}
```

```
allow-configuration = "<allow-configuration-regex>"
deny-commands = "<deny-commands-regex>"
deny-configuration = "<deny-configuration-regex>"
}
```

This **service** statement can appear in a **user** or **group** statement.

**Related
Documentation**

- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1722](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1734](#)
- [Juniper Networks Vendor-Specific TACACS+ Attributes on page 1670](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)
- [Junos OS User Authentication Methods on page 1681](#)

Configuring TACACS+ System Accounting

You can use TACACS+ to track and log software logins, configuration changes, and interactive commands. To audit these events, include the following statements at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
destination {
  tacplus {
    server {
      server-address {
        port port-number;
        secret password;
        single-connection;
        timeout seconds;
      }
    }
  }
}
```

Tasks for configuring TACACS+ system accounting are:

1. [Specifying TACACS+ Auditing and Accounting Events on page 1714](#)
2. [Configuring TACACS+ Server Accounting on page 1715](#)

Specifying TACACS+ Auditing and Accounting Events

To specify the events you want to audit when using a TACACS+ server for authentication, include the **events** statement at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
```

events is one or more of the following:

- **login**—Audit logins

- **change-log**—Audit configuration changes
- **interactive-commands**—Audit interactive commands (any command-line input)

Configuring TACACS+ Server Accounting

To configure TACACS+ server accounting, include the **server** statement at the **[edit system accounting destination tacplus]** hierarchy level:

```
[edit system accounting destination tacplus]
server {
  server-address {
    port port-number;
    secret password;
    single-connection;
    timeout seconds;
  }
}
```

server-address specifies the address of the TACACS+ server. To configure multiple TACACS+ servers, include multiple **server** statements.



NOTE: If no TACACS+ servers are configured at the **[edit system accounting destination tacplus]** statement hierarchy level, Junos OS uses the TACACS+ servers configured at the **[edit system tacplus-server]** hierarchy level.

port-number specifies the TACACS+ server port number.

You must specify a secret (password) by using the **secret** statement. The local router or switch passes the **secret** to the TACACS+ client. If the password contains spaces, enclose the entire password in quotation marks (" "). The password used by the local router or switch must match that used by the server.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a TACACS+ server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

Optionally, you can maintain one open TCP connection to the server for multiple requests, rather than opening a connection for each connection attempt, by including the **single-connection** statement.

To ensure that start and stop requests for accounting of login events are correctly logged in the Accounting file instead of the Administration log file on a TACACS+ server, include either the **no-cmd-attribute-value** statement or the **exclude-cmd-attribute** at the **[edit system tacplus-options]** hierarchy level.

If you use the **no-cmd-attribute-value** statement, the value of the **cmd** attribute is set to a null string in the start and stop requests. If you use the **exclude-cmd-attribute** statement, the **cmd** attribute is totally excluded from the start and stop requests. Both statements

support the correct logging of accounting requests in the Accounting file, instead of the Administration file.

```
[edit system tacplus-options]
(no-cmd-attribute-value | exclude-cmd-attribute);
```

Related Documentation

- [Configuring TACACS+ Authentication on page 1711](#)

Defining Junos OS Login Classes

To define a login class and its access privileges, include the **class** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
class class-name {
  allow-commands "regular-expression";
  allow-configuration "regular-expression";
  deny-commands "regular-expression";
  deny-configuration "regular-expression";
  idle-timeout minutes;
  permissions [ permissions ];
}
```

Related Documentation

- [Junos OS Login Classes Overview on page 1683](#)
- [Junos OS User Accounts Overview on page 1681](#)
- [Example: Creating Login Classes with Specific Privileges on page 1736](#)
- [Configuring the Junos OS to Display a System Login Announcement on page 149](#)

Limiting the Number of User Login Attempts for SSH and Telnet Sessions

You can limit the number of times a user can attempt to enter a password while logging in through SSH or Telnet. The connection is terminated if a user fails to log in after the number of attempts specified. You can also specify a delay, in seconds, before a user can try to enter a password after a failed attempt. In addition, you can specify the threshold for the number of failed attempts before the user experiences a delay in being able to enter a password again.

To specify the number of times a user can attempt to enter a password while logging in, include the **retry-options** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
retry-options {
  tries-before-disconnect number;
  backoff-threshold number;
  backoff-factor seconds;
  maximum-time seconds;
  minimum-time seconds;
}
```

You can configure the following options:

- **tries-before-disconnect**—Number of times a user can attempt to enter a password when logging in. The connection closes if a user fails to log in after the number specified. The range is from 1 through 10, and the default is 10.
- **backoff-threshold**—Threshold for the number of failed login attempts before the user experiences a delay in being able to enter a password again. Use the **backoff-factor** option to specify the length of the delay in seconds. The range is from 1 through 3, and the default is 2.
- **backoff-factor**—Length of time, in seconds, before a user can attempt to log in after a failed attempt. The delay increases by the value specified for each subsequent attempt after the threshold. The range is from 5 through 10, and the default is 5 seconds.
- **maximum-time seconds**—Maximum length of time, in seconds, that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the **maximum-time** value, the connection is closed. The range is from 20 through 300 seconds, and the default is 120 seconds.
- **minimum-time**—Minimum length of time, in seconds, that a connection remains open while a user is attempting to enter a correct password. The range is from 20 through 60, and the default is 40.

Related Documentation

- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions on page 1746](#)
- [Configuring Junos OS User Accounts on page 1692](#)

Recovering the Root Password

If you forget the root password for the QFX3500 switch, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the QFX3500 switch power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the switch into the RJ-45-to-DB-9 serial port adapter supplied with the switch.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.

5. Connect the other end of the Ethernet rollover cable to the console port on the switch.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).

8. Configure the port settings as follows:

- Bits per second: 9600
- Data bits: 8
- Parity: None
- Stop bits: 1
- Flow control: None

9. Power on the switch by (if necessary) plugging the power cords into the QFX3500 switch power supply, or turning on the power to the device or switch by switching on the AC power outlet the device is plugged into

The terminal emulation screen on your management device displays the switch's boot sequence.

10. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN  
for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.

14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1  
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit  
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the switch.

```
Reboot the system? [y/n] y
```

- Related Documentation**
- [Configuring the Root Password on page 1702](#)

Specifying Access Privileges for Junos OS Configuration Mode Hierarchies

You can specify extended regular expressions with the **allow-configuration** and **deny-configuration** statements to define user access privileges to parts of the configuration hierarchy. Doing so overrides login class permission bits set for a user. You can also use wildcards to restrict access. When you define access privileges to parts of the configuration hierarchy, do the following tasks:

- Specify the full paths in the extended regular expressions with the **allow-configuration** and **deny-configuration** statements.
- Put parentheses around an extended regular expression that connects two or more expressions with the pipe | symbol. For example:

```
[edit system login class class-name]
user@switch# set deny-configuration "(system login class) | (system services)"
```



NOTE: Each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses (). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol. You cannot define access to keywords such as **set**, **edit**, or **activate**.

When you explicitly provide access to configuration mode hierarchies or regular expressions using the **allow-configuration** statement, you add to the regular permissions set with the **permissions** statement. If you explicitly deny access to configuration mode hierarchies or regular expressions using the **deny-configuration** statement, you remove permissions for the specified configuration mode hierarchy from the default permissions provided by the **permissions** statement.

To explicitly provide access to an individual configuration mode hierarchy that would otherwise be denied, include the **allow-configuration** statement at the **[edit system login class *class-name*]** hierarchy level:

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

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

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

You can include one **deny-configuration** and one **allow-configuration** statement in each login class.

If you allow and deny the same set of configuration hierarchy levels, regular expressions, or commands, the **allow-configuration** statement permissions take precedence over the

permissions specified by the **deny-configuration** statement. For example, if you include **allow-configuration "system services"** and **deny-configuration "system services"**, the login class user can continue to edit the configuration or issue commands at the **edit system services** hierarchy level.

**Related
Documentation**

- [Defining Access Privileges Using allow/deny-configuration Statements on page 1745](#)
- [Configuring Access Privilege Levels on page 1692](#)

Specifying Access Privileges for Junos OS Operational Mode Commands

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

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

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

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

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

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

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



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

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

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

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

includes the **request system software add** command, the specified login class user can perform all operations allowed by the **all** permissions flag, except installing software using the **request system software add** command.

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

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

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

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

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

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

Related Documentation

- [Example: Configuring Access Privileges for Operational Mode Commands on page 1729](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1686](#)
- *allow-commands*
- *deny-commands*

Using Junos OS to Configure Logical System Administrators

Using Junos OS, you can partition a single router or switch into multiple logical devices that perform independent routing or switching tasks. When creating logical systems, you must configure logical system administrators and interfaces, assign logical interfaces to logical systems, and configure various other logical system statements.

The master administrator can assign one or more logical system administrators to each logical system. Once assigned to a logical system, administrators are restricted to viewing

only configurations of the logical system to which they are assigned and accessing only the operational commands that apply to that particular logical system. This restriction means that these administrators cannot access global configuration statements, and all command output is restricted to the logical system to which the administrators are assigned.

To configure logical system administrators, include the **logical-system *logical-system-name*** statement at the **[edit system login class *class-name*]** hierarchy level and apply the class to the user. For example:

```
[edit]
system {
  login {
    class admin1 {
      permissions all;
      logical-system logical-system-LS1;
    }
    class admin2 {
      permissions view; # Gives users assigned to class admin2 the ability to view
                        # but not to change the configuration.
      logical-system logical-system-LS2;
    }
    user user1 {
      class admin1;
    }
    user user2 {
      class admin2;
    }
  }
}
```

Fully implementing logical systems requires that you also configure any protocols, routing statements, switching statements, and policy statements for the logical system.

- Related Documentation**
- [Defining Junos OS Login Classes](#)
 - [Defining Junos OS Login Classes on page 1716](#)

Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands

Use regular expressions to specify which operational or configuration mode commands are allowed or denied when you use a RADIUS or TACACS+ server for user authentication. You can specify the regular expressions using the appropriate Juniper Networks vendor-specific RADIUS or TACACS+ attributes in your authentication server configuration.

You can specify **allow-configuration**, **deny-configuration**, **allow-commands**, or **deny-commands** in a single extended regular expression, enclosing multiple commands in parentheses and separating them using the pipe symbol. For example, you can specify multiple **allow-commands** parameters using: **allow-commands= (cmd1 | cmd2 | cmdn)**. You can specify **user-permissions** as a list of comma-separated values, and not as a regular expression.

On a RADIUS or TACACS+ server, you can also use a simplified version for regular expressions where you specify each individual expression on a separate line. The simplified version is valid for **allow-commands**, **deny-commands**, **allow-configuration**, **deny-configuration**, and **permissions** vendor-specific attributes.

For a RADIUS server, specify the individual regular expressions using the following syntax:

```
Juniper-Allow-Commands+= "cmd1"
Juniper-Allow-Commands+= "cmd2"
Juniper-Allow-Commands+= "cmdn"
Juniper-Deny-Commands+= "cmd1"
Juniper-Deny-Commands+= "cmd2"
Juniper-Deny-Commands+= "cmdn"
Juniper-Allow-Configuration+= "regex1"
Juniper-Allow-Configuration+= "regex2"
Juniper-Allow-Configuration+= "regexn"
Juniper-Deny-Configuration+= "regex1"
Juniper-Deny-Configuration+= "regex2"
Juniper-Deny-Configuration+= "regexn"
Juniper-User-Permissions+= "permission-flag1"
Juniper-User-Permissions+= "permission-flag2"
Juniper-User-Permissions+= "permission-flagn"
```

For a TACACS+ server, specify the individual regular expressions using the following syntax:

```
allow-commands1="cmd1"
allow-commands2="cmd2"
allow-commandsn="cmdn"
deny-commands1="cmd1"
deny-commands2="cmd2"
deny-commandsn="cmdn"
allow-configuration1="regex1"
allow-configuration2="regex2"
allow-configurationn="regexn"
deny-configuration1="regex1"
deny-configuration2="regex2"
deny-configurationn="regexn"
user-permissions1="permission-flag1"
user-permissions2="permission-flag2"
user-permissionsn="permission-flagn "
```

**NOTE:**

- Numeric values 1 to *n* in the syntax (for a TACACS+ server) must be unique but need not be sequential. For example, the following syntax is valid:

```
allow-commands1="cmd1"  
allow-commands3="cmd3"  
allow-commands2="cmd2"  
deny-commands3="cmd3"  
deny-commands2="cmd2"  
deny-commands1="cmd1"
```

- The limit on the number of lines of individual regular expressions is imposed by the TACACS+ or RADIUS server.
- When you issue the `show cli authorization` command, the command output displays the regular expression in a single line, even if you specify each individual expression on a separate line.

For more information about Juniper Networks vendor-specific RADIUS and TACACS+ attributes, see [“Juniper Networks Vendor-Specific RADIUS Attributes” on page 1668](#) and [“Juniper Networks Vendor-Specific TACACS+ Attributes” on page 1670](#).



NOTE: When RADIUS or TACACS+ authentication is configured for a router, regular expressions configured on the RADIUS or TACACS+ server merge with any regular expressions configured on the local router at the [edit system login class] hierarchy level using the `allow-commands`, `deny-commands`, `allow-configuration`, `deny-configuration`, or `permissions` statements. If the final expression has a syntax error, the overall result is an invalid regular expression.

Related Documentation

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1676](#)

VSA Match Conditions and Actions

EX Series switches and the QFX Series support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs). They are configured on RADIUS servers and work in combination with 802.1X authentication. Using VSAs, you can apply port firewall filter attributes as a subset of match conditions and actions sent from the RADIUS server to the switch as a result of successful 802.1X authentication.

Each term in a VSA configured through the RADIUS server consists of *match conditions* and an *action*. Match conditions are the values or fields that the packet must contain. You can define single, multiple, or no match conditions. If no match conditions are specified for the term, the packet is accepted by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Allowed actions are to accept a packet or to discard a packet.

The following guidelines apply when you specify match conditions and actions for VSAs:

- Both **match** and **action** statements are mandatory.
- Any or all options (separated by commas) may be included in each **match** and **action** statement.
- Fields separated by commas will be ANDed if they are of a different type. The same types cannot be repeated.
- For OR cases (for example, match **10.1.1.0/24 OR 11.1.1.0/24**), apply multiple VSAs to the 802.1X supplicant.
- In order for the **forwarding-class** option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored.

[Table 170 on page 1725](#) describes the match conditions you can specify when configuring a VSA using the **match** command on the RADIUS server. The string that defines a match condition is called a *match statement*.

Table 170: Match Conditions

Option	Description
destination-mac <i>mac-address</i>	Destination media access control (MAC) address of the packet.
source-vlan <i>source-vlan</i>	Name of the source VLAN.
source-dot1q-tag <i>tag</i>	Tag value in the 802.1Q header, in the range 0 through 4095.
destination-ip <i>ip-address</i>	Address of the final destination node.
ip-protocol <i>protocol-id</i>	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms: ah , egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), tcp (6), or udp (17)
source-port <i>port</i>	TCP or User Datagram Protocol (UDP) source port field. Normally, you specify this match statement in conjunction with the ip-protocol match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text options listed under destination-port .

Table 170: Match Conditions (*continued*)

Option	Description
destination-port <i>port</i>	<p>TCP or UDP destination port field. Normally, you specify this match in conjunction with the ip-protocol match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cvspserver (2401), cmd (514), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), telnet (23), tacacs-ds (65), talk (517), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), zephyr-hm (2104)</p>

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. [Table 171 on page 1726](#) shows the actions that you can specify in a term.

Table 171: Actions for VSAs

Option	Description
(allow deny)	Accept a packet or discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.
forwarding-class <i>class-of-service</i>	<p>(Optional) Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> assured-forwarding best-effort expedited-forwarding network-control
loss-priority (low medium high)	(Optional) Set the packet loss priority (PLP) to low , medium , or high . Specify both the forwarding class and loss priority.

Related Documentation

- [Filtering 802.1X Supplicants Using RADIUS Server Attributes](#)
- [Understanding 802.1X and VSAs on EX Series Switches](#)
- [Understanding VSAs on the QFX Series on page 1668](#)

Configuration Examples

- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 1727](#)
- [Example: Configuring Access Privilege Levels on page 1729](#)
- [Example: Configuring Access Privileges for Operational Mode Commands on page 1729](#)

- [Example: Configuring a Plain-Text Password for Root Logins on page 1730](#)
- [Example: Configuring RADIUS Authentication on page 1730](#)
- [Example: Configuring RADIUS Authentication on a QFabric System on page 1731](#)
- [Example: Configuring RADIUS System Accounting on page 1732](#)
- [Example: Configuring the Root Password on page 1733](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1733](#)
- [Example: Configuring User Accounts on page 1733](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1734](#)
- [Example: Creating Login Classes with Specific Privileges on page 1736](#)
- [Example: Configuring QFabric System Login Classes on page 1737](#)
- [Example: Configuring User Login Accounts on page 1744](#)
- [Example: Configuring RADIUS Template Accounts on page 1745](#)
- [Defining Access Privileges Using allow/deny-configuration Statements on page 1745](#)
- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions on page 1746](#)

Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 1727](#)
- [Overview on page 1727](#)
- [Configuration on page 1727](#)

Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**, **minimum-punctuations**, or **minimum-upper-cases**.

Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

set system login password minimum-length 12

set system login password maximum-length 22

set system login password minimum-numeric 1

set system login password minimum-upper-cases 1

set system login password minimum-lower-cases 1

set system login password minimum-punctuations 1

Configuring Requirements for Plain-Text Passwords

Step-by-Step Procedure

This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the [system login password] hierarchy level.

```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.

```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.

```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.

```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

- Related Documentation**
- [Special Requirements for Junos OS Plain-Text Passwords on page 1687](#)
 - [password \(Login\) on page 278](#)

Example: Configuring Access Privilege Levels

Create two access privilege classes on the router or switch, one for configuring and viewing user accounts only and the second for configuring and viewing SNMP parameters only:

```
[edit]
system {
  login {
    class user-accounts {
      permissions [ configure admin admin-control ];
    }
    class network-mgmt {
      permissions [ configure snmp snmp-control ];
    }
  }
}
```

- Related Documentation**
- [Configuring Access Privilege Levels on page 1692](#)

Example: Configuring Access Privileges for Operational Mode Commands

The following example shows how to configure access privileges for different login classes for individual operational mode commands:

```
[edit]
system {
  # This login class has operator privileges and the additional ability
  # to reboot the router.
  login {
    # This login class has operator privileges and the additional ability to reboot the
    # router or switch.
    class operator-and-boot {
      permissions [ clear network reset trace view ];
      allow-commands "request system reboot";
    }
    # This login class has operator privileges but can't use any commands beginning
    # with "set" .
    # This login class has operator privileges
    # but cannot use any commands beginning with "set"
    class operator-no-set {
      permissions [ clear network reset trace view ];
      deny-commands "^set";
    }
  }
}
```

```
}
# This login class has operator privileges and can install software but not view
# BGP information, and can issue the show route command, without specifying
# commands or arguments under it.
class operator-and-install-but-no-bgp {
  permissions [ clear network reset trace view ];
  allow-commands "(request system software add)|(show route$)";
  deny-commands "show bgp";
}
}
```

Related Documentation • [Specifying Access Privileges for Junos OS Operational Mode Commands on page 178](#)

Example: Configuring a Plain-Text Password for Root Logins

The following example shows how to set a plain-text password for root logins:

```
[edit]
user@switch# set system root-authentication plain-text-password
New password: type root password
Retype new password: retype root password
[edit]
user@switch# show
system {
  root-authentication {
    encrypted-password "$1$14c5.$sBopasddsd0";
  }
}
```

Related Documentation • [Configuring the Root Password on page 158](#)

Example: Configuring RADIUS Authentication

The Junos OS supports two protocols for central authentication of users on multiple routers: RADIUS and TACACS+. We recommend RADIUS because it is a multivendor IETF standard, and its features are more widely accepted than those of TACACS+ or other proprietary systems. In addition, we recommend using a one-time-password system for increased security, and all vendors of these systems support RADIUS.

The Junos OS uses one or more template accounts to perform user authentication. You create the template account or accounts, and then configure the user access to use that account. If the RADIUS server is unavailable, the fallback is for the login process to use the local account that set up on the router or switch.

The following example shows how to configure RADIUS authentication:

```
[edit]
system {
  authentication-order [ radius password ];
  root-authentication {
    encrypted-password "$9$aH1j8gqQ1gjyghgijgiiii"; # SECRET-DATA
  }
}
```

```

name-server {
  10.1.1.1;
  10.1.1.2;
}

```

The following example shows how to enable RADIUS authentication and define the shared secret between the client and the server. The secret enables the client and server to determine that they are talking to the trusted peer.

Define a timeout value for each server, so that if there is no response within the specified number of seconds, the router can try either the next server or the next authentication mechanism.

```

[edit]
system {
  radius-server {
    10.1.2.1 {
      secret "$9$aHlj8gqQ1sdjerrhser"; # SECRET-DATA
      timeout 5;
    }
    10.1.2.2 {
      secret "$9$aHlj8gqQ1csdoiuardwefoiud"; # SECRET-DATA
      timeout 5;
    }
  }
}

```

Related Documentation

- *Configuring RADIUS Authentication*

Example: Configuring RADIUS Authentication on a QFabric System

RADIUS authentication is a method of authenticating users who are attempting to access a network device. On a QFabric system, users are load balanced on each of the Director devices. Each Director device needs to be able to communicate with the RADIUS server. Packets sent to the RADIUS server originate from the Director device IP addresses.

The following example shows how to configure RADIUS authentication on the QFabric system:

Perform the following steps to configure RADIUS authentication on the QFabric system:

1. Configure the order in which the authentication methods are used.

For example:

```
user@switch # set system authentication-order [radius password]
```

In this example, RADIUS authentication is the first authentication method that Junos OS will use when a user logs into the system.

2. Configure the IP address of the RADIUS server and the secret password. The secret password on the switch must match the secret password on the RADIUS server.

For example:

```
user@switch # set system radius-server 172.28.36.108 secret testing123
```

3. Assign the login class and the template account for the user.

For example:

```
user@switch # set system login user remote class super-user
```

Here are the results of your configuration:

```
[edit]
system {
  authentication-order [ radius password ];
  login {
    user remote {
      class super-ruser;
    }
  }
  radius-server {
    172.28.36.108 {
      secret test123
    }
  }
}
```

Related Documentation • [Configuring RADIUS Authentication on page 1699](#)

Example: Configuring RADIUS System Accounting

The following example shows three servers (10.5.5.5, 10.6.6.6, and 10.7.7.7) configured for RADIUS accounting.

```
system {
  accounting {
    events [ login change-log interactive-commands ];
    destination {
      radius {
        server {
          10.5.5.5 {
            accounting-port 3333;
            secret $9$dkafeqwrew;
            source-address 10.1.1.1;
            retry 3;
            timeout 3;
          }
          10.6.6.6 secret $9$fe3erqwrez;
          10.7.7.7 secret $9$f34929ftby;
        }
      }
    }
  }
}
```

Related Documentation • [Configuring RADIUS System Accounting on page 1697](#)

Example: Configuring the Root Password

The following example shows how to configure the root password:

```
[edit]
user@switch# set system root-authentication encrypted-password
"$1$14c5.$sBopasddsdfs0"
[edit]
user@switch# show
system {
  root-authentication {
    encrypted-password "$1$14c5.$sBopasddsdfs0";
  }
}
```

Related Documentation

- [Configuring the Root Password on page 158](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 188](#)
- [Configuring the Root Password on page 1702](#)

Example: Configuring SSH Authentication for Root Logins

The following example shows how to configure two public DSA keys for SSH authentication of root logins:

```
[edit system]
root-authentication {
  encrypted-password "$1$1wp5tqMX$uy/u5H7OdXTwfWTmeJWXe/";
  ## SECRET-DATA;
  ssh-dsa "2354 95 9304@boojum.per";
  ssh-dsa "0483 02 8362@ecbatana.per";
}
```

Related Documentation

- [Configuring the Root Password on page 158](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1687](#)

Example: Configuring User Accounts

The following example shows how to create accounts for four router or switch users, and create an account for the template user **remote**. All users use one of the default system login classes. User **alexander** also has two digital signal algorithm (DSA) public keys configured for SSH authentication.

```
[edit]
system {
  login {
    user philip {
      full-name "Philip of Macedonia";
      uid 1001;
      class super-user;
      authentication {
        encrypted-password "$1$poPPeY";
      }
    }
  }
}
```

```
}
user alexander {
  full-name "Alexander the Great";
  uid 1002;
  class view;
  authentication {
    encrypted-password "$1$14c5.$sBopasdFFdssdfFFdsdfs0";
    ssh-dsa "8924 37 5678 5678@gaugamela.per";
    ssh-dsa "6273 94 9283@boojum.per";
  }
}
user darius {
  full-name "Darius King of Persia";
  uid 1003;
  class operator;
  authentication {
    ssh-rsa "1024 37 12341234@ecbatana.per";
  }
}
user anonymous {
  class unauthorized;
}
user remote {
  full-name "All remote users";
  uid 9999;
  class read-only;
}
}
```

- Related Documentation**
- *Junos OS User Accounts Overview*
 - *Limiting the Number of User Login Attempts for SSH and Telnet Sessions*

Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication

The following example shows how to configure system authentication for RADIUS, TACACS+, and password authentication.

In this example, only the user Philip and users authenticated by a remote RADIUS server can log in. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router or switch. If the RADIUS server is not available, the user is authenticated using the **password** authentication method and allowed access to the router or switch. For more information about the password authentication method, see ["Using Local Password Authentication" on page 1677](#).

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the **super-user** class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the privileges associated with the **operator** class.

[edit]

```

system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class operator;
    }
  }
}

```



NOTE: For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see [“Overview of Template Accounts for RADIUS and TACACS+ Authentication” on page 1665](#).

When a user logs in to a device, the user's login name is used by the RADIUS or TACACS+ server for authentication. If the user is authenticated successfully by the authentication server and the user is not configured at the **[edit system login user]** hierarchy level, the device uses the default remote template user account for the user, provided a remote template account is configured at the **edit system login user remote** hierarchy level. The remote template account serves as a default template user account for all users that are authenticated by the authentication server but not having a locally configured user account on the device. Such users share the same login class and UID.

To configure an alternate template user, specify the **user-name** parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample Junos OS configuration:

```

[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user operator {
      full-name "All operators";
      uid 9990;
      class operator;
    }
    user remote {
      full-name "All remote users";
    }
  }
}

```

```
        uid 9999;
        class read-only;
    }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (**super-user**) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

Related Documentation • [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1694](#)

Example: Creating Login Classes with Specific Privileges

The following example shows how to create several user classes, each with specific privileges. In this example, you configure timeouts to disconnect the class members after a period of inactivity. Users' privilege levels, and therefore the classes of which they are members, should be dependent on their responsibilities within the organization, and the permissions shown here are only examples.

The first class of users (called “observation”) can only view statistics and configuration. They are not allowed to modify any configuration. The second class of users (called “operation”) can view and modify the configuration. The third class of users (called “engineering”) has unlimited access and control.

```
[edit]
system {
  login {
    class observation {
      idle-timeout 5;
      permissions [ view ];
    }
    class operation {
      idle-timeout 5;
      permissions [ admin clear configure interface interface-control network
        reset routing routing-control snmp snmp-control trace-control
        firewall-control rollback ];
    }
    class engineering {
      idle-timeout 5;
      permissions all;
    }
  }
}
```

}

Related Documentation

- [Defining Junos OS Login Classes](#)

Example: Configuring QFabric System Login Classes

This example shows you how to assign the correct login class to users so they can access components within a QFabric system.

- [Requirements on page 1737](#)
- [Overview on page 1737](#)
- [Configuration on page 1738](#)
- [Verification on page 1740](#)

Requirements

This example uses the following hardware and software components:

- One QFX3000-G QFabric system containing:
 - Two QFX3100 Director devices
 - Two QFX3008-I Interconnect devices
 - Eight QFX3500 Node devices
 - Junos OS Release 12.2 for these QFX Series components
- Eight EX4200 switches, used to make two redundant Virtual Chassis with four members apiece
- Junos OS Release 12.1R1.9 for the EX Series switches used in the Virtual Chassis

Before you begin:

- Perform the initial setup of the QFabric system on the Director group, which includes the creation of a username and password for the QFabric system components. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#).

Overview

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system (such as Node devices and Interconnect devices). The *qfabric-admin* class provides the ability to log in to individual QFabric system components and manage them. The *qfabric-operator* class enables the user to log in to individual components and view component-level operations and configurations. The *qfabric-user* class prevents access to individual QFabric system components.

You include these classes in your configuration at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level. The key task is to decide which class you should apply to users based on their need to access QFabric system components.



NOTE: To set QFabric system login classes for a root user, include the `remote-debug-permission` statement at the `[edit system root-authentication]` hierarchy level and specify the `qfabric-admin` class.

If you assign the `qfabric-admin` or the `qfabric-operator` class to a user, the QFabric system maps the user to a list of authorized users who are permitted to access components. To facilitate ease of use, the QFabric system uses the component password you specified during the initial setup of the Director group. When users assigned the `qfabric-admin` or the `qfabric-operator` class log in to a component by issuing the **request component login** operational mode command, the QFabric system verifies the class and sends the username and password to the component. The component accepts these credentials and permits access.



NOTE:

- The three QFabric system login classes give access to the components only. To provide access to the QFabric system as a whole through the default partition command-line interface (CLI), you must configure the usual Junos OS login classes or permissions (such as the *super-user* class). For more information about login classes, see [“Junos OS Login Classes Overview” on page 1683](#).
- If you have completed the QFabric system initial setup and the system is operational, you can change the component password by issuing the `device-authentication` statement at the `[edit system]` hierarchy level in the QFabric default partition CLI.

Topology

This example defines three users: Adam, Oscar, and Ulf. Adam needs to manage QFabric system components, Oscar needs limited access, and Ulf should not have any access to the components. As a result, assign the `qfabric-admin` class to Adam, the `qfabric-operator` class to Oscar, and the `qfabric-user` class to Ulf. However, all three users should have all permissions to access the QFabric system CLI.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login class all-qfabric permissions all
set system login user Adam class all-qfabric
set system login user Adam authentication encrypted-password
"$1$aoYSFkve$G/dYqsTV5iSvVW2sND69U."
set system login user Adam authentication remote-debug-permission qfabric-admin
set system login user Oscar class all-qfabric
set system login user Oscar authentication encrypted-password
"$1$3e.3wJQ8$3ISrzV0.efdRbk.ZJncKm0"
```

```

set system login user Oscar authentication remote-debug-permission qfabric-operator
set system login user Ulf class all-qfabric
set system login user Ulf authentication encrypted-password
"$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"
set system login user Ulf authentication remote-debug-permission qfabric-user

```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To provide the same access to the QFabric system CLI for all users, but different QFabric system component-level access to different users:

1. Define and provide all-qfabric access and passwords to all three users. This administrator-defined class provides full permissions, enabling the users to log in to the QFabric system default partition and use the CLI. Alternatively, you can assign the super-user class to these users to accomplish the same goal.

```

[edit]
user@qfabric# set system login class all-qfabric permissions all
user@qfabric# set system login user Adam class all-qfabric
user@qfabric# set system login user Adam authentication encrypted-password
"$1$aoYSFkvESG/dYqsTV5iSvVW2sND69U."
user@qfabric# set system login user Oscar class all-qfabric
user@qfabric# set system login user Oscar authentication encrypted-password
"$1$3e.3wJQ8$31SrZV0.efdRbk.ZJncKm0"
user@qfabric# set system login user Ulf class all-qfabric
user@qfabric# set system login user Ulf authentication encrypted-password
"$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"

```

2. Provide qfabric-admin component access to Adam so he can manage QFabric system components.

```

[edit]
user@qfabric# set system login user Adam authentication remote-debug-permission
qfabric-admin

```

3. Provide qfabric-operator component access to Oscar so he can view the CLI at the QFabric system components.

```

[edit]
user@qfabric# set system login user Oscar authentication remote-debug-permission
qfabric-operator

```

4. Assign qfabric-user component restrictions to Ulf to prevent him from accessing the QFabric system components.

```

[edit]
user@qfabric# set system login user Ulf authentication remote-debug-permission
qfabric-user

```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example.

```
[edit]
system {
  login {
    class all-qfabric {
      permissions all;
    }
    user Adam {
      class all-qfabric;
      authentication {
        encrypted-password "$1$aoYSFkvE$G/dYqsTV5iSvVW2sND69U."; ##
        SECRET-DATA
        remote-debug-permission qfabric-admin;
      }
    }
    user Oscar {
      class all-qfabric;
      authentication {
        encrypted-password "$1$3e.3wJQ8$31SrZV0.efdBk.ZJncKm0"; ## SECRET-DATA
        remote-debug-permission qfabric-operator;
      }
    }
    user Ulf {
      class all-qfabric;
      authentication {
        encrypted-password "$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"; ##
        SECRET-DATA
        remote-debug-permission qfabric-user;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the QFabric system and component-level access configuration is working properly for all three users. Adam, Oscar, and Ulf should have equivalent, full-permission access to the QFabric system CLI. Adam should have management-level access to components. Oscar should have read-only access to components. Ulf should have no component-level access.

- [Verifying qfabric-admin Access on page 1740](#)
- [Verifying qfabric-operator Access on page 1742](#)
- [Verifying qfabric-user Access on page 1743](#)

Verifying qfabric-admin Access

Purpose Verify that Adam can access the QFabric system CLI at the default partition and manage QFabric system components.

Action From a management station on your network, issue the `ssh user@qfabric` command and enter the password to open an SSH session for Adam to the QFabric system. Issue the `?` command to view the CLI operational mode commands that Adam has permission to use on the QFabric system default partition.

```
> ssh Adam@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Adam@qfabric.network.net's password:
Last login: Sun Nov 20 14:12:29 2011 from 192.168.28.19
Juniper QFabric Director 11.3.5510 2011-10-21 16:31:44 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg0
Adam@qfabric>
```

```
Adam@qfabric> ?
Possible completions:
clear          Clear information in the system
configure      Manipulate software configuration information
file           Perform file operations
help           Provide help information
load           Load information from file
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
telnet         Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host
```

Issue the `request component login ?` command to view the components that Adam can access. Next, issue the `request component login component-name` command to log in to a Node device without being prompted for a username or password.

```
Adam@qfabric> request component login ?
Possible completions:
<[Enter]>      Execute this command
<node-name>    Inventory name for the remote node
BBAK0372       Node device
BBAK0394       Node device
DRE-0          Diagnostic routing engine
EE3093         Node device
FC-0           Fabric control
FC-1           Fabric control
FM-0           Fabric manager
NW-NG-0        Node group
WS001/RE0      Interconnect device control board
WS001/RE1      Interconnect device control board
|              Pipe through a command
```

```
Adam@qfabric> request component login EE3093
Warning: Permanently added 'qnode-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the `?` command to view the CLI operational mode commands that Adam has the permission to use on the Node device. Notice that the CLI prompt now indicates Adam's component access level (**qfabric-admin**) as the username and the Node device identifier (**EE3093**) as the host.

```
qfabric-admin@EE3093> ?
Possible completions:
clear          Clear information in the system
file           Perform file operations
help           Provide help information
load           Load information from file
monitor        Show real-time debugging information
mtrace         Trace multicast path from source to receiver
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
ssh            Start secure shell on another host
start          Start shell
telnet         Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host
```

Meaning The output shows that Adam has received the proper permissions to access the QFabric system CLI and log in to individual components with management-level access.

Verifying qfabric-operator Access

Purpose Verify that Oscar can access the QFabric system CLI at the default partition and view the CLI on the QFabric system components.

Action From a management station on your network, issue the `ssh user@qfabric` command and enter the password to open an SSH session for Oscar to the QFabric system. Issue the `?` command to view the CLI operational mode commands that Oscar has permission to use on the QFabric system default partition. Notice that these permissions are the same as those given to Adam.

```
> ssh Oscar@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Oscar@qfabric.network.net's password:
Last login: Sun Nov 19 19:21:29 2011 from 192.168.28.14
Juniper QFabric Director 11.3.5510 2011-10-22 18:33:41 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg1
Oscar@qfabric>
```

```
Oscar@qfabric> ?
Possible completions:
clear          Clear information in the system
configure      Manipulate software configuration information
file           Perform file operations
help           Provide help information
load           Load information from file
```

op	Invoke an operation script
ping	Ping remote target
quit	Exit the management session
request	Make system-level requests
restart	Restart software process
save	Save information to file
set	Set CLI properties, date/time, craft interface message
show	Show system information
telnet	Telnet to another host
test	Perform diagnostic debugging
tracert	Trace route to remote host

Issue the **request component login *component-name*** command to log in to a Node device without being prompted for a username or password.

```
Oscar@qfabric> request component login EE3093
Warning: Permanently added 'qfnode-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the **?** command to view the CLI operational mode commands that Oscar has permission to use on the Node device. Notice that the CLI prompt now indicates Oscar's component access level (**qfabric-operator**) as the username and the Node device identifier (**EE3093**) as the host. Additionally, Oscar has fewer CLI commands available than Adam because of Oscar's read-only qfabric-operator login class.

```
qfabric-operator@EE3093> ?
Possible completions:
file          Perform file operations
help          Provide help information
load          Load information from file
op            Invoke an operation script
quit          Exit the management session
request       Make system-level requests
save          Save information to file
set           Set CLI properties, date/time, craft interface message
show          Show system information
start         Start shell
test          Perform diagnostic debugging
```

Meaning The output shows that Oscar has full permissions to access the QFabric system CLI, but only read-only access when he logs in to individual components. Oscar's permissions on the QFabric system are the same as Adam's, but Oscar has fewer permissions than Adam on the Node device.

Verifying qfabric-user Access

Purpose Verify that Ulf has full access to the QFabric system CLI at the default partition but cannot access the QFabric system components.

Action From a management station on your network, issue the **ssh *user*@qfabric** command and enter the password to open an SSH session for Ulf to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Ulf has permission to use

on the QFabric system default partition. Notice that these permissions are the same as those given to Adam and Oscar.

```
> ssh Ulf@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Ulf@qfabric.network.net's password:
Last login: Sun Nov 17 17:12:24 2011 from 192.168.28.22
Juniper QFabric Director 11.3.5510 2011-10-23 19:23:31 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg0
Ulf@qfabric>
```

```
Ulf@qfabric> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  help           Provide help information
  load           Load information from file
  op             Invoke an operation script
  ping           Ping remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart software process
  save           Save information to file
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  telnet         Telnet to another host
  test           Perform diagnostic debugging
  traceroute     Trace route to remote host
```

When Ulf issues the **request component login *component-name*** command, the Node device denies his access attempt.

```
Ulf@qfabric> request component login EE3093
error: User Ulf does not have sufficient permissions to login to device EE3093
```

Meaning The output shows that Ulf has full permissions to access the QFabric system CLI in the same way as Adam and Oscar. However, unlike Adam and Oscar, Ulf cannot access individual components because of the qfabric-user login class assigned to him.

- Related Documentation**
- [Understanding QFabric System Login Classes on page 1230](#)
 - [remote-debug-permission on page 1401](#)
 - [request component login on page 1444](#)
 - [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
 - [Junos OS Login Classes Overview on page 1683](#)

Example: Configuring User Login Accounts

The following example shows how to configure the local administrator account (**user admin**). If RADIUS fails or becomes unreachable, the login process reverts to password authentication on the local accounts on the router or switch.

```
[edit]
system {
  login {
    user admin {
      uid 1000;
      class engineering;
      authentication {
        encrypted-password "<PASSWORD>"; # SECRET-DATA
      }
    }
  }
}
```

Related Documentation

- [Configuring Junos OS User Accounts](#)

Example: Configuring RADIUS Template Accounts

The following example shows how to configure RADIUS template accounts for different users or groups of users:

```
[edit]
system {
  login {
    user observation {
      uid 1001;
      class observation;
    }
    user operation {
      uid 1002;
      class operation;
    }
    user engineering {
      uid 1003;
      class engineering;
    }
  }
}
```

Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1665](#)

Defining Access Privileges Using allow/deny-configuration Statements

The following examples show how to configure access privileges for individual configuration mode hierarchy levels.

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot configure telnet parameters:

```
[edit system login class class-name]
user@switch# set deny-configuration "system services telnet"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]
user@switch# set deny-configuration "system login class m.*"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot edit a configuration or issue commands (such as **commit**) at the login class or system services hierarchy levels:

```
[edit system login class class-name]
user@switch# set deny-configuration "(system login class) | (system services)"
```

The following example shows how to configure permissions for individual configuration mode hierarchies:

```
[edit]
system {
  login { # This login class has operator privileges and the additional ability to edit
    # configuration at the system services hierarchy level.
    class only-system-services {
      permissions [ configure ];
      allow-configuration "system services";
    }
    # services commands.
    class all-except-system-services { # This login class has operator privileges but
      # cannot edit any system services configuration.
      permissions [ all ];
      deny-configuration "system services";
    }
  }
}
```

**Related
Documentation**

- [Specifying Access Privileges Using allow/deny-configuration Statements](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1719](#)

Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions

The following example shows how to limit the user to four attempts when the user enters a password while logging in through SSH or Telnet. Set the **backoff-threshold** to 2, the **back-off-factor** to 5 seconds, and the **minimum-time** to 40 seconds. The user experiences a delay of 5 seconds after the second attempt to enter a correct password fails. After each subsequent failed attempt, the delay increases by 5 seconds. After the fourth and final failed attempt to enter a correct password, the user experiences an additional 10-second delay, and the connection closes after a total of 40 seconds.

The additional variables **maximum-time** and **lockout-period** are not set in this example.

```
[edit]
system {
  login {
    retry-options {
      backoff-threshold 2;
      backoff-factor 5;
      minimum-time 40;
      tries-before-disconnect 4;
    }
  }
}
```

```

    password {
    }
  }
}

```



NOTE: This sample only shows the portion off the [edit system login] hierarchy level being modified.

Related Documentation

- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [login](#)
- [login on page 264](#)

Configuration Statements

- [access on page 1749](#)
- [accounting on page 1750](#)
- [accounting-options on page 1751](#)
- [accounting-server on page 1753](#)
- [accounting-stop-on-access-deny on page 1754](#)
- [accounting-stop-on-failure on page 1755](#)
- [advertisement-interval on page 1756](#)
- [agent-address on page 1757](#)
- [archival on page 1758](#)
- [archive-sites \(Configuration File\) on page 1759](#)
- [authentication-order on page 1760](#)
- [authentication-server on page 1761](#)
- [authorization on page 1762](#)
- [categories on page 1763](#)
- [client-list on page 1763](#)
- [client-list-name on page 1764](#)
- [clients on page 1764](#)
- [commit-delay on page 1765](#)
- [community \(SNMP\) on page 1766](#)
- [configuration on page 1767](#)
- [connection-limit on page 1768](#)
- [contact on page 1769](#)
- [disable \(LLDP\) on page 1769](#)
- [ethernet-switching-options on page 1770](#)
- [falling-threshold \(Health Monitor\) on page 1772](#)

- [filter-duplicates](#) on page 1772
- [full-name](#) on page 1773
- [health-monitor](#) on page 1773
- [hold-multiplier](#) on page 1774
- [idle-timeout \(Access\)](#) on page 1775
- [interface \(LLDP\)](#) on page 1776
- [interval \(Health Monitor\)](#) on page 1777
- [lldp](#) on page 1778
- [lldp-configuration-notification-interval](#) on page 1779
- [location](#) on page 1779
- [management-address](#) on page 1780
- [name](#) on page 1780
- [nas-ip-address](#) on page 1781
- [nonvolatile](#) on page 1781
- [oid](#) on page 1782
- [order](#) on page 1783
- [port \(RADIUS Server\)](#) on page 1784
- [profile](#) on page 1785
- [protocols](#) on page 1786
- [protocol-version](#) on page 1799
- [ptopo-configuration-maximum-hold-time](#) on page 1799
- [ptopo-configuration-trap-interval](#) on page 1800
- [radius](#) on page 1801
- [radius-options \(edit system\)](#) on page 1802
- [radius-server](#) on page 1803
- [rate-limit](#) on page 1804
- [remote-debug-permission](#) on page 1805
- [retry](#) on page 1806
- [rising-threshold \(Health Monitor\)](#) on page 1807
- [root-login](#) on page 1808
- [services \(Switches\)](#) on page 1809
- [snmp](#) on page 1810
- [ssh](#) on page 1814
- [system](#) on page 1815
- [tacplus-options](#) on page 1821
- [targets](#) on page 1822
- [traceoptions \(LLDP\)](#) on page 1823

- [transfer-interval \(Configuration\)](#) on page 1825
- [transfer-on-commit](#) on page 1826
- [trap-group](#) on page 1827
- [trap-options](#) on page 1828
- [user \(Access\)](#) on page 1829
- [version](#) on page 1830

access

Syntax

```
access {
  address-assignment
  pool pool-name
  address-pool pool-name
  profile profile-name {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      (authentication-order (ldap radius | none);
      order (radius | none);
    }
    radius {
      accounting-server [server-addresses];
      authentication-server [server-addresses];
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure authentication, authorization, and accounting (AAA) services.

The statements are explained separately.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Default Not enabled

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\)](#)

accounting

Syntax	<pre>accounting { accounting-stop-on-access-deny; accounting-stop-on-failure; order (radius none); }</pre>
Hierarchy Level	[edit access profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the authentication order for authentication, authorization, and accounting (AAA) services.
Default	Not enabled
Options	none —Use no authentication for specified subscribers. radius —Use RADIUS authentication for specified subscribers. The remaining statements are explained separately.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i>• <i>Configuring RADIUS Accounting</i>• Understanding RADIUS Accounting on page 1667

accounting-options

```

Syntax  accounting-options {
            class-usage-profile profile-name {
                destination-classes {
                    destination-class-name;
                }
                file filename;
                interval minutes;
                source-classes {
                    source-class-name;
                }
            }
            file filename {
                archive-sites {
                    site-name;
                }
                files number;
                nonpersistent;
                size bytes;
                start-time time;
                transfer-interval minutes;
            }
            filter-profile profile-name {
                counters {
                    counter-name;
                }
                file filename;
                interval minutes;
            }
            interface-profile profile-name {
                fields {
                    input-bytes;
                    input-errors;
                    input-multicast;
                    input-packets;
                    input-unicast;
                    output-bytes;
                    output-errors;
                    output-multicast;
                    output-packets;
                    output-unicast;
                    rpf-check-bytes;
                    rpf-check-packets;
                    rpf-check6-bytes;
                    rpf-check6-packets;
                    unsupported-protocol;
                }
                file filename;
                interval minutes;
            }
            mib-profile profile-name {
                file filename;
                interval minutes;
            }
        }

```

```
object-names {  
    mib-object-name;  
}  
operation (get | get-next | walk);  
}  
policy-decision-statistics-profile profile-name {  
    application-aware-access-list-fields {  
        address;  
        application;  
        application-group;  
        input-bytes;  
        input-interface;  
        input-packets;  
        mask;  
        output-bytes;  
        output-packets;  
        subscriber-name;  
        timestamp;  
        vrf-name;  
    }  
    file filename;  
}  
routing-engine-profile profile-name {  
    fields {  
        field-name;  
    }  
    file filename;  
    interval minutes;  
}  
}
```

Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure options for accounting statistics collection.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.

accounting-server

Syntax	<code>accounting-server[server-addresses];</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
Default	Not enabled
Options	<i>server-addresses</i> —One or more addresses of RADIUS authentication servers.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>show network-access aaa statistics authentication</i> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i> • Understanding RADIUS Accounting on page 1667

accounting-stop-on-access-deny


Syntax	accounting-stop-on-access-deny;
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if the AAA server denies access to a supplicant.



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Default	Not enabled
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i>• <i>show network-access aaa statistics authentication</i>• <i>Configuring RADIUS Accounting</i>

accounting-stop-on-failure

Syntax	accounting-stop-on-failure;
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if a supplicant fails AAA authorization, but the RADIUS server grants access. For example, a supplicant might fail AAA authentication because of an internal error such as a timeout.
<div>  NOTE: The [edit access] hierarchy is not available on QFabric systems. </div>	
Default	Not enabled
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i> • <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i> • <i>Configuring RADIUS Accounting</i> • Understanding RADIUS Accounting on page 1667

advertisement-interval

Syntax	<code>advertisement-interval seconds;</code>
Hierarchy Level	[edit protocols lldp], [edit routing-instances <i>routing-instance-name</i> protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>For MX Series and T Series routers, configure an interval for LLDP advertisement.</p> <p>For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.</p> <p>The advertisement-interval value must be greater than or equal to four times the transmit-delay value, or an error will be returned when you attempt to commit the configuration.</p>





NOTE: The default value of **transmit-delay** is 2 seconds. If you configure the **advertisement-interval** as less than 8 seconds and you do not configure a value for **transmit-delay**, the default value of **transmit-delay** is automatically changed to 1 second in order to satisfy the requirement that the **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value.

Default	Disabled.
Options	seconds —Interval between LLDP advertisement. Default: 30 Range: 5 through 32768
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LLDP• show lldp on page 1842• Configuring LLDP (CLI Procedure)• Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches• transmit-delay• Understanding LLDP on page 1666



agent-address

Syntax	agent-address outgoing-interface;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is outgoing-interface , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
Options	outgoing-interface —Value of the agent address of all SNMPv1 traps generated by this router or switch. The outgoing-interface option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. Default: Disabled (the agent address is not specified in SNMPv1 traps).
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Agent Address for SNMP Traps</i>

archival

Syntax	<pre> archival { configuration { archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } transfer-interval interval; transfer-on-commit; } } </pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.
<div>  <p>NOTE: The <code>edit system archival</code> hierarchy is not available on QFabric systems.</p> </div>	
Options	The remaining statements are explained separately.
<div>  <p>NOTE: The <code>[edit system archival]</code> hierarchy is not available on QFabric systems.</p> </div>	
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611

archive-sites (Configuration File)

Syntax	<pre>archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; }</pre>
Hierarchy Level	[edit system archival configuration]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([]). For example, "scp://username<:password>@[ipv6-host-address]<:port>/url-path"</p> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <p><i>router-name_juniper.conf.n.gz_YYYYMMDD_HHMMSS.</i></p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
Options	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p>file:// —transfer on a path to a named file</p> <p>ftp:// —transfer using active FTP server</p> <p>pasvftp:// —transfer to a device that only accepts passive FTP services</p>

scp:// —transfer to a known host using background SCP file transfers

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Archive Sites for Transfer of Active Configuration Files on page 1612](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 50](#)
- [configuration on page 1617](#)
- [transfer-on-commit on page 1619](#)

authentication-order

Syntax authentication-order [none | password | radius];

Hierarchy Level [edit [access profile](#) *profile-name*],
 [edit [system](#)]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending authentication messages.

Default Not enabled

Options **none**—No authentication for specified subscribers.

password—Password authentication.

radius—RADIUS authentication.




NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

authentication-server

Syntax	<code>authentication-server [server-addresses];</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the RADIUS server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
Options	server-addresses —Configure one or more RADIUS server addresses.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i>• <i>show network-access aaa statistics authentication</i>

authorization

Syntax	<code>authorization <i>authorization</i>;</code>
Hierarchy Level	[edit snmp community <i>community-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the access authorization for SNMP Get , GetBulk , GetNext , and Set requests.
Options	<p><i>authorization</i>—Access authorization level:</p> <ul style="list-style-type: none">• read-only—Enable Get, GetNext, and GetBulk requests.• read-write—Enable all requests, including Set requests. You must configure a view to enable Set requests.
	<div> NOTE: The read-write option is not supported on the QFX3000 QFabric system.</div>
	Default: read-only
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMP Community String on page 6194

categories

Syntax	<code>categories { category; }</code>
Hierarchy Level	<code>[edit snmp trap-group group-name]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the types of traps that are sent to the targets of the named trap group.
Default	If you omit the categories statement, all trap types are included in trap notifications.
Options	category —Name of a trap type: authentication , chassis , configuration , link , remote-operations , rmon-alarm , or startup .
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 6195

client-list

Syntax	<code>client-list client-list-name { ip-addresses; }</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define a list of SNMP clients.
Options	client-list-name —Name of the client list. ip-addresses —IP addresses of the SNMP clients to be added to the client list,
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Adding a Group of Clients to an SNMP Community on page 6196

client-list-name

Syntax	<code>client-list-name <i>client-list-name</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a client list or prefix list to an SNMP community.
Options	<i>client-list-name</i> —Name of the client list or prefix list.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Group of Clients to an SNMP Community on page 6196

clients

Syntax	<pre>clients { <i>address</i> <restrict>; }</pre>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
Default	If you omit the clients statement, all SNMP clients using this community string are authorized to access the switch.
Options	<i>address</i> —Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <i>address</i> options. <i>restrict</i> —(Optional) Do not allow the specified SNMP client to access the switch.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMP Community String

commit-delay

Syntax	commit-delay <i>seconds</i> ;
Hierarchy Level	[edit snmp nonvolatile]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the timer for the SNMP Set reply and start of the commit.
Options	seconds —Delay between an affirmative SNMP Set reply and start of the commit operation. Default: 5 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Commit Delay Timer</i>

community (SNMP)

Syntax `community community-name {
 authorization authorization;
 client-list-name client-list-name;
 clients {
 address restrict;
 }
 view view-name;
 }`

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.



NOTE: The **authorization read-write** option is not supported on the QFX3000 QFabric system.

The SNMP client application specifies an SNMP community name in **Get**, **GetBulk**, **GetNext**, and **Set** SNMP requests.

Default If you omit the **community** statement, all SNMP requests are denied.


Options **community-name**—Community string. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
 snmp-control—To add this statement to the configuration.

Related Documentation • [Configuring the SNMP Community String on page 6194](#)

configuration

Syntax	<pre>configuration { transfer-interval interval; transfer-on-commit; archive-sites { file://<path>/<filename>; ftp://username@host:<port>url-path password password; http://username@host:<port>url-path password password; pasvftp://username@host:<port>url-path password password; scp://username@host:<port>url-path password password; } }</pre>
Hierarchy Level	[edit system archival]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the router or switch to periodically transfer its currently active configuration (or after each commit).
<div>  <p>NOTE: The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>	
Options	The remaining statements are explained separately.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1611 • archive on page 6366 • archive-sites on page 1615 • transfer-interval on page 1618 • transfer-on-commit on page 1619

connection-limit

Syntax	<code>connection-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
Options	<p>limit—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p>Range: 1 through 250</p> <p>Default: 75</p>



NOTE: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring clear-text or SSL Service for Junos XML Protocol Client Applications • Configuring DTCP-over-SSH Service for the Flow-Tap Application • Configuring Finger Service for Remote Access to the Router • Configuring FTP Service for Remote Access to the Router or Switch • Configuring SSH Service for Remote Access to the Router or Switch on page 1709 • Configuring Telnet Service for Remote Access to a Router or Switch

contact

Syntax	<code>contact <i>contact</i>;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysContact object, which is the contact person for the managed system.
Options	contact —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the System Contact on a Device Running Junos OS</i>

disable (LLDP)

Syntax	<code>disable;</code>
Hierarchy Level	<code>[edit protocols lldp],</code> <code>[edit protocols interface lldp]</code>
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable the LLDP configuration on the switch or on one or more interfaces.
Default	If you do not configure LLDP, it is disabled on the switch and on specific switch interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 1842 • <i>Configuring LLDP (CLI Procedure)</i> • <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i> • Configuring LLDP on page 1693 • Understanding LLDP on page 1666

ethernet-switching-options

```
Syntax ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
      output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds {
  }
  port-error-disable {
    disable-timeout timeout;
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      fcoe-trusted;
      mac-limit limit action action;
      no-allowed-mac-log;
    }
    vlan (all | vlan-name) {
      (arp-inspection | no-arp-inspection) [
        forwarding-class (for DHCP Snooping or DAI Packets) class-name;
      ]
    }
  }
}
```

```

dhcp-option82 {
  circuit-id {
    prefix (Circuit ID for Option 82) hostname;
    use-interface-description;
    use-vlan-id;
  }
  remote-id {
    prefix (Remote ID for Option 82) hostname | mac | none;
    use-interface-description;
    use-string string;
  }
  vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp) {
  forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
  examine-vn2vn {
    beacon-period milliseconds;
  }
  fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
 - [Overview of Access Port Protection on page 4674](#)
 - [Understanding Storm Control on page 4694](#)

falling-threshold (Health Monitor)

Syntax	<code>falling-threshold <i>percentage</i>;</code>
Hierarchy Level	<code>[edit snmp health-monitor]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —Lower threshold for the alarm entry. Range: 1 through 100 Default: 70 percent of the maximum possible value
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• rising-threshold on page 1807• Configuring Health Monitoring on page 6200

filter-duplicates

Syntax	<code>filter-duplicates;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Filter duplicate Get , GetNext , or GetBulk SNMP requests.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.

full-name

Syntax	<code>full-name <i>complete-name</i>;</code>
Hierarchy Level	[edit system login user]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the complete name of a user.
Options	<i>complete-name</i> —Full name of the user. If the name contains spaces, enclose it in quotation marks.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS User Accounts</i> • <i>user</i> • user on page 314


health-monitor

Syntax	<pre>health-monitor { falling-threshold <i>percentage</i>; interval <i>seconds</i>; rising-threshold <i>percentage</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure health monitoring. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Health Monitoring on page 6200 • Understanding Health Monitoring on page 6123

hold-multiplier

Syntax	<code>hold-multiplier <i>number</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series.
Description	Specify the multiplier used in combination with the advertisement-interval value to determine the length of time LLDP information is held before it is discarded. The default value is 4 (or 120 seconds).
Default	Disabled.
Options	<i>number</i> —A number used as a multiplier. Range: 2 through 10 Default: 4 (or 120 seconds)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show lldp on page 1842• <i>Configuring LLDP (CLI Procedure)</i>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i>• Configuring LLDP on page 1693• Understanding LLDP on page 1666

idle-timeout (Access)

Syntax	<code>idle-timeout seconds;</code>
Hierarchy Level	<code>[edit access group-profile <i>profile-name</i> ppp],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> ppp]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:</p> <ul style="list-style-type: none"> • There is no ingress traffic on the PPP session. • There is no egress traffic. • There is neither ingress or egress traffic on the PPP session. • There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.
Options	<p>seconds—Number of seconds a user can remain idle before the session is terminated.</p> <p>Range: 0 through 4,294,967,295 seconds</p> <p>Default: 0</p>
<div>  NOTE: The [edit access] hierarchy is not available on QFabric systems. </div>	
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Group Profile for Defining L2TP Attributes</i> • <i>Configuring PPP Properties for a Client-Specific Profile</i> • <i>Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile</i>

interface (LLDP)

Syntax	<pre>interface (all <i>interface-name</i>) { disable; power-negotiation { disable; } }</pre>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a specific interface.
Default	None
Options	<p>all—All interfaces on the switch.</p> <p><i>interface-name</i>—Name of a specific interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring LLDP (CLI Procedure)</i>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i>• Configuring LLDP on page 1693• Understanding LLDP on page 1666

interval (Health Monitor)

Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the interval between sampling of the object being monitored by the health monitor.
Options	<p><i>seconds</i>—Time between samples, in seconds.</p> <p>Range: 1 through 2147483647 seconds</p> <p>Default: 300 seconds</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 6200

lldp

```
Syntax  lldp {
    advertisement-interval seconds;
    disable;
    hold-multiplier number;
    interface (all | [interface-name]) {
        disable;
        power-negotiation {
            disable;
        }
    }
    lldp-configuration-notification-interval seconds;
    management-address ip-management-address;
    netbios-snooping;
    ptopo-configuration-maximum-hold-time seconds;
    ptopo-configuration-trap-interval seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <no-stamp> <replace>;
        flag flag <disable>;
    }
    transmit-delay seconds;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for QFX Series.

Description Configure Link Layer Discovery Protocol (LLDP). The switch uses LLDP to advertise its identity and capabilities on a LAN, as well as receive information about other network devices. LLDP is defined in the IEEE standard 802.1AB-2005.

The remaining statements are explained separately.



NOTE: The `transmit-delay` and `netbios-snooping` options are not available on QFabric systems.

Default LLDP is enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [show lldp on page 1842](#)
- *Configuring LLDP (CLI Procedure)*
- *Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches*

- [Configuring LLDP on page 1693](#)
- [Understanding LLDP on page 1666](#)

lldp-configuration-notification-interval

Syntax	lldp-configuration-notification-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify how often SNMP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, trap notifications of database changes are disabled.
Default	SNMP trap notifications of LLDP database changes are disabled.
Options	<i>seconds</i> —Interval between trap notifications about LLDP database changes. Range: 0 through 3600
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 1842

location

Syntax	location <i>location</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysLocation object, which is the physical location of the managed system.
Options	<i>location</i> —Location of the local system. You must enclose the name within quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the System Location for a Device Running Junos OS

management-address

Syntax	<code>management-address <i>ip-management-address</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the management address of the switch to be used in the LLDP Management type, length, and value (TLV) .
Default	LLDP Management TLV uses the IP address of the switch's management Ethernet interface (me0) or the IP address of the virtual management Ethernet (VME) interface if the switch is a Virtual Chassis.
Options	<i>ip-management-address</i> —You can specify either an IPv4 or IPv6 management address for the switch.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show lldp on page 1842• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i>• <i>EX Series Switches Interfaces Overview</i>• Understanding LLDP on page 1666

name

Syntax	<code>name <i>name</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the system name from the command-line interface.
Options	<i>name</i> —System name override.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Name</i>

nas-ip-address

Syntax	<code>nas-ip-address <i>ip-address</i>;</code>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the NAS-IP address for outgoing RADIUS packets.
Options	ip-address —IP address of the network access server (NAS) that requests user authentication.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RADIUS Authentication</i> • Configuring RADIUS Authentication on page 1699

nonvolatile

Syntax	<code>nonvolatile { commit-delay <i>seconds</i>; }</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure options for SNMP Set requests. The statement is explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Commit Delay Timer</i> • <i>commit-delay</i>

oid

Syntax	<code>oid <i>object-identifier</i> (exclude include);</code>
Hierarchy Level	<code>[edit snmp view <i>view-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects.
Options	<p>exclude—Exclude the subtree of MIB objects represented by the specified OID.</p> <p>include—Include the subtree of MIB objects represented by the specified OID.</p> <p><i>object-identifier</i>—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring MIB Views on page 6197

order

Syntax	<code>order (radius [<i>accounting-order-data-list</i>]);</code>
Hierarchy Level	[edit access profile <i>profile-name</i> accounting]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending accounting messages and updates.
Default	No order specified
Options	<p>radius—RADIUS accounting for specified subscribers.</p> <p>[<i>accounting-order-data-list</i>]— Set of data listing the authentication order to be used, enclosed by brackets. This can be any combination of the authentication methods, up to and including a full list of the entire authentication order.</p>



NOTE: The [edit access] hierarchy is not available on QFabric systems.

Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i> • <i>Configuring RADIUS Accounting</i>

port (RADIUS Server)


Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	[edit system radius-server <i>address</i>], [edit system accounting destination radius server <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the port number on which to contact the RADIUS server.
Options	<i>number</i> —Port number on which to contact the RADIUS server. Default: 1812 (as specified in RFC 2865)



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RADIUS Authentication</i>• Configuring RADIUS Authentication on page 1699

profile

Syntax	<pre> profile <i>profile-name</i> { accounting { accounting-stop-on-access-deny; accounting-stop-on-failure; order (radius [<i>accounting-order-data-list</i>]); } authentication-order [<i>authentication-method</i>]; radius { accounting-server [<i>server-addresses</i>]; authentication-server [<i>server-addresses</i>]; } } </pre>
Hierarchy Level	[edit access]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure an access profile. The access profile contains the entire authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, including the authentication method and order, AAA server addresses, and AAA accounting.
Default	Not enabled
Options	<p><i>profile-name</i>—Profile name of up to 32 characters.</p> <p>The remaining statements are explained separately.</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>NOTE: The [edit access] hierarchy is not available on QFabric systems.</p> </div> </div>	
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i> • <i>Configuring RADIUS Accounting</i>

protocols

```
Syntax protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            hold-down-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family family-name {
            ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
        }
        graceful-restart {
            disable;
            restart-time seconds;
            stale-routes-time seconds;
        }
        group group-name {
            ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
        }
        hold-time seconds;
        import [ policy-names ];
        include-mp-next-hop;
        keep (all | none);
        local-address address;
```

```

local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tll-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}

```

```

    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (Liveness Detection) (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable {
        interface interface-name;
    }
    interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
            group ip-address;
        }
    }
}

```



```

    }
    robust-count number;
  }
}
isis {
  disable;
  export [ policy-names ];
  ignore-attached-bit;
  interface interface-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
    }
    version (1 | automatic);
  }
  checksum;
  csnp-interval (seconds | disable);
  disable;
  hello-padding (adaptive | loose | strict);
  level (1 | 2) {
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    metric metric;
    passive;
    priority number;
  }
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-unicast-topology;
  passive;
  point-to-point;
}
level (1 | 2) {
  disable;
  authentication-key key;
  authentication-type authentication;
  external-preference preference;

```

```
no-csnp-authentication;
no-hello-authentication;
no-psnp-authentication;
preference preference;
prefix-export-limit number;
wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
}
topologies {
    ipv4-multicast;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
```

```

forward-delay seconds;
hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number> <size size> <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}

```

```
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
}
(summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
```

```

authentication {
    md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
    simple-password key-string;
}
dead-interval seconds;
demand-circuit;
flood-reduction;
hello-interval seconds;
ipsec-sa sa-name;
no-neighbor-down-notification;
retransmit-interval seconds;
topology (name | default | ipv4-multicast) {
    disable;
    metric metric;
}
transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
}

```

```
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
}
```

```

bootstrap-import [ policy-names ];
bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
prefix-list prefix-list-addresses {
    source source-address {
        next-hop next-hop-address;
    }
    wildcard-source {
        next-hop next-hop-address;
    }
}

```

```

    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    tunnel-devices [ mt-fpc/pic/port ];
}
rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
    }
    export [ policy-names ];
    import [ policy-names ];
    metric-out metric;
    neighbor neighbor-name {
        any-sender;
        authentication-key password;
        authentication-type type;
        bfd-liveness-detection {
            ... same statements as at the [edit protocols rip group group-name
                bfd-liveness-detection] hierarchy level ...
        }
        (check-zero | no-check-zero);
        import [ policy-names ];
        message-size number;
        metric-in metric;
        receive (both | none | version-1 | version-2);
        route-timeout seconds;
        send (broadcast | multicast | none | version-1);
        update-interval seconds;
    }
    preference preference;
    route-timeout seconds;
    update-interval seconds;
}

```



```

holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number> <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
    }
}

```

```

    priority priority;
  }
  max-age seconds;
}
traceoptions {
  file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
  flag flag;
}
uplink-failure-detection {
  group group-name {
    link-to-monitor interface-name;
    link-to-disable interface-name;
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number > <size size > <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
}
}

```

Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure protocols. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- Related Documentation**
- [Junos OS Routing Protocols Configuration Guide](#)

protocol-version

Syntax	<code>protocol-version <i>version</i>;</code>
Hierarchy Level	[edit system services ssh]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the secure shell (SSH) protocol version.
Default	v2—SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
Options	<i>version</i> —SSH protocol version: v1, v2, or both.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the SSH Protocol Version on page 1710


ptopo-configuration-maximum-hold-time

Syntax	<code>ptopo-configuration-maximum-hold-time <i>seconds</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure how long to maintain the physical topology database entries. The physical topology identifies the devices on the network and their physical interconnections.
Options	<i>seconds</i> —Time to maintain physical topology database entries. Default: 300 Range: 1 through 2147483647
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show lldp on page 1842 • Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches • Understanding LLDP on page 1666


ptopo-configuration-trap-interval

Syntax	<code>ptopo-configuration-trap-interval <i>seconds</i>;</code>
Hierarchy Level	[edit protocols lldp]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify how often SNMP trap notifications are sent regarding changes in physical topology global statistics.
Default	SNMP trap notifications of changes in physical topology global statistics are disabled.
Options	<i>seconds</i> —Interval between SNMP trap notifications about physical topology global statistics. Range: 0 through 3600
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.


radius

Syntax	radius { accounting-server [server-addresses]; authentication-server [server-addresses]; }
Hierarchy Level	[edit access profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the RADIUS servers for authentication and for accounting. To configure multiple RADIUS servers, include multiple radius statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p> <p>The statements are explained separately.</p>
<div>  NOTE: The [edit access] hierarchy is not available on QFabric systems. </div>	
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i> • <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i> • <i>Filtering 802.1X Supplicants Using RADIUS Server Attributes</i> • <i>Configuring RADIUS Accounting</i>

radius-options (edit system)

Syntax	<pre>radius-options { attributes { nas-ip-address <i>ip-address</i>; } password-protocol <i>mschap-v2</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
	<div> NOTE: The <code>radius-options</code> statement is not available on QFabric systems.</div>
Description	Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.
Options	<p><code>nas-ip-address <i>ip-address</i></code>—IP address of the network access server (NAS) that requests user authentication.</p> <p><code>password-protocol <i>mschap-v2</i></code>—Protocol MS-CHAPv2, used for password authentication and password changing.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Authentication• Configuring RADIUS Authentication on page 1699

radius-server

Syntax	<pre>radius-server server-address { accounting-port port-number; port number; retry number; secret password; source-address source-address; timeout seconds; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a RADIUS server for Point-to-Point Protocol (PPP).</p> <p>To configure multiple RADIUS servers, include multiple radius-server statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p>
Options	<p>server-address—Address of the RADIUS authentication server.</p> <p>The remaining statements are explained separately.</p>
<div>  <p>NOTE: The accounting-port and source-address options are not available on QFabric systems.</p> </div>	
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring RADIUS Authentication on page 1699 • accounting-port on page 234 • port on page 1784 • retry on page 284 • secret on page 288 • source-address on page 292 • timeout on page 307

rate-limit

Syntax	<code>rate-limit <i>limit</i>;</code>
Hierarchy Level	[edit system services finger], [edit system services ftp], [edit system services netconf ssh], [edit system services ssh], [edit system services telnet], [edit system services xnm-clear-text], [edit system services xnm-ssl]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the maximum number of connections attempts per protocol (either IPv6 or IPv4) on an access service.
Default	150 connections
Options	rate-limit <i>limit</i> —(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6). Range: 1 through 250 Default: 150
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i>

remote-debug-permission

Syntax	remote-debug-permission (qfabric-admin qfabric-operator qfabric-user);
Hierarchy Level	[edit system login user <i>username</i> authentication] [edit system root-authentication]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Configure authentication classes that permit or deny user access to individual components of the QFabric system.
Default	qfabric-user
Options	<p>qfabric-admin—Permits a user to log in to individual QFabric system components, view operations, and change component configurations.</p> <p>qfabric-operator—Permits a user to log in to individual QFabric system components and view component operations.</p> <p>qfabric-user—Prevents a user from logging in to individual QFabric system components.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring QFabric System Login Classes on page 1345 • request component login on page 1444 • Understanding QFabric System Login Classes on page 1230

retry

Syntax	<code>retry number;</code>
Hierarchy Level	[edit system radius server <i>server-address</i>], [edit system accounting destination radius server <i>server-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.
Options	<i>number</i> —Number of retries allowed for contacting a RADIUS server. Range: 1 through 10 Default: 3



NOTE: The [edit system accounting] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RADIUS Authentication on page 1699• Configuring RADIUS Accounting• timeout on page 307

rising-threshold (Health Monitor)

Syntax	<code>rising-threshold <i>percentage</i>;</code>
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<p><i>percentage</i>—Upper threshold for the alarm entry.</p> <p>Range: 1 through 100</p> <p>Default: 80 percent of the maximum possible value</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Health Monitoring on page 6200• falling-threshold on page 1772

root-login

Syntax	root-login (allow deny deny-password);
Hierarchy Level	[edit system services ssh]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Control user access through SSH.
Default	Allow user access through SSH.
Options	allow —Allow users to log in to the router or switch as root through SSH. deny —Disable users from logging in to the router or switch as root through SSH. deny-password —Allow users to log in to the router or switch as root through SSH when the authentication method (for example, RSA authentication) does not require a password.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Root Login Through SSH on page 1710

services (Switches)

Syntax

```
services {  
  service-deployment {  
    servers address {  
      port-number port-number;  
    }  
    source-address address;  
  }  
  ssh {  
    connection-limit limit;  
    protocol-version [v1 v2];  
    rate-limit limit;  
    root-login (allow | deny | deny-password);  
  }  
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the switch so that users on remote systems can access the local switch through SSH.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

snmp

```
Syntax  snmp {
    client-list client-list-name {
        ip-addresses;
    }
    community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                    addresses;
                }
            }
        }
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
        }
    }
}
```

```

    variable oid-variable;
}
event index {
    community community-name;
    description description;
    type type;
}
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance routing-instance-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
}
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance routing-instance-name;
        tag-list tag-list;
        target-parameters target-parameters-name;
    }
}

```

```
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
  usm {
    local-engine {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none;
      }
    }
    remote-engine engine-id {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none {
          privacy-password privacy-password;
        }
      }
    }
  }
}
```



```

}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure SNMP.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Understanding the Implementation of SNMP on page 6107](#)
- [Configuring SNMP on page 1703](#)

ssh

Syntax	<pre>ssh { ciphers [<i>cipher-1 cipher-2 cipher-3 ...</i>]; client-alive-count-max <i>seconds</i>; client-alive-interval <i>seconds</i>; connection-limit <i>limit</i>; hostkey-algorithm <<i>algorithm</i> no-<i>algorithm</i>>; key-exchange <<i>algorithm</i>>; macs <<i>algorithm</i>>; max-sessions-per-connection <<i>number</i>>; no-tcp-forwarding; protocol-version [<i>v1 v2</i>]; rate-limit <i>limit</i>; root-login (<i>allow</i> <i>deny</i> <i>deny-password</i>); }</pre>
Hierarchy Level	[edit system services]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>client-alive-interval and client-alive-max-count statements introduced in Junos OS Release 12.2.</p>
Description	<p>Allow SSH requests from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SSH Service for Remote Access to the Router or Switch on page 1709

system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```

```
altitude feet;  
building name;  
country-code code;  
floor number;  
hcoord horizontal-coordinate;  
lata service-area;  
latitude degrees;  
longitude degrees;  
npa-nxx number;  
postal-code postal-code;  
rack number;  
vcoord vertical-coordinate;  
}  
login {  
  announcement text;  
  class class-name {  
    access-end;  
    access-start;  
    allow-configuration "regular-expression";  
    allowed-days "regular-expression";  
    deny-commands "regular-expression";  
    deny-configuration "regular-expression";  
    idle-timeout minutes;  
    login-tip;  
    permissions [ permissions ];  
  }  
  message text;  
  password {  
    change-type (set-transitions | character-set);  
    format (md5 | sha1 | des);  
    maximum-length length;  
    minimum-changes number;  
    minimum-length length;  
  }  
  retry-options {  
    backoff-factor seconds;  
    backoff-threshold number;  
    minimum-time seconds;  
    tries-before-disconnect number;  
  }  
  user username {  
    authentication {  
      (encrypted-password "password" | plain-text-password);  
      load-key-file URL;  
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);  
      ssh-rsa "public-key";  
      ssh-dsa "public-key";  
    }  
    uid uid-value;  
    class class-name;  
    full-name complete-name;  
  }  
}  
name-server {  
  address;  
}
```

```

no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}

```

```
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}
```

```

}
console {
    facility severity;
}
file filename {
    archive {
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine | scc-master) {
    explicit-priority;
    facility-override facility;
    facility severity;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}
tacplus-options {
    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure system management properties.



.....

NOTE: The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

.....

Required Privilege system—To view this statement in the configuration.

Level system-control—To add this statement to the configuration.

tacplus-options

Syntax	<pre> tacplus-options { (exclude-cmd-attribute no-cmd-attribute-value); service-name <i>service-name</i>; timestamp-and-timezone; } </pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options for no-cmd-attribute-value and exclude-cmd-attribute introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Option for timestamp-and-timezone introduced in Junos OS Release 12.2.</p>
Description	Configure TACACS+ options for authentication and accounting.
Options	<p>exclude-cmd-attribute—Exclude the cmd attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p>no-cmd-attribute-value—Set the cmd attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p>service-name <i>service-name</i>—Name of the authentication service used when you configure multiple TACACS+ servers to use the same authentication service.</p> <p>Default: junos-exec</p> <p>timestamp-and-timezone—Include this statement if you want start time, stop time, and timezone attributes included in start/stop accounting records.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring TACACS+ Authentication • Configuring TACACS+ System Accounting • Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1676 • Configuring TACACS+ Authentication on page 1711 • Configuring TACACS+ System Accounting on page 1714

targets

Syntax	<pre>targets { address; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure one or more systems to receive SNMP traps.
Options	address —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 6195

traceoptions (LLDP)

Syntax `traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable> <no-stamp>
 <replace>;
 flag flag <disable>;
}`

Hierarchy Level [edit protocols [lldp](#)]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Define tracing operations for the Link Layer Discovery Protocol (LLDP). You can trace messages under LLDP for LLDP and PTOPO MIBs.



NOTE: The `traceoptions` statement is not supported on the QFX3000 QFabric system.

Default Tracing operations are disabled.

Options **file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **configuration**—Trace configuration operations.
- **interface**—Trace interface update events.
- **netbios**—Trace NetBIOS events.
- **packet**—Trace packet events.
- **rtsock**—Trace routing socket operations.
- **snmp**—Trace SNMP configuration operations.

- **vlan**—Trace VLAN update events.

no-stamp—(Optional) Do not timestamp the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending output to it.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

Default: If you do not include this option, tracing output is appended to an existing trace file.

world-readable—(Optional) Enable unrestricted file access.



NOTE: The **traceoptions** statement is not supported on the QFX3000 QFabric system.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring LLDP-MED (CLI Procedure)• Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches• Configuring LLDP on page 1693• Understanding LLDP on page 1666
------------------------------	---

transfer-interval (Configuration)

Syntax	<code>transfer-interval <i>interval</i>;</code>
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to periodically transfer its currently active configuration to an archive site.



NOTE: The `edit system archival` hierarchy is not available on QFabric systems.

Options *interval*—Interval at which to transfer the current configuration to an archive site.
Range: 15 through 2880 minutes



NOTE: The `[edit system archival]` hierarchy is not available on QFabric systems.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1611](#)
- [archive on page 6366](#)
- [configuration on page 1617](#)
- [transfer-on-commit on page 1619](#)

transfer-on-commit

Syntax	transfer-on-commit;
Hierarchy Level	[edit system archival configuration]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path".



NOTE: The [edit system archival] hierarchy is not available on QFabric systems.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1612• archive on page 6366• configuration on page 1617• transfer-interval on page 1618

trap-group

Syntax	<pre> trap-group <i>group-name</i> { categories { <i>category</i>; } destination-port <i>port-number</i>; routing-instance <i>instance</i>; targets { <i>address</i>; } version (all v1 v2); } </pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
Options	<p><i>group-name</i>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring SNMP Trap Groups

trap-options

Syntax	<pre>trap-options { agent-address outgoing-interface; source-address address; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
Default	Disabled
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Options</i>

user (Access)

Syntax	<pre> user username { authentication { (encrypted-password "password" plain-text-password); load-key-file URL; remote-debug-permission (qfabric-admin qfabric-operator qfabric-user); ssh-dsa "public-key" <from hostname>; ssh-rsa "public-key" <from hostname>; } class class-name; full-name "complete-name"; uid uid-value; } </pre>
Hierarchy Level	[edit system login]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure access permission for individual users.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Junos OS User Accounts on page 1692 • class on page 247

version

Syntax	version (all v1 v2);
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the version number of SNMP traps.
Default	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.
Options	all—Send an SNMPv1 and SNMPv2 trap for every trap condition. v1—Send SNMPv1 traps only. v2—Send SNMPv2 traps only.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 6195

Administration

- [Routine Monitoring on page 1831](#)
- [Monitoring Commands on page 1832](#)

Routine Monitoring

- [Monitoring SNMP on page 1831](#)

Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
```

```
Alarm
```

Index	Variable description	Value	State
32768	Health Monitor: root file system utilization jnxHrStoragePercentUsed.1	58	active
32769	Health Monitor: /config file system utilization jnxHrStoragePercentUsed.2	0	active
32770	Health Monitor: RE 0 CPU utilization jnxOperatingCPU.9.1.0.0	0	active
32773	Health Monitor: RE 0 Memory utilization jnxOperatingBuffer.9.1.0.0	35	active
32775	Health Monitor: jkernel daemon CPU utilization Init daemon	0	active

Chassis daemon	50 active
Firewall daemon	0 active
Interface daemon	5 active
SNMP daemon	11 active
MIB2 daemon	42 active
...	

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system
```

```
sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel  
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:  
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx
```

```
Build date: 2010-09-26 06:00:10 U  
sysObjectID.0 = jnxProductQFX3500  
sysUpTime.0   = 24444184  
sysContact.0  = J Smith  
sysName.0     = Lab QFX3500  
sysLocation.0 = Lab  
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics
```

SNMP statistics:

Input:

```
Packets: 0, Bad versions: 0, Bad community names: 0,  
Bad community uses: 0, ASN parse errors: 0,  
Too bigs: 0, No such names: 0, Bad values: 0,  
Read onlys: 0, General errors: 0,  
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,  
Throttle drops: 0, Duplicate request drops: 0
```

Output:

```
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

- Related Documentation
- [health-monitor on page 1773](#)
 - [show snmp mib on page 6452](#)
 - [show snmp statistics on page 1859](#)

Monitoring Commands

- [clear lldp neighbors](#)
- [clear lldp statistics](#)
- [request component login](#)
- [show ethernet-switching interfaces](#)

- `show lldp`
- `show lldp local-information`
- `show lldp neighbors`
- `show lldp statistics`
- `show route instance`
- `show snmp statistics`
- `ssh`

clear lldp neighbors

Syntax	clear lldp neighbors <interface <i>interface</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear the learned remote neighbor information on all or selected interfaces.
Options	none —Clear the remote neighbor information on all interfaces. interface <i>interface</i> —(Optional) Clear the remote neighbor information from the selected interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show lldp• Configuring LLDP on page 1693• Understanding LLDP on page 1666
List of Sample Output	clear lldp neighbors on page 1834 clear lldp neighbors interface on page 1834

Sample Output

clear lldp neighbors

```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```

clear lldp statistics

Syntax	<code>clear lldp statistics</code> <code><interface <i>interface</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear LLDP statistics on one or more interfaces.
Options	<p>none—Clears LLDP statistics on all interfaces.</p> <p>interface <i>interface-names</i>—(Optional) Clear LLDP statistics on an interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring LLDP on page 1693 • Understanding LLDP on page 1666
List of Sample Output	clear lldp statistics on page 1835 clear lldp statistics interface on page 1835

Sample Output

clear lldp statistics

```
user@switch> clear lldp statistics
```

clear lldp statistics interface

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

request component login

Syntax	<code>request component login <i>component-name</i></code>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the request component login command, you must first provide the qfabric-admin or qfabric-operator class privilege to your user (for more information, see: remote-debug-permission).
Options	<i>component-name</i> —Specify the QFabric system component to which you wish to log in.
Required Privilege Level	admin
Related Documentation	<ul style="list-style-type: none">• Example: Configuring QFabric System Login Classes on page 1345• remote-debug-permission on page 1401• Understanding QFabric System Login Classes on page 1230
List of Sample Output	request component login (with qfabric-admin Privileges) on page 1836 request component login (with qfabric-operator Privileges) on page 1837 request component login (with qfabric-user Privileges) on page 1837

Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
clear          Clear information in the system
file           Perform file operations
help           Provide help information
load           Load information from file
monitor        Show real-time debugging information
mtrace         Trace multicast path from source to receiver
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
ssh            Start secure shell on another host
start          Start shell
```



```

telnet          Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host{master}
qfabric-admin@node-ee3093>

```

request component login (with qfabric-operator Privileges)

```

operator@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-operator@node-ee3093> ?
Possible completions:
file          Perform file operations
help          Provide help information
load          Load information from file
op            Invoke an operation script
quit          Exit the management session
request       Make system-level requests
save          Save information to file
set           Set CLI properties, date/time, craft interface message
show          Show system information
start         Start shell
test          Perform diagnostic debugging
{master}
qfabric-operator@node-ee3093>

```

request component login (with qfabric-user Privileges)

```

user0@qfabric> request component login EE3093
error: User user0 does not have sufficient permissions to login to device ee3093

```

show ethernet-switching interfaces

Syntax	show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about switched Ethernet interfaces.
Options	<p>none—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display Ethernet-switching information for a specific interface.</p>
Required Privilege Level	view
List of Sample Output	show ethernet-switching interfaces on page 1839 show ethernet-switching interfaces summary on page 1840 show ethernet-switching interfaces brief on page 1840 show ethernet-switching interfaces detail on page 1840 show ethernet-switching interfaces interface-name on page 1841
Output Fields	Table 172 on page 1838 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 172: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are up or down .	none, brief , detail , summary
VLAN members	Name of a VLAN.	none, brief , detail , summary

Table 172: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Blocking	Forwarding state of the interface: <ul style="list-style-type: none"> • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface. • MAC limit exceeded—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • MAC move limit exceeded—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control in effect —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control shutdown in effect —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires. 	none, brief , detail , summary
Index	VLAN index internal to Junos OS software.	detail
untagged tagged	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	detail

Sample Output

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
xe-0/0/0.0	up	T1122	unblocked
xe-0/0/1.0	down	default	– MAC limit exceeded
xe-0/0/2.0	down	default	– MAC move limit exceeded
xe-0/0/3.0	down	default	– Storm control in effect
xe-0/0/4.0	down	default	unblocked
xe-0/0/5.0	down	default	unblocked
xe-0/0/6.0	down	default	unblocked
xe-0/0/7.0	down	default	unblocked
xe-0/0/8.0	down	default	unblocked
xe-0/0/9.0	up	T111	unblocked
xe-0/0/10.0	down	default	unblocked
xe-0/0/11.0	down	default	unblocked
xe-0/0/12.0	down	default	unblocked
xe-0/0/13.0	down	default	unblocked
xe-0/0/14.0	down	default	unblocked
xe-0/0/15.0	down	default	unblocked
xe-0/0/16.0	down	default	unblocked
xe-0/0/17.0	down	default	unblocked
xe-0/0/18.0	down	default	unblocked
xe-0/0/19.0	up	T111	unblocked
xe-0/1/0.0	down	default	unblocked
xe-0/1/1.0	down	default	unblocked
xe-0/1/2.0	down	default	unblocked
xe-0/1/3.0	down	default	unblocked

show ethernet-switching interfaces summary

```
user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

show ethernet-switching interfaces brief

```
user@switch> show ethernet-switching interfaces brief
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down   default       unblocked
xe-0/0/1.0  down   employee-vlan unblocked
xe-0/0/2.0  down   employee-vlan unblocked
xe-0/0/3.0  down   employee-vlan unblocked
xe-0/0/8.0  down   employee-vlan unblocked
xe-0/0/10.0 down   default       unblocked
xe-0/0/11.0 down   employee-vlan unblocked
```

show ethernet-switching interfaces detail

```
user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/1.0 Index: 66
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/2.0 Index: 67
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/3.0 Index: 68
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/8.0 Index: 69
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/10.0 Index: 70
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/11.0 Index: 71
State: down
VLANs:
    employee-vlan          tagged     unblocked
```

show ethernet-switching interfaces interface-name

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
  Interface  State   VLAN members  Blocking
xe-0/0/0.0  down    default        unblocked
```

show lldp

Syntax `show lldp`
`<detail>`

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol–Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.



NOTE: LLDP-MED is not available on the QFX Series.

Options **none**—Display LLDP information for all interfaces.
detail—(Optional) Display detailed LLDP information for all interfaces.



NOTE: **fast-start** is not available on the QFX Series.

Required Privilege Level view

Related Documentation

- *Configuring LLDP (CLI Procedure)*
- *Configuring LLDP-MED (CLI Procedure)*
- *Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches*
- [Configuring LLDP on page 1693](#)
- [Understanding LLDP on page 1666](#)

List of Sample Output [show lldp on page 1845](#)
[show lldp detail on page 1845](#)

Output Fields [Table 173 on page 1843](#) lists the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 173: show lldp Output Fields

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be enabled or disabled . NOTE: If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as disabled .	All levels
Advertisement interval	Frequency, in seconds, at which LLDP advertisements are sent. This value is set by the advertisement-interval configuration statement.	All levels
Transmit delay	Seconds of delay before advertisements are sent to neighbors following a change to a TLV (type, length, or value) element in the LLDP protocol or to the state of the local system, such as a change in hostname or management address. You can set this value to reduce the delay in notifying neighbors of a change in the local system. This value is set by the transmit-delay configuration statement.	All levels
Hold timer	Multiplier used in combination with the advertisement-interval value to determine the length of time LLDP information is held before it is discarded. This value is set by the hold-multiplier configuration statement.	All levels
Notification interval	How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications of database changes are disabled. This value is set by the lldp-configuration-notification-interval configuration statement.	All levels
Config Trap Interval	How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications of topology changes are disabled. This value is set by the ptopo-configuration-trap-interval configuration statement.	All levels
Connection Hold timer	Amount of time the system maintains dynamic topology entries. This value is set by the ptopo-configuration-maximum-hold-time configuration statement.	All levels
LLDP-MED	LLDP-MED operating state. The state can be enabled or disabled .	All levels
MED fast start count	Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second. This value is set by the fast-start configuration statement.	All levels
Interface	Name of the interface for which LLDP configuration information is being reported.	All levels
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.	All levels

Table 173: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be enabled or disabled .	All levels
Power Negotiation	LLDP power negotiation operating state. The state can be Enabled or Disabled .	All levels
Neighbor count	Total number of new LLDP neighbors detected since the last switch reboot.	detail
Interface	Name of the interface that is advertising VLAN information.	All levels
Vlan-id	VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.	detail
Vlan-name	VLAN name associated with the VLAN ID.	detail
LLDP basic TLVs supported	<p>Basic TLVs supported on the switch:</p> <ul style="list-style-type: none"> • Chassis identifier—TLV that advertises the MAC address associated with the local system. • Port identifier—TLV that advertises the port identification for the specified port in the local system. • Port description—Interface name for the port. • System name—TLV that advertises the user-configured name of the local system. • System description—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable. • System capabilities—TLV that advertises the primary functions performed by the system—for example, bridge or router. • Management address—TLV that advertises the IP management address of the local system. 	detail
Supported LLDP 802 TLVs	<p>802.3 TLVs supported on the switch:</p> <ul style="list-style-type: none"> • MAC/PHY configuration status—TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is based on the physical interface structure and is not configurable. • Power via MDI—TLV that advertises MDI power support, PSE power pair, and power class information. • Link aggregation—TLV that advertises if the interface is aggregated and its aggregated interface ID. • Maximum frame size—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames. • Port VLAN tag—TLV that advertises the VLAN tag configured on the interface. • Port VLAN name—TLV that advertises the VLAN name configured on the interface. 	detail

Table 173: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Supported LLDP MED TLVs	<p>LLDP-MED TLVs supported on the switch:</p> <ul style="list-style-type: none"> • LLDP MED capabilities—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15: <ul style="list-style-type: none"> • 0—Capabilities • 1—Network Policy • 2—Location Identification • 3—Extended Power via MDI-PSE • 4—Inventory • 5–15—Reserved • Network policy—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points. • Endpoint location—TLV that advertises the physical location of the endpoint. • Extended power Via MDI—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port. 	detail

Sample Output

show lldp

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                         : 4 seconds
Notification interval              : 0 Second(s)
Config Trap Interval               : 0 seconds
Connection Hold timer              : 300 seconds

LLDP MED                           : Disabled
MED fast start count               : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
all	-	Enabled	Enabled	Enabled

show lldp detail

```

user@switch> show lldp detail
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                         : 4 seconds
Notification interval              : 0 Second(s)
Config Trap Interval               : 0 seconds
Connection Hold timer              : 300 seconds

LLDP MED                           : Disabled
MED fast start count               : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
Neighbor count				
all	-	Enabled	Enabled	Enabled
8				

Interface	Parent Interface	Vlan-id	Vlan-name
xe-3/0/0.0	ae31.0	100	v100
xe-3/0/0.0	ae31.0	101	v101
xe-3/0/0.0	ae31.0	4000	v4000
xe-3/0/1.0	ae31.0	100	v100
xe-3/0/1.0	ae31.0	101	v101
xe-3/0/1.0	ae31.0	4000	v4000
xe-3/0/2.0	ae31.0	100	v100
xe-3/0/2.0	ae31.0	101	v101
xe-3/0/2.0	ae31.0	4000	v4000

LLDP basic TLVs supported:

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

Supported LLDP 802 TLVs:

MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port VLAN name.

Supported LLDP MED TLVs:

LLDP MED capabilities, Network policy, Endpoint location, Extended power Via MDI.

show lldp local-information

Syntax	show lldp local-information
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring LLDP (CLI Procedure)</i> • <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i> • management-address on page 1780 • Configuring LLDP on page 1693 • Understanding LLDP on page 1666
List of Sample Output	show lldp local-information (EX Series Switch) on page 1848
Output Fields	Table 174 on page 1847 lists the output fields for the show lldp local-information command. Output fields are listed in the approximate order in which they appear.

Table 174: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	<p>Information about the local system (the switch):</p> <ul style="list-style-type: none"> • Chassis ID—MAC address associated with the switch. • System name—User-configured name of the switch. • System descr—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.
System Capabilities	Capabilities (such as bridge or router) that are supported or enabled on the system.
Management Information	<p>Details of the management information: Port Name, Port Address (such as 10.204.34.35), Address Type (such as ipv4 or ipv6), Port ID (SNMP interface index), Port ID Subtype, and Port Subtype.</p> <p>The Port Subtype displays:</p> <ul style="list-style-type: none"> • ifindex(2)—IP address of the switch's management Ethernet interface (me0) or virtual management Ethernet (VME) interface address (for a virtual chassis) is used to manage the switch. • unknown(1)—IP management address has been configured with set protocols lldp management-address.

Table 174: show lldp local-information Output Fields (*continued*)

Field Name	Field Description
Interface name	Name of the local interface which is configured for either LLDP or LLDP-MED.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the local interface belongs.
SNMP Index	SNMP interface index.
Interface description	User-configured port description.
Status	Administrative status of the interface: either up or down .
Tunneling	Status of tunneling on the interface: either enabled or disabled .

Sample Output

show lldp local-information (EX Series Switch)

```
user@switch> show lldp local-information
```

LLDP Local Information details

```
Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.4I0 [builder] Build
               date: 2010-11-17 12:38:30 UTC
```

System Capabilities

```
Supported   : Bridge Router
Enabled     : Bridge Router
```

Management Information

```
Port Name    : -
Port Address : 10.93.54.6
Address Type  : IPv4
Port ID       : 34
Port ID Subtype : local(7)
Port Subtype  : ifIndex(2)
```

Interface name	Parent Interface	SNMP Index	Interface description	Status	Tunneling
me0.0	-	34	-	Down	Disabled
xe-3/0/0.0	ae31.0	769	xe-3/0/0.0	Up	Disabled
xe-3/0/1.0	ae31.0	770	xe-3/0/1.0	Up	Disabled
xe-3/0/2.0	ae31.0	771	xe-3/0/2.0	Up	Disabled
xe-3/0/3.0	ae31.0	772	xe-3/0/3.0	Up	Disabled
xe-3/0/4.0	ae31.0	577	xe-3/0/4.0	Up	Disabled
xe-3/0/5.0	ae31.0	578	xe-3/0/5.0	Up	Disabled
xe-3/0/6.0	ae31.0	579	xe-3/0/6.0	Up	Disabled
xe-3/0/7.0	ae31.0	581	xe-3/0/7.0	Up	Disabled

show lldp neighbors

Syntax <show lldp *neighbors*>
<interface *interface-ids*>

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display learned information about Link Layer Discovery Protocol (LLDP) on all neighboring interfaces or on selected interfaces.

Options **none**—Display learned LLDP information on all neighboring interfaces and devices.

interface *interface-ids*—(Optional) Display learned LLDP information on the selected interfaces or devices.



NOTE: When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors in order to interoperate with a wider variety of converged network adapters (CNAs). As a result, information for those ports will not be listed in the output for this command.

Required Privilege Level view

Related Documentation

- [Configuring LLDP on page 1693](#)
- [Understanding LLDP on page 1666](#)

List of Sample Output [show lldp neighbors on page 1851](#)
[show lldp neighbors interface on page 1852](#)

Output Fields [Table 175 on page 1849](#) lists the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 175: show lldp neighbors Output Fields

Field Name	Field Description
Local Interface	List of local interfaces for which neighbor information is available.
Parent Interface	List of aggregated Ethernet interfaces, if any, to which the local interfaces belong.
Chassis ID	List of chassis identifiers for neighbors.
Port info	List of port information gathered from neighbors. This could be the port identifier or port description.
System name	List of system names gathered from neighbors.

Table 175: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
LLDP Neighbor Information	Information about both the local system (the switch) and a neighbor system on the interface (appears when the interface option is used).
Local Information	Information about the local system (appears when the interface option is used).
Index	Local interface index (appears when the interface option is used).
Time to live	Number of seconds for which this information is valid (appears when the interface option is used).
Time mark	Date and timestamp of information (appears when the interface option is used).
Local Interface	Name of the local physical interface (appears when the interface option is used).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the interface option is used).
Local Port ID	Local interface SNMP index (appears when the interface option is used).
Ageout Count	Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval has expired (appears when the interface option is used).
Neighbor Information	Information about a neighbor system on the interface (appears when the interface option is used).
Chassis type	Type of chassis identifier supplied, such as MAC address (appears when the interface option is used).
Chassis ID	Chassis identifier of the chassis type listed (appears when the interface option is used).
Port type	Type of port identifier supplied, such as locally assigned (appears when the interface option is used).
Port ID	Port identifier of the port type listed (appears when the interface option is used).
Port description	Port description (appears when the interface option is used).
System name	Name supplied by the system on the interface (appears when the interface option is used).
System Description	Description supplied by the system on the interface (appears when the interface option is used).

Table 175: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
System capabilities	Capabilities (such as Bridge , Router , and Telephone) that are supported or enabled by the system on the interface (appears when the interface option is used).
Management Info	<p>Details of management information: Type (such as ipv4 or ipv6), Address (such as 10.204.34.35), Port ID, Subtype, Interface Subtype, and organization identifier (OID) (appears when the interface option is used).</p> <p>The Interface Subtype displays:</p> <ul style="list-style-type: none"> • ifindex(2)—IP address of the neighbor's management Ethernet interface (me0) or virtual management Ethernet (VME) interface address (for a virtual chassis) is used to manage the switch. • unknown(1)—Neighbor's IP management address has been configured with set protocols lldp management-address.
Media Info	Additional details about the endpoint device appear when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details might include Media endpoint class (such as Class 3 for communication devices such as IP phones), MED Hardware revision , MED Firmware revision , MED Software revision , MED Serial number , MED Manufacturer name , or MED Model name .
Organization Info	One or more entries listing remote information by organizationally unique identifier (OUI), Subtype , Index , and Info (appears when the interface option is used).
Age	How long the neighbor has been identified (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Local Interface	Name of the local physical interface (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Chassis ID	Chassis identifier of the chassis type listed (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
Port description	Port description (appears when the interface option is used and NetBIOS snooping is enabled on the switch).
System name	NetBIOS name of the host (appears when the interface option is used and NetBIOS snooping is enabled on the switch).

Sample Output

show lldp neighbors

```
user@switch> show lldp neighbors
```

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-3/0/4.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/0.0	newyork31
xe-3/0/5.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/1.0	newyork31
xe-3/0/6.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/2.0	newyork31
xe-3/0/7.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/3.0	newyork31
xe-3/0/0.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/0.0	newyork31
xe-3/0/1.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/1.0	newyork31
xe-3/0/2.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/2.0	newyork31
xe-3/0/3.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/3.0	newyork31

show lldp neighbors interface

```
user@switch> show lldp neighbors interface ge-0/0/2
```

LLDP Neighbor Information:

Local Information:

```
Index: 1 Time to live: 240 Time mark: Wed Dec 1 10:23:24 2010 Age: 29 secs
Local Interface   : ge-0/0/2.0
Parent Interface  : -
Local Port ID     : 507
Ageout Count      : 0
```

Neighbour Information:

```
Chassis type      : Mac address
Chassis ID        : 00:1f:12:38:7f:c0
Port type         : Locally assigned
Port ID           : 507
Port description  : ge-0/0/2.0
System name       : bng-148p5-dev
```

```
System Description : Juniper Networks, Inc. ex4200-48p , version 10.4I0 Build
date: 2010-11-30 09:32:17 UTC
```

System capabilities

```
Supported : Bridge Router
Enabled   : Bridge Router
```

Management Info

```
Type           : IPv4
Address         : 10.204.96.235
Port ID        : 34
Subtype        : 1
Interface Subtype : ifIndex(2)
OID            : 1.3.6.1.2.1.31.1.1.1.1.34
```

```
Media endpoint class: Network Connectivity
```

Organization Info

```
OUI       : 0.12.f
Subtype   : 1
Index     : 1
Info      : 22A8360000
```

Organization Info

```
OUI       : 0.12.f
Subtype   : 2
Index     : 2
Info      : 030100
```


show lldp statistics

Syntax	<code>show lldp statistics</code> <code><interface <i>interface-ids</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display LLDP statistics on all or selected interfaces.
Options	<p>none—Display LLDP statistics on all interfaces and devices.</p> <p>interface <i>interface-ids</i>—(Optional) Display LLDP statistics on the selected devices.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring LLDP on page 1693 • Understanding LLDP on page 1666
List of Sample Output	show lldp statistics on page 1853
Output Fields	Table 176 on page 1853 lists the output fields for the show lldp statistics command. Output fields are listed in the approximate order in which they appear.

Table 176: show lldp statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Name of an interface.	All levels
Received	Total number of LLDP frames received on an interface.	All levels
Unknown-TLVs	Number of unrecognized LLDP TLVs received on an interface.	All levels
With Errors	Number of LLDP frames received that contain errors.	All levels
Discarded TLVs	Number of LLDP TLVs received and then discarded on an interface.	All levels
Transmitted	Total number of LLDP frames transmitted on an interface.	All levels
Untransmitted	Total number of LLDP frames not transmitted on an interface.	All levels

Sample Output

show lldp statistics

```
user@switch> show lldp statistics
```

```

Interface  Received  Unknown TLVs  With Errors  Discarded TLVs  Transmitted
Untransmitted
me0.0      0         0             0           0               8003         0

```

ge-0/0/0.0 8002	0	0	0	8003	0
ge-0/0/1.0 8002	0	0	0	8003	0

show route instance

Syntax	show route instance <brief detail summary> <instance-name> <operational>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric systems only) Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for a specified routing instance.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show route instance on page 1856 show route instance detail on page 1856 show route instance operational on page 1857 show route instance summary on page 1857
Output Fields	Table 177 on page 1855 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.

Table 177: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding or virtual-router .	All levels
State	State of the routing instance: active or inactive .	detail
Interfaces	Name of interfaces belonging to this routing instance.	detail
Tables	Tables (and number of routes) associated with this routing instance.	detail
Router ID	Identifier for the router.	detail

Table 177: show route instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Primary RIB	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@switch> show route instance
Instance          Type
      Primary RIB
master            forwarding
      inet.0
                  4/0/1

__juniper_private1__ forwarding
      __juniper_private1__.inet.0
                  1/0/3

__juniper_private2__ forwarding
      __juniper_private2__.inet.0
                  0/0/1

__juniper_private3__ forwarding
      __juniper_private3__.inet.0
                  1/0/2

__juniper_private4__ forwarding
      __juniper_private4__.inet.0
                  4/0/2

__master.anon__    forwarding

r1                 virtual-router

r2                 virtual-router

```

show route instance detail

```

user@switch> show route instance detail
master:
  Router ID: 3.3.3.7
  Type: forwarding      State: Active
  Tables:
    inet.0              : 5 routes (4 active, 0 holddown, 1 hidden)

__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16385
    bme0.0
  Tables:
    __juniper_private1__.inet.0: 6 routes (1 active, 0 holddown, 3 hidden)

__juniper_private2__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16384

```

```

Tables:
  __juniper_private2__.inet.0: 1 routes (0 active, 0 holddown, 1 hidden)

__juniper_private3__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.1
Tables:
  __juniper_private3__.inet.0: 4 routes (1 active, 0 holddown, 2 hidden)

__juniper_private4__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.2
Tables:
  __juniper_private4__.inet.0: 8 routes (4 active, 0 holddown, 2 hidden)

__master.anon__:
Router ID: 0.0.0.0
Type: forwarding      State: Active

r1:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/0.0

r2:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/3.0

```

show route instance operational

```

user@switch> show route instance operational
Operational Routing Instances:

__juniper_private1__
__juniper_private2__
__juniper_private3__
__juniper_private4__
r1---qfabric
r2---qfabric
master

```

show route instance summary

```

user@switch> show route instance summary

```

Instance	Type	Primary RIB	Active/holddown/hidden
master	forwarding	inet.0	4/0/1
__juniper_private1__	forwarding	__juniper_private1__.inet.0	1/0/3
__juniper_private2__	forwarding	__juniper_private2__.inet.0	0/0/1

__juniper_private3__ forwarding	
__juniper_private3__.inet.0	1/0/2
__juniper_private4__ forwarding	
__juniper_private4__.inet.0	4/0/2
__master.anon__ forwarding	
r1	virtual-router
r2	virtual-router

show snmp statistics

Syntax	show snmp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear snmp statistics on page 6403
List of Sample Output	show snmp statistics on page 1862
Output Fields	Table 178 on page 1859 describes the output fields for the show snmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 178: show snmp statistics Output Fields

Field Name	Field Description
Input	<p>Information about received packets:</p> <ul style="list-style-type: none"> • Packets(snmplnPkts)—Total number of messages delivered to the SNMP entity from the transport service. • Bad versions—(snmplnBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • Bad community names—(snmplnBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • Bad community uses—(snmplnBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • ASN parse errors—(snmplnASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • Too big—(snmplnTooBig) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig. • No such names—(snmplnNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmplnBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue. • Read only—(snmplnReadOnly) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

Table 178: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> • General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr. • Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs. • Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs. • Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity. • Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity. • Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity. • Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity. • Traps—(snmpInTraps) Total number of SNMP traps generated by the SNMP entity. • Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • Commit pending drops—Number of SNMP packets for Set requests dropped because of a previous pending SNMP Set request on the committed configuration. • Throttle drops—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 178: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> • Unknown security models—(snmpUnknownSecurityModels) Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • Invalid messages—(snmpInvalidMsgs) Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • Unknown pdu handlers—(snmpUnknownPDUHandlers) Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • Unavailable contexts—(snmpUnavailableContexts) Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • Unknown contexts—(snmpUnknownContexts) Total number of requests received for a context that is unknown to the SNMP engine. • Unsupported security levels—(usmStatsUnsupportedSecLevels) Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • Not in time windows—(usmStatsNotInTimeWindows) Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • Unknown user names—(usmStatsUnknownUserNames) Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • Unknown engine ids—(usmStatsUnknownEngineIDs) Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • Wrong digests—(usmStatsWrongDigests) Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • Decryption errors—(usmStatsDecryptionErrors) Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 178: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Output	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> • Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service. • Too big—(snmpOutTooBigs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig. • No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue. • General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of genErr. • Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs generated by the SNMP entity. • Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs generated by the SNMP entity. • Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs generated by the SNMP entity. • Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs generated by the SNMP entity. • Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Sample Output

show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read only: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0

```

ssh

List of Syntax [Syntax on page 1863](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 1863](#)

Syntax `ssh host`
 `<bypass-routing>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<logical-system logical-system-name>`
 `<routing-instance routing-instance-name>`
 `<source address>`
 `<v1 | v2>`

Syntax (EX Series Switch and the QFX Series) `ssh host`
 `<bypass-routing>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<routing-instance routing-instance-name>`
 `<source address>`
 `<v1 | v2>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the **ssh** command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

```
ssh address 'cli-command1 ; cli-command2 '
```

Options **host**—Name or address of the remote system.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

inet | inet6—(Optional) Create an IPv4 or IPv6 connection, respectively.

interface interface-name—(Optional) Interface name for the SSH session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

logical-system logical-system-name—(Optional) Name of a particular logical system for the SSH attempt.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the SSH attempt.

source address—(Optional) Source address of the SSH connection.

v1 | v2—(Optional) Use SSH version 1 or 2, respectively, when connecting to a remote host.

Additional Information To configure an SSH (version 1) key for your user account, include the **authentication ssh-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. To configure an SSH (version 2) key for your user account, include the **authentication dsa-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. For details, see the .

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the **retry-options** statement at the **[edit system login]** hierarchy level. For details, see the .

Required Privilege Level network

Related Documentation

- [Configuring SSH Host Keys for Secure Copying of Data on page 1707](#)

List of Sample Output [ssh on page 1864](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

ssh

```
user@switch> ssh cree
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?cree' added to the list of known hosts.
boojun@cree's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```

PART 8

Ethernet Features

- [Overview on page 1867](#)
- [Configuration on page 1917](#)
- [Administration on page 2173](#)
- [Troubleshooting on page 2261](#)

CHAPTER 20

Overview

- [Software Features Overview on page 1867](#)
- [Bridging and VLANs on page 1873](#)
- [Spanning Trees Overview on page 1899](#)
- [Q-in-Q Tunneling on page 1906](#)
- [Layer 2 Protocol Tunneling on page 1910](#)
- [Proxy ARP on page 1914](#)

Software Features Overview

- [Overview of Layer 2 Networking on page 1867](#)
- [Understanding Bridging on page 1869](#)
- [Understanding Unicast on page 1871](#)
- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 1871](#)
- [Understanding Layer 2 Broadcasting on page 1872](#)

Overview of Layer 2 Networking

Layer 2, also known as the Data Link Layer, is the second level in the seven-layer OSI reference model for network protocol design. Layer 2 is equivalent to the link layer (the lowest layer) in the TCP/IP network model. Layer 2 is the network layer used to transfer data between adjacent network nodes in a wide area network or between nodes on the same local area network.

A *frame* is a protocol data unit, the smallest unit of bits on a Layer 2 network. Frames are transmitted to and received from devices on the same local area network (LAN). Unlike bits, frames have a defined structure and can be used for error detection, control plane activities and so forth. Not all frames carry user data. The network uses some frames to control the data link itself..

At Layer 2, *unicast* refers to sending frames from one node to a single other node, whereas *multicast* denotes sending traffic from one node to multiple nodes, and *broadcasting* refers to the transmission of frames to all nodes in a network. A *broadcast domain* is a logical division of a network in which all nodes of that network can be reached at Layer 2 by a broadcast.

Segments of a LAN can be linked at the frame level using *bridges*. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN.

Forwarding is the relaying of packets from one network segment to another by nodes in the network. On a VLAN, a frame whose origin and destination are in the same VLAN are forwarded only within the local VLAN. A network segment is a portion of a computer network wherein every device communicates using the same physical layer.

Layer 2 contains two sublayers:

- Logical link control (LLC) sublayer, which is responsible for managing communications links and handling frame traffic.
- Media access control (MAC) sublayer, which governs protocol access to the physical network medium. By using the MAC addresses that are assigned to all ports on a switch, multiple devices on the same physical link can uniquely identify one another.

The ports, or interfaces, on a switch operate in either access mode, tagged-access, or trunk mode:

- *Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode.
- *Tagged-Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:
- *Trunk mode* ports handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to other devices or switches.

With native VLAN configured, frames that do not carry VLAN tags are sent over the trunk interface. If you have a situation where packets pass from a device to a switch in access mode, and you want to then send those packets from the switch over a trunk port, use native VLAN mode. Configure the single VLAN on the switch's port (which is in access mode) as a native VLAN. The switch's trunk port will then treat

those frames differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, frames on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag). There is another native VLAN option. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Including the sublayers, Layer 2 on the QFX Series supports the following functionality:

- Unicast, multicast, and broadcast traffic.
- Bridging.
- VLAN 802.1Q—Also known as *VLAN tagging*, this protocol allows multiple bridged networks to transparently share the same physical network link by adding VLAN tags to an Ethernet frame.
- Extension of Layer 2 VLANs across multiple switches using Spanning Tree Protocol (STP) prevents looping across the network.
- *MAC learning*, including per-VLAN MAC learning and Layer 2 learning suppression—This process obtains the MAC addresses of all the nodes on a network
- Link aggregation—This process groups of Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or LAG bundle
- Storm control on the physical port for unicast, multicast, and broadcast
- STP support, including 802.1d, RSTP, MSTP, and Root Guard

**Related
Documentation**

- [Understanding VLANs on page 1873](#)
- [Understanding Bridging on page 1869](#)

Understanding Bridging

Bridging is a frame-forwarding technique used in Layer 2. With bridging, there are no assumptions about where a particular address is located in a network, so a mixed collection of interface types and speeds can be logically grouped together.

Layer 2 on a switch uses the *transparent bridging* protocol. This protocol enables a node to learn information about all the nodes on the LAN, including nodes on all the different VLANs. This information is used to create address-lookup tables, called *Ethernet switching tables*. The switch uses these tables when forwarding traffic to or toward a destination on the LAN.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the node:

- *Learning* is the process used to obtain the MAC addresses of all the nodes on a network. When a node is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. To learn MAC addresses, a node reads all detected packets on the LAN or on the local VLAN, looking for the MAC addresses of sending nodes. These addresses are added to the Ethernet switching table, along with the interface (or port) on which the traffic was received and the time when the address was learned.
- *Forwarding* describes how traffic passes from an incoming to an outgoing interface that leads to or toward the destination. To forward frames, the Ethernet switching table is read to determine if the table contains the MAC address corresponding to the frame destination. If the Ethernet switching table contains an entry for the desired destination address, the traffic is sent out to the interface associated with the MAC address. The Ethernet switching table is read in the same way when frames are transmitted that originate on nodes connected directly to the initiating node.
- *Flooding* is the mechanism by which a node learns about destinations not in its Ethernet switching table. If the table has no entry for a particular destination MAC address, traffic floods out to all interfaces except the interface on which it was received. (If traffic originates on the node, the device floods it out all interfaces.) When the destination node receives the flooded traffic, it returns an acknowledgment packet, allowing the node to learn the MAC address and add it to its Ethernet switching table.
- *Filtering* refers to the limiting of traffic to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, an increasingly complete picture of the VLAN and the larger LAN is built, showing which nodes are in the local VLAN and which are on other network segments. This information is used to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents traffic forwarding to other network segments.
- *Aging* keeps the entries in the Ethernet switching table current. For each MAC address in the table, a timestamp shows when the information about the network node was learned. Each time traffic from a MAC address is detected, the timestamp is updated. A timer on the node periodically checks the timestamp, and if it is older than a user-configured value, that MAC address is removed from the Ethernet switching table. This aging process ensures that the device tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

**Related
Documentation**

- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 1871](#)
- [Understanding MAC Learning on page 1876](#)
- [Understanding VLANs on page 1873](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)

Understanding Unicast

Unicasting is the act of sending data from one node of the network to another. In contrast, multicast transmissions send traffic from one data node to multiple other data nodes.

Unknown unicast traffic consists of unicast frames with unknown destination MAC addresses. By default, the switch floods these unicast frames that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic to interfaces on the switch can trigger a security issue. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm.

To prevent a storm, you can disable the flooding of unknown unicast packets to all interfaces by configuring one VLAN or all VLANs to forward any unknown unicast traffic to a specific trunk interface. (This channels the unknown unicast traffic to a single interface.)

Related Documentation

- [Overview of Layer 2 Networking on page 1867](#)
- [Understanding VLANs on page 1873](#)

Introduction to the Media Access Control (MAC) Layer 2 Sublayer

This topic provides an introduction to the MAC sublayer of the data link layer (Layer 2).

In Layer 2 of a network, the Media Access Control (MAC) sublayer provides addressing and channel access control mechanisms that enable several terminals or network nodes to communicate in a network.

The MAC sublayer acts as an interface between the logical link control (LLC) Ethernet sublayer and Layer 1 (the physical layer). The MAC sublayer emulates a full-duplex logical communication channel in a multipoint network. This channel may provide unicast, multicast, or broadcast communication service. The MAC sublayer uses MAC protocols to prevent collisions.

In Layer 2, multiple devices on the same physical link can uniquely identify one another at the data link layer, by using the MAC addresses that are assigned to all ports on a switch. A MAC algorithm accepts as input a secret key and an arbitrary-length message to be authenticated, and outputs a MAC address.

A MAC address is a 12-digit hexadecimal number (48 bits in long). MAC addresses are usually written in one of these formats:

- MM:MM:MM:SS:SS:SS
- MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the adapter manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the adapter by the manufacturer.

Contrast MAC addressing, which works at Layer 2, with IP addressing, which runs at Layer 3 (networking and routing). One way to remember the difference is that the MAC addresses apply to a physical or virtual node, whereas IP addresses apply to the software implementation of that node. MAC addresses are typically fixed on a per-node basis, whereas IP addresses change when the node moves from one part of the network to another.

IP networks maintain a mapping between the IP and MAC addresses of a node using the Address Resolution Protocol (ARP) table. DHCP also typically uses MAC addresses when assigning IP addresses to nodes.

- Related Documentation**
- [Overview of Layer 2 Networking on page 1867](#)
 - [Understanding MAC Learning on page 1876](#)

Understanding Layer 2 Broadcasting

In a Layer 2 network, *broadcasting* refers to sending traffic to all nodes on a network.

Layer 2 broadcast traffic stays within a local area network (LAN) boundary; known as the *broadcast domain*. Layer 2 broadcast traffic is sent to the broadcast domain using a MAC address of FF:FF:FF:FF:FF:FF. Every device in the broadcast domain recognizes this MAC address and passes the broadcast traffic on to other devices in the broadcast domain, if applicable. Broadcasting can be compared to unicasting (sending traffic to a single node) or multicasting (delivering traffic to a group of nodes simultaneously).

Layer 3 broadcast traffic, however, is sent to all devices in a network using a broadcast network address. For example, if your network address is 192.0.0.0, the broadcast network address is 192.255.255.255. In this case, only devices that belong to the 192.0.0.0 network receive the Layer 3 broadcast traffic. Devices that do not belong to this network drop the traffic.

Broadcasting is used in the following situations:

- Address Resolution Protocol (ARP) uses broadcasting to map MAC addresses to IP addresses. ARP dynamically binds the IP address (the logical address) to the correct MAC address. Before IP unicast packets can be sent, ARP discovers the MAC address used by the Ethernet interface where the IP address is configured.
- Dynamic Host Configuration Protocol (DHCP) uses broadcasting to dynamically assign IP addresses to hosts on a network segment or subnet.
- Routing protocols use broadcasting to advertise routes.

Excessive broadcast traffic can sometimes create a broadcast storm. A broadcast storm occurs when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses that create a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service.

- Related Documentation**
- [Overview of Layer 2 Networking on page 1867](#)
 - [Understanding Storm Control on page 4694](#)
 - [Understanding Bridging on page 1869](#)
 - [Understanding VLANs on page 1873](#)

Bridging and VLANs

- [Understanding VLANs on page 1873](#)
- [Understanding Routed VLAN Interfaces on page 1875](#)
- [Understanding MAC Learning on page 1876](#)
- [Understanding Reflective Relay for Use with VEPA Technology on page 1877](#)
- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887](#)
- [Understanding Egress Firewall Filters with PVLANS on page 1896](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on page 1897](#)

Understanding VLANs

A virtual local area network (virtual LAN, or *VLAN*) is a collection of network nodes that are grouped together to form separate *broadcast domains*.

- [Introduction to VLANs on page 1873](#)
- [VLAN Tagging on page 1874](#)
- [Assigning Traffic to a VLAN on page 1874](#)
- [Ethernet Switching Tables on page 1874](#)
- [Layer 2 and Layer 3 Forwarding of VLAN Traffic on page 1874](#)

Introduction to VLANs

On an Ethernet network that is a single LAN, all traffic is forwarded to all nodes on the LAN. A VLAN has the same attributes as a physical LAN, but allows end stations to be grouped together. An advantage to a VLAN is that you can perform network reconfiguration through the software rather than having to physically relocate a device.

On VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN. Frames that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. VLANs limit the amount of traffic flowing across the entire LAN, reducing the number of possible collisions and packet retransmissions within a VLAN and on the LAN as a whole.

On an Ethernet LAN, all network nodes must be physically connected to the same network. On VLANs, the physical location of the nodes is not important, so you can group network nodes in any way that makes sense for your organization, such as by department or

business function, types of network nodes, or even physical location. Each VLAN is identified by a single IP subnetwork and by standardized IEEE 802.1Q encapsulation.

In networks that use Fibre Channel over Ethernet (FCoE), a VLAN that uses FCoE frames must be dedicated to FCoE traffic only; you cannot mix standard Ethernet and FCoE traffic on the same VLAN.

QFX Series systems support a maximum of 4089 VLANs, which includes the default VLAN. You can assign a VLAN ID in the range of 1 to 4094.



NOTE: You cannot mix Fibre Channel and Ethernet ports on the same VLAN. For further information about VLANs and Fibre Channel, see [“Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric”](#) on page 5099.

VLAN Tagging

When you divide an Ethernet LAN into multiple VLANs, each VLAN is assigned a unique IEEE 802.1Q tag. This tag is associated with each frame in the VLAN, and the network nodes receiving the traffic can use the tag to identify which VLAN a frame is associated with. For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

Assigning Traffic to a VLAN

You can assign traffic to a VLAN in one of the following ways:

- *Interface (port)*—You can specify that all traffic received on a particular interface is assigned to a specific VLAN. By default, all traffic received on an access interface is untagged. This traffic is part of a default VLAN, but does not have an 802.1Q tag. During configuration, you specify which VLAN to assign the traffic to. You configure the VLAN either by using a VLAN ID number or by using a name, which is translated into a numeric VLAN ID.
- *MAC address*—You can specify that all traffic received from a specific MAC address be forwarded to a specific egress port (the next hop port). This method is cumbersome to configure manually, but it can be useful when you are using automated databases to manage the devices on your network.

Ethernet Switching Tables

Learned MAC addresses on local VLANs are stored in a [“bridge”](#) on page 1869. With each MAC address, the Ethernet switching table stores and associates the name of the learned port address. The information in this table is used when packets are forwarded toward their destination.

Layer 2 and Layer 3 Forwarding of VLAN Traffic

To pass traffic within a VLAN, Layer 2 forwarding protocols are used, including IEEE 802.1Q and Spanning Tree Protocol (STP),

To pass traffic between two VLANs, standard Layer 3 routing protocols are used, such as static routing. On the QFX Series, the same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Understanding FCoE on page 4968](#)
- [Interfaces Overview on page 2415](#)

Understanding Routed VLAN Interfaces

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). VLANs limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you normally you need a router that connects the VLANs. However, you can accomplish this forwarding on a switch without using a router by configuring a routed VLAN interface (RVI). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

An RVI is a special type of Layer 3 virtual interface named **vlan**. Like normal Layer 3 interfaces, the **vlan** interface needs a logical unit number with an IP address. In fact, to be useful an RVI needs at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your RVI must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs. [Table 179 on page 1875](#) shows values you might use when configuring an RVI:

Table 179: Sample RVI Values

Property	Settings
VLAN names and tags (IDs)	blue , ID 100 red , ID 200
Subnets associated with VLANs	blue : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) red : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
RVI name	interface vlan
RVI units and addresses	logical unit 100: 192.0.2.1/25 logical unit 200: 192.0.2.129/25

For the sake of consistency and to avoid confusion, [Table 179 on page 1875](#) shows RVI logical unit numbers that match the IDs of the corresponding VLANs. However, you do not have to assign logical unit numbers that match the VLAN IDs—you can use any values for the units. To bind the logical units of the RVI to the appropriate VLANs, you use the [l3-interface](#) statement.

Because RVIs operate at Layer 3, you can use Layer 3 services such as firewall filters or CoS rewriting with them. RVIs are similar to integrated routing and bridging (IRB) interfaces supported on Juniper routers and switch virtual interfaces (SVIs) and bridge-group virtual interfaces (BVI) supported on other vendors' devices.

[Table 180 on page 1876](#) shows the number of RVIs that each QFX platform supports.

Table 180: Number of Supported RVIs by Platform

Platform	Number of Supported RVIs
QFX3500	1200
QFX3000-G	2000
QFX3000-M	2000

Related Documentation

- [Example: Configuring Routing Between VLANs on One Switch on page 1976](#)

Understanding MAC Learning

MAC learning is the process of obtaining the MAC addresses of all the nodes on a network.

When a node is first connected to an Ethernet LAN or VLAN, it has no information about the other nodes on the network. As data is sent through the network, data packets include a data frame listing their source and destination MAC addresses. The data frame is forwarded to a target port, which is connected to the second device. The MAC address is learned locally at the target port, which facilitates communications for frames that later enter the target port and contain addresses previously learned from a received frame.

MAC learning can also be enabled on a per-VLAN basis. See [“Example: Disabling MAC Learning in a VLAN” on page 2003](#) for further information.

By default, MAC learning is enabled on the QFX Series.

Related Documentation

- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 1871](#)
- [Overview of Layer 2 Networking on page 1867](#)

Understanding Reflective Relay for Use with VEPA Technology

Virtual Ethernet Port Aggregator (VEPA) technology aggregates packets generated by virtual machines located on the same server and relays them to a physical switch. The physical switch then provides connectivity between the virtual machines located on the server, so the virtual machines do not communicate with one another. Offloading switching activities from a virtual switch to a physical switch reduces the computing overhead on the virtual servers and takes advantage of the security, filtering, and management features of the physical switch. Reflective relay, also known as “hairpin turn,” enables the physical switch to receive aggregated packets from the virtual machines hosted on the server through the VEPA on the downstream port and send those packets out the same downstream port from which the physical switch received them.

- [VEPA on page 1877](#)
- [Reflective Relay on page 1877](#)

VEPA

Even though virtual machines are capable of sending packets directly to one another, it is more efficient to pass these aggregated packets from the VEPA to a physical switch. The switch can then send any packets destined for a virtual machine located on the same server to the VEPA.

Reflective Relay

Reflective relay, also known as a “hairpin turn” or “hairpin mode,” returns aggregated packets to the VEPA by using the same downstream port that initially delivered the aggregated packets from the VEPA to the switch. Reflective relay must be configured on the interface located on the physical switch that receives aggregated packets, such as VEPA packets, because some of these packets might need to be sent back to the server if they are destined for another virtual machine on the same server.

Reflective relay only occurs in two situations:

- When the destination address of the packet was learned on that downstream port
- When the destination has not yet been learned

Reflective relay does not otherwise change the operation of the switch. If the interface to which the virtual machine is connected and the MAC address of the virtual machine packet are not yet included in the Ethernet switching table for the virtual machine's associated VLAN, an entry is added. If the source MAC address of an incoming packet under the respective VLAN is not yet present in the Ethernet switching table, the switch floods the packet on all the other ports that are members of the same VLAN, including the port on which the packet arrived.

Related Documentation

- [Understanding Bridging on page 1869](#)
- [Understanding VLANs on page 1873](#)
- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958](#)

Understanding Private VLANs

VLANs limit broadcasts to specified users. Private VLANs (PVLANS) take this concept a step further by splitting the broadcast domain into multiple isolated broadcast subdomains and essentially putting secondary VLANs inside a primary VLAN. PVLANS restrict traffic flows through their member switch ports (called “private ports”) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network. Each PVLAN typically contains many private ports that communicate only with a single uplink, thereby preventing the ports from communicating with each other.

Just like regular VLANs, PVLANS are isolated on Layer 2 and require that a Layer 3 device be used to route traffic among them. PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts. Service providers use PVLANS to keep their customers isolated from one another.


This topic explains the following concepts regarding PVLANS on the QFX Series:

- [Typical Structure and Primary Application of PVLANS on page 1878](#)
- [Using 802.1Q Tags to Identify Packets on page 1880](#)
- [Efficient Use of IP Addresses on page 1881](#)
- [PVLAN Port Types on page 1881](#)
- [Limitations of Private VLANs on page 1883](#)

Typical Structure and Primary Application of PVLANS

A PVLAN can be created on a single switch or can be configured to span multiple switches. The PVLAN shown in [Figure 41 on page 1879](#) includes two switches, with a primary PVLAN domain and various subdomains.

Figure 41: Subdomains in a PVLAN



As shown in [Figure 41 on page 1879](#), a PVLAN has only one primary domain and multiple secondary domains. The types of domains are:


- Primary VLAN—VLAN used to forward frames downstream to isolated and community VLANs.
- Secondary isolated VLAN—VLAN that receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.
- Secondary interswitch isolated VLAN—VLAN used to forward isolated VLAN traffic from one switch to another through PVLAN trunk ports. 802.1Q tags are required for interswitch isolated VLANs because IEEE 802.1Q uses an internal tagging mechanism

by which a trunking device inserts a 4-byte VLAN frame identification tab into the packet header.

- Secondary community VLAN—VLAN used to transport frames among members of a community (a subset of users within the VLAN) and to forward frames upstream to the primary VLAN.

Figure 42 on page 1880 shows a PVLAN spanning multiple switches, where the primary VLAN (100) contains two community domains (300 and 400) and one interswitch isolated domain.

Figure 42: PVLAN Spanning Multiple Switches



NOTE: Primary and secondary VLANs count against the limit of 4089 VLANs supported on the QFX Series. For example, each VLAN in Figure 42 on page 1880 counts against this limit.

Using 802.1Q Tags to Identify Packets

When packets are marked with a customer-specific 802.1Q tag, that tag identifies ownership of the packets for any switch or router in the network. Sometimes, 802.1Q tags are needed within PVLANS to keep track of packets from different subdomains. Table 181 on page 1881 indicates when a VLAN 802.1Q tag is needed on the primary VLAN or on secondary VLANs.

Table 181: PVLAN Requirements for 802.1Q Tags

	On a Single Switch	On Multiple Switches
Primary VLAN	Specify an 802.1Q tag by setting a VLAN ID.	Specify an 802.1Q tag by setting a VLAN ID.
Secondary VLAN	No tag needed on VLANs.	VLANs need 802.1Q tags: <ul style="list-style-type: none"> Specify an 802.1Q tag for each community VLAN by setting a VLAN ID. Specify the 802.1Q tag for an isolation VLAN ID by setting an isolation ID.

Efficient Use of IP Addresses

PVLANS provide IP address conservation and efficient allocation of IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In PVLANS, the hosts in all secondary VLANs belong to the same IP subnet because the subnet is allocated to the primary VLAN. Hosts within the secondary VLAN are assigned IP addresses based on IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet. However, each secondary VLAN is a separate broadcast domain.

PVLAN Port Types

PVLANS can use six different port types. The network depicted in [Figure 42 on page 1880](#) uses a promiscuous port to transport information to the router, community ports to connect the finance and HR communities to their respective switches, isolated ports to connect the servers, and a PVLAN trunk port to connect the two switches. PVLAN ports have different restrictions:

- Promiscuous trunk port—A promiscuous port is an upstream trunk port connected to a router, firewall, server, or provider network. A promiscuous trunk port can communicate with all interfaces, including the isolated and community ports within a PVLAN.
- PVLAN trunk port—A PVLAN trunk port is required in multiswitch PVLAN configurations to span the switches. The PVLAN trunk port is a member of all VLANs within the PVLAN (that is, the primary VLAN, the community VLANs, and the interswitch isolated VLAN), and it carries traffic from the primary VLAN and all secondary VLANs. It can communicate with all ports.

Communication between a PVLAN trunk port and an isolated port is usually unidirectional. A PVLAN trunk port's membership in the interswitch isolated VLAN is egress-only, meaning that an isolated port can forward packets to a PVLAN trunk port, but a PVLAN trunk port does not forward packets to an isolated port (unless the packets ingress on a promiscuous access port and are therefore being forwarded to all the secondary VLANs in the same primary VLAN as the promiscuous port).

- Secondary VLAN trunk port (not shown)—Secondary trunk ports carry secondary VLAN traffic. For a given private VLAN, a secondary VLAN trunk port can carry traffic for only one secondary VLAN. However, a secondary VLAN trunk port can carry traffic for multiple secondary VLANs as long as each secondary VLAN is a member of a different

primary VLAN. For example, a secondary VLAN trunk port can carry traffic for a community VLAN that is part of primary VLAN pvlan100 and also carry traffic for an isolated VLAN that is part of primary VLAN pvlan400.

- **Community port**—Community ports communicate among themselves and with their promiscuous ports. Community ports serve only a select group of users. These interfaces are separated at Layer 2 from all other interfaces in other communities or isolated ports within their PVLAN.
- **Isolated access port**—Isolated ports have Layer 2 connectivity only with promiscuous ports and PVLAN trunk ports—an isolated port cannot communicate with another isolated port even if these two ports are members of the same isolated VLAN (or interswitch isolated VLAN) domain. Typically, a server, such as a mail server or a backup server, is connected on an isolated port. In a hotel, each room would typically be connected on an isolated port, meaning that room-to-room communication is not possible, but each room can access the Internet on the promiscuous port.
- **Promiscuous access port (not shown)**—These ports carry untagged traffic. Traffic that ingresses on a promiscuous access port is forwarded to all secondary VLAN ports on the device. If traffic ingresses into the device on a VLAN-enabled port and egresses on a promiscuous access port, the traffic is untagged on egress. If tagged traffic ingresses on a promiscuous access port, the traffic is discarded.

Table 182 on page 1882 summarizes whether Layer 2 connectivity exists between the different types of ports.

Table 182: PVLAN Ports and Layer 2 Connectivity

Port Type	Promiscuous Trunk	PVLAN Trunk	Secondary Trunk	Community	Isolated Access	Promiscuous access
Promiscuous trunk	Yes	Yes	Yes	Yes	Yes	Yes
PVLAN trunk	Yes	Yes	Yes	Yes—same community only	Yes	Yes
Secondary Trunk	Yes	Yes	No	Yes	No	Yes
Community	Yes	Yes	Yes	Yes—same community only	No	Yes
Isolated access	Yes	Yes—unidirectional only	No	No	No	Yes
Promiscuous access	Yes	Yes	Yes	Yes	Yes	No



NOTE: If you enable the `no-mac-learning` statement on a primary VLAN, all isolated VLANs in the PVLAN inherit that setting. However, if you want to disable MAC address learning on any community VLANs, you must configure the `no-mac-learning` statement on each of those VLANs.

Limitations of Private VLANs

The following constraints apply to private VLAN configurations:

- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.

Related Documentation

- [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)

Understanding PVLAN Traffic Flows Across Multiple Switches

This topic illustrates and explains three different traffic flows on a sample multiswitch network configured with a private VLAN (PVLAN). PVLANS restrict traffic flows through their member switch ports (which are called “private ports”) so that they communicate only with a specific uplink trunk port or with specified ports within the same VLAN.

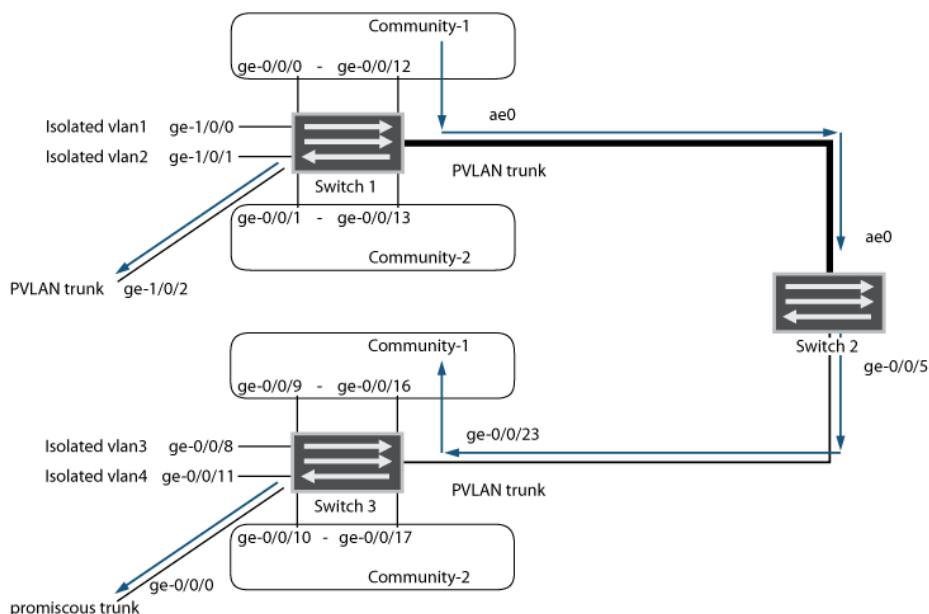
This topic describes:

- [Community VLAN Sending Untagged Traffic on page 1883](#)
- [Isolated VLAN Sending Untagged Traffic on page 1884](#)
- [PVLAN Tagged Traffic Sent on a Promiscuous Port on page 1885](#)

Community VLAN Sending Untagged Traffic

In this scenario, a VLAN in Community-1 of Switch 1 at interface ge-0/0/0 sends untagged traffic. The arrows in [Figure 43 on page 1884](#) represent this traffic flow.

Figure 43: Community VLAN Sends Untagged Traffic



g041099

In this scenario, the following activity takes place on Switch 1:

- Community-1 VLAN on interface ge-0/0/0: Learning
- pvlan100 on interface ge-0/0/0: Replication
- Community-1 VLAN on interface ge-0/0/12: Receives traffic
- PVLAN trunk port: Traffic exits from ge-1/0/2 and from ae0 with tag 10
- Community-2: Interface receives no traffic
- Isolated VLANs: Interfaces receive no traffic


In this scenario, this activity takes place on Switch 3:

- Community-1 VLAN on interface ge-0/0/23 (PVLAN trunk): Learning
- pvlan100 on interface ge-0/0/23: Replication
- Community-1 VLAN on interface ge-0/0/9 and ge-0/0/16: Receives traffic
- Promiscuous trunk port: Traffic exits from ge-0/0/0 with tag 100
- Community-2: Interface receives no traffic
- Isolated VLANs: Interfaces receive no traffic

Isolated VLAN Sending Untagged Traffic

In this scenario, isolated VLAN1 on Switch 1 at interface ge-1/0/0 sends untagged traffic. The arrows in Figure 44 on page 1885 represent this traffic flow.

Figure 44: Isolated VLAN Sends Untagged Traffic



In this scenario, the following activity takes place on Switch 1:

- Isolated VLAN1 on interface ge-1/0/0: Learning
- pvlan100 on interface ge-1/0/0: Replication
- Traffic exits from pvlan-trunk ge-1/0/2 and ae0 with tag 50
- Community-1 and Community-2: Interfaces receive no traffic
- Isolated VLANs: Interfaces receive no traffic

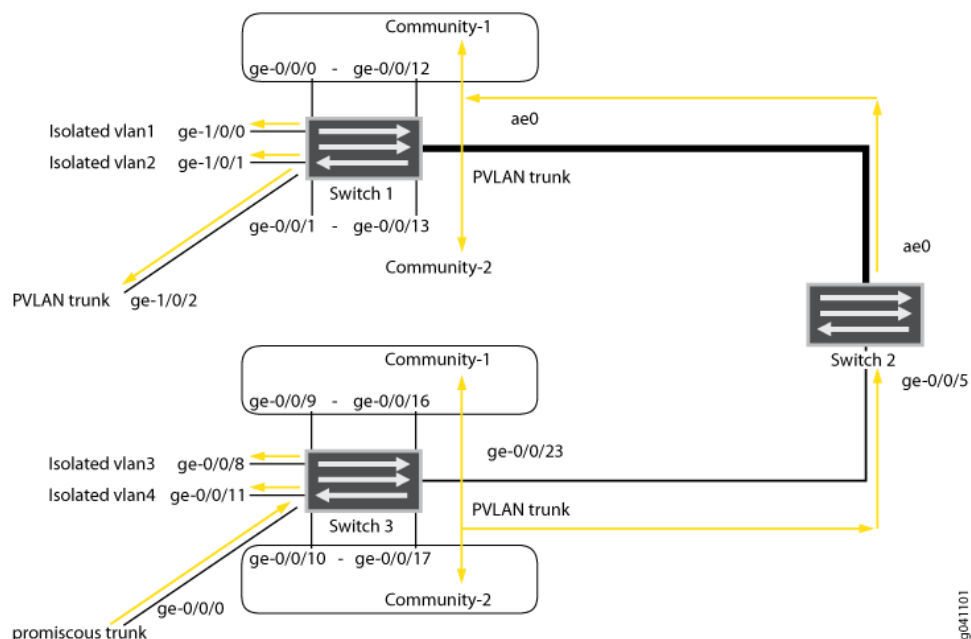
In this scenario, this activity takes place on Switch 3:

- VLAN on interface ge-0/0/23 (PVLAN trunk port): Learning
- pvlan100 on interface ge0/0/23: Replication
- Promiscuous trunk port: Traffic exits from ge-0/0/0 with tag 100
- Community-1 and Community-2: Interfaces receive no traffic
- Isolated VLANs: Receive no traffic

PVLAN Tagged Traffic Sent on a Promiscuous Port

In this scenario, PVLAN tagged traffic is sent on a promiscuous port. The arrows in [Figure 45 on page 1886](#) represent this traffic flow.

Figure 45: PVLAN Tagged Traffic Sent on a Promiscuous Port



In this scenario, the following activity takes place on Switch 1:

- pvlan100 VLAN on interface ae0 (PVLAN trunk): Learning
- Community-1, Community-2, and all isolated VLANs on interface ae0: Replication
- VLAN on interface ae0: Replication
- Traffic exits from pvlan-trunk ge-1/0/2 with tag 100
- Community-1 and Community-2: Interfaces receive traffic
- Isolated VLANs: Receive traffic

In this scenario, this activity takes place on Switch 3:

- pvlan100 on interface ge-0/0/0: Learning
- Community-1, Community-2 and all isolated VLANs on interface ge-0/0/0: Replication
- VLAN on interface ge-0/0/0: Replication
- Community-1 and Community-2: Interfaces receive traffic
- Isolated VLANs: Receive traffic

Related Documentation

- *Understanding Private VLANs on EX Series Switches*
- *Example: Configuring a Private VLAN on a Single EX Series Switch*
- *Example: Configuring a Private VLAN Spanning Multiple EX Series Switches*
- *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)*
- *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*

- [Understanding Private VLANs on page 1878](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Example: Configuring a Private VLAN on a Single Switch on page 2005](#)
- [Example: Configuring a Private VLAN Spanning Multiple Switches on page 2010](#)

Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS

VLANs limit broadcasts to specified users. Private VLANs (PVLANS) take this concept a step further by splitting a VLAN into multiple broadcast subdomains and essentially putting secondary VLANs inside a primary VLAN. PVLANS restrict traffic flows through their member ports so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network. A PVLAN typically contains many private ports that communicate only with a single uplink, thereby preventing the ports from communicating with each other.

Secondary trunk ports and promiscuous access ports extend the functionality of PVLANS for use in complex deployments, such as:

- Enterprise VMWare Infrastructure environments
- Multitenant cloud services with VM management
- Web hosting services for multiple customers

For example, you can use secondary VLAN trunk ports to connect QFX devices to VMware servers that are configured with private VLANs. You can use promiscuous access ports to connect QFX devices to systems that do not support trunk ports but do need to participate in private VLANs.

This topic explains the following concepts regarding PVLANS on the QFX Series:

- [PVLAN Port Types on page 1887](#)
- [Secondary VLAN Trunk Port Details on page 1888](#)
- [Use Cases on page 1889](#)

PVLAN Port Types

PVLANS can use the following different port types:

- Promiscuous trunk port—A promiscuous port is an upstream trunk port connected to a router, firewall, server, or provider network. A promiscuous trunk port can communicate with all interfaces, including the isolated and community ports within a PVLAN.
- PVLAN trunk port—A PVLAN trunk port is required in multiswitch PVLAN configurations to span the switches. The PVLAN trunk port is a member of all VLANs within the PVLAN (that is, the primary VLAN, the community VLANs, and the interswitch isolated VLAN), and it carries traffic from the primary VLAN and all secondary VLANs. It can communicate with all ports.

Communication between a PVLAN trunk port and an isolated port is usually unidirectional. A PVLAN trunk port's membership in the interswitch isolated VLAN is egress-only, meaning that an isolated port can forward packets to a PVLAN trunk port, but a PVLAN trunk port does not forward packets to an isolated port (unless the packets ingressed on a promiscuous access port and are therefore being forwarded to all the secondary VLANs in the same primary VLAN as the promiscuous port).

- Secondary VLAN trunk port—Secondary VLAN trunk ports carry secondary VLAN traffic. For a given private (primary) VLAN, a secondary VLAN trunk port can carry traffic for only one secondary VLAN. However, a secondary VLAN trunk port can carry traffic for multiple secondary VLANs as long as each secondary VLAN is a member of a different primary VLAN. For example, a secondary VLAN trunk port can carry traffic for a community VLAN that is part of primary VLAN pvlan100 and also carry traffic for an isolated VLAN that is part of primary VLAN pvlan400.



NOTE: When traffic egresses from a secondary VLAN trunk port, it normally carries the tag of the primary VLAN that the secondary port is a member of. If you want traffic that egresses from a secondary VLAN trunk port to retain its secondary VLAN tag, use the [extend-secondary-vlan-id](#) statement.

- Community port—Community ports communicate among themselves and with their promiscuous ports. Community ports serve only a select group of users. These interfaces are separated at Layer 2 from all other interfaces in other communities or isolated ports within their PVLAN.
- Isolated access port—Isolated ports have Layer 2 connectivity only with promiscuous ports and PVLAN trunk ports. An isolated access port cannot communicate with another isolated port even if these two ports are members of the same isolated VLAN.
- Promiscuous access port—These ports carry untagged traffic and can be a member of only one primary VLAN. Traffic that ingresses on a promiscuous access port is forwarded to the ports of the secondary VLANs that are members of the primary VLAN that the promiscuous access port is a member of. In this case, the traffic carries the appropriate secondary VLAN tag when it egresses from the secondary VLAN port if the secondary VLAN port is a trunk port. If traffic ingresses on a secondary VLAN port and egresses on a promiscuous access port, the traffic is untagged on egress. If tagged traffic ingresses on a promiscuous access port, the traffic is discarded.

Secondary VLAN Trunk Port Details

When using a secondary VLAN trunk port, be aware of the following:

- You must configure an isolation VLAN ID for each primary VLAN that the secondary VLAN trunk port will participate in. This is true even if the secondary VLANs that the secondary VLAN trunk port will carry are confined to a single device.
- If you configure a port to be a secondary VLAN trunk port for a given primary VLAN, you can also configure the same physical port to be any of the following:

- Secondary VLAN trunk port for another primary VLAN
 - PVLAN trunk for another primary VLAN
 - Promiscuous trunk port
 - Access port for a non-private VLAN
- Traffic that ingresses on a secondary VLAN trunk port (with a secondary VLAN tag) and egresses on a PVLAN trunk port retains the secondary VLAN tag on egress.
 - Traffic that ingresses on a secondary VLAN trunk port and egresses on a promiscuous trunk port has the appropriate primary VLAN tag on egress.
 - Traffic that ingresses on a secondary VLAN trunk port and egresses on a promiscuous access port is untagged on egress.
 - Traffic that ingresses on a promiscuous trunk port with a primary VLAN tag and egresses on a secondary VLAN trunk port carries the appropriate secondary VLAN tag on egress. For example, assume that you have configured the following on a switch:
 - Primary VLAN 100
 - Community VLAN 200 as part of the primary VLAN
 - Promiscuous trunk port
 - Secondary trunk port that carries community VLAN 200

If a packet ingresses on the promiscuous trunk port with primary VLAN tag 100 and egresses on the secondary VLAN trunk port, it carries tag 200 on egress.

Use Cases

On the same physical interface, you can configure multiple secondary VLAN trunk ports (in different primary VLANs) or combine a secondary VLAN trunk port with other types of VLAN ports. The following use cases provide examples of doing this and show how traffic would flow in each case:

- [Secondary VLAN Trunks In Two Primary VLANs on page 1889](#)
- [Secondary VLAN Trunk and Promiscuous Trunk on page 1891](#)
- [Secondary VLAN Trunk and PVLAN Trunk on page 1892](#)
- [Secondary VLAN Trunk and Non-Private VLAN Interface on page 1894](#)
- [Traffic Ingressing on Promiscuous Access Port on page 1895](#)

Secondary VLAN Trunks In Two Primary VLANs

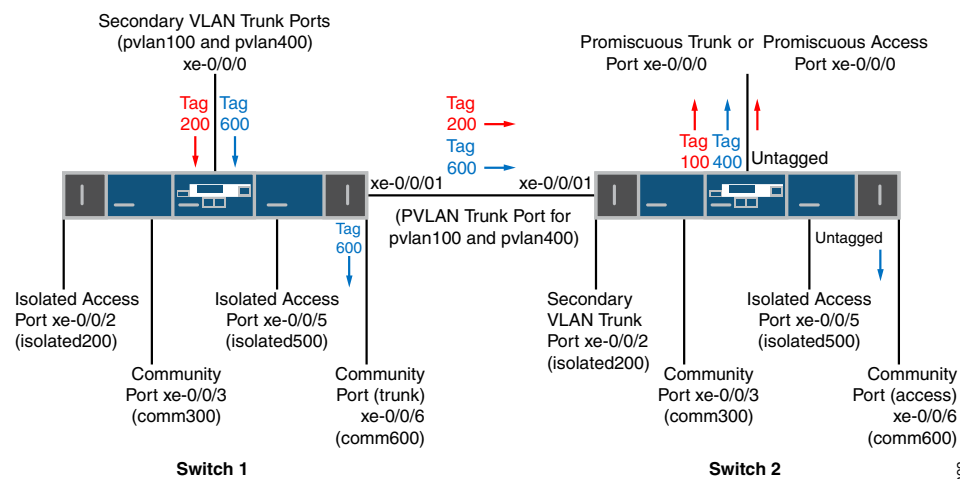
For this use case, assume you have two switches with the following configuration:

- Primary VLAN pvlan100 with tag 100.
 - Isolated VLAN isolated200 with tag 200 is a member of pvlan100.
 - Community VLAN comm300 with tag 300 is a member of pvlan100.
- Primary VLAN pvlan400 with tag 400.

- Isolated VLAN isolated500 with tag 500 is a member of pvlan400.
- Community VLAN comm600 with tag 600 is a member of pvlan400.
- Interface xe-0/0/0 on Switch 1 connects to a VMware server (not shown) that is configured with the private VLANs used in this example. This interface is configured with secondary VLAN trunk ports to carry traffic for secondary VLAN comm600 and the isolated VLAN (tag 200) that is a member of pvlan100.
- Interface xe-0/0/0 on Switch 2 is shown configured as a promiscuous trunk port or promiscuous access port. In the latter case, you can assume that it connects to a system (not shown) that does not support trunk ports but is configured with the private VLANs used in this example.
- On Switch 1, xe-0/0/6 is a member of comm600 and is configured as a trunk port.
- On Switch 2, xe-0/0/6 is a member of comm600 and is configured as an access port.

Figure 46 on page 1890 shows this topology and how traffic for isolated200 and comm600 would flow after ingressing on xe-0/0/0 on Switch 1. Note that traffic would flow only where the arrows indicate. For example, there are no arrows for interfaces xe-0/0/2, xe-0/0/3, and xe-0/0/5 on Switch 1 because no packets would egress on those interfaces.

Figure 46: Two Secondary VLAN Trunk Ports on One Interface



Here is the traffic flow for VLAN isolated200:

1. After traffic for isolated200 ingresses on the secondary VLAN trunk port on Switch 1, it egresses on the PVLAN trunk port because the PVLAN trunk port is a member of all the VLANs. The packets keep the secondary VLAN tag (200) when egressing.
2. After traffic for isolated200 ingresses on the secondary VLAN trunk port on Switch 2, it egresses on xe-0/0/0, which is configured as a promiscuous trunk port or promiscuous access port.
 - If xe-0/0/0 on Switch 2 is configured as a promiscuous trunk port, the packets egress on this port with the primary VLAN tag (100).

- If xe-0/0/0 on Switch 2 is configured as a promiscuous access port, the packets egress on this port untagged.

Note that traffic for VLAN isolated200 does not egress on isolated access port xe-0/0/2 on Switch 1 or secondary VLAN trunk port xe-0/0/2 on Switch 2 even though these two ports are members of the same isolated VLAN.

Here is the traffic flow for VLAN comm600:

1. After traffic for comm600 ingresses on the secondary VLAN trunk port on Switch 1, it egresses on the PVLAN trunk port because the PVLAN trunk port is a member of all the VLANs. The packets keep the secondary VLAN tag (600) when egressing.
2. Traffic for comm600 also egresses on community port xe-0/0/6 on Switch 1. The traffic is tagged because the port is configured as a trunk.
3. After traffic for comm600 ingresses on the PVLAN trunk port on Switch 2, it egresses on xe-0/0/0, if this interface is configured as a promiscuous trunk port.



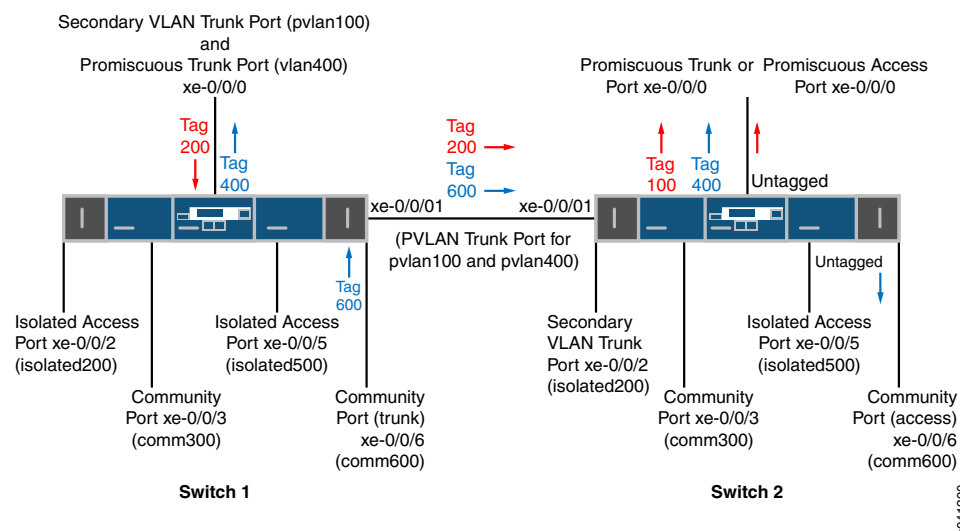
NOTE: If xe-0/0/0 on Switch 2 is configured as a promiscuous access port, the port can participate in only one primary VLAN. In this case, the promiscuous access port is part of pvlan100, so traffic for comm600 does not egress from it

4. Traffic for comm600 also egresses on community port xe-0/0/6 on Switch 2. In this case, the traffic is untagged because the port mode is access.

Secondary VLAN Trunk and Promiscuous Trunk

For this use case, assume you have two switches configured with the same ports and VLANs as in the previous use case, with one exception: In this case, xe-0/0/0 on Switch 1 is configured as a secondary VLAN trunk port for VLAN pvlan100 and is also configured as a promiscuous trunk port for pvlan400.

Figure 47 on page 1892 shows this topology and how traffic for isolated200 (member of pvlan100) and comm600 (member of pvlan400) would flow after ingressing on Switch 1.

Figure 47: Secondary VLAN Trunk and Promiscuous Trunk on One Interface

The traffic flow for VLAN isolated200 is the same as in the previous use case, but the flow for comm600 is different. Here is the traffic flow for VLAN comm600:

1. After traffic for comm600 ingresses on community VLAN port xe-0/0/6 on Switch 1, it egresses on promiscuous trunk port xe-0/0/0 on Switch 1. In this case it carries the primary VLAN tag (400).
2. Traffic for comm600 also egresses on the PVLAN trunk port because the PVLAN trunk port is a member of all the VLANs. The packets keep the secondary VLAN tag (600) when egressing.

3. After traffic for comm600 ingresses on the PVLAN trunk port on Switch 2, it egresses on xe-0/0/0, if this interface is configured as a promiscuous trunk port.

It does not egress on xe-0/0/0 if this interface is configured as a promiscuous access port because the port can participate only in pvlan100.


4. Traffic for comm600 also egresses on community port xe-0/0/6 on Switch 2.

Secondary VLAN Trunk and PVLAN Trunk

For this use case, assume you have two switches configured with the same ports and VLANs as in the previous use cases except that xe-0/0/0 on Switch 1 is configured as a secondary VLAN trunk port for VLAN pvlan100 and is also configured as a PVLAN trunk port for pvlan400.

Figure 48 on page 1893 shows this topology and how traffic for comm300 (member of pvlan100) and comm600 (member of pvlan400) would flow after ingressing on Switch 1.

Figure 48: Secondary VLAN Trunk and PVLAN Trunk on One Interface



Here is the traffic flow for VLAN comm300:

1. After traffic for comm300 ingresses on community port xe-0/0/3 on Switch 1, it egresses on PVLAN trunk port xe-0/0/1 because that PVLAN trunk port is a member of all the VLANs. The packets keep the secondary VLAN tag (300) when egressing.



NOTE: Traffic for comm300 does not egress on xe-0/0/0 because the secondary VLAN trunk port on this interface carries isolated200, not comm300.

2. After traffic for comm300 ingresses on the PVLAN trunk port on Switch 2, it egresses on xe-0/0/0, which is configured as a promiscuous trunk port or promiscuous access port.
 - If xe-0/0/0 on Switch 2 is configured as a promiscuous trunk port, the packets egress on this port with the primary VLAN tag (100).
 - If xe-0/0/0 on Switch 2 is configured as a promiscuous access port, the packets egress on this port untagged.
3. Traffic for comm300 also egresses on community port xe-0/0/3 on Switch 2.

Here is the traffic flow for VLAN comm600:

1. After traffic for comm600 ingresses on the PVLAN port xe-0/0/0 on Switch 1, it egresses on the community port xe-0/0/6 on Switch 1. The packets keep the secondary VLAN tag (600) when egressing because xe-0/0/6 is a trunk port.
2. Traffic for comm600 also egresses on PVLAN trunk port xe-0/0/1 because that PVLAN trunk port is a member of all the VLANs. The packets keep the secondary VLAN tag (600) when egressing.

- After traffic for comm600 ingresses on the PVLAN trunk port on Switch 2, it egresses on xe-0/0/0, if this interface is configured as a promiscuous trunk port.

It does not egress on xe-0/0/0 if this interface is configured as a promiscuous access port because the port can participate only in pvlan100.

- Traffic for comm600 also egresses on community port xe-0/0/6 on Switch 2. This traffic is untagged on egress because xe-0/0/6 is an access port.

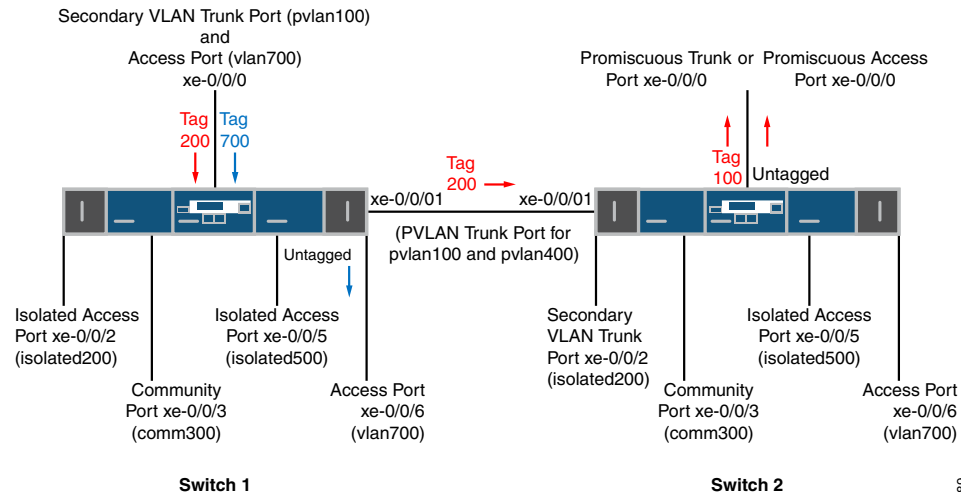
Secondary VLAN Trunk and Non-Private VLAN Interface

For this use case, assume you have two switches configured with the same ports and VLANs as in the previous use cases except for these differences:

- Configuration for xe-0/0/0 on Switch 1:
 - Secondary VLAN trunk port for VLAN pvlan100
 - Access port for vlan700
- Port xe-0/0/6 on both switches is an access port for vlan700.

Figure 49 on page 1894 shows this topology and how traffic for isolated200 (member of pvlan100) and vlan700 would flow after ingressing on Switch 1.

Figure 49: Secondary VLAN Trunk and Non-Private VLAN Port on One Interface



Here is the traffic flow for VLAN isolated200:

- After traffic for isolated200 ingresses on the secondary VLAN trunk port on Switch 1, it egresses on the PVLAN trunk port. The packets keep the secondary VLAN tag (200) when egressing.
- After traffic for isolated200 ingresses on the PVLAN trunk port on Switch 2, it egresses on xe-0/0/0, which is configured as a promiscuous trunk port or promiscuous access port.

- If xe-0/0/0 on Switch 2 is configured as a promiscuous trunk port, the packets egress on this port with the primary VLAN tag (100).
- If xe-0/0/0 on Switch 2 is configured as a promiscuous access port, the packets egress on this port untagged.

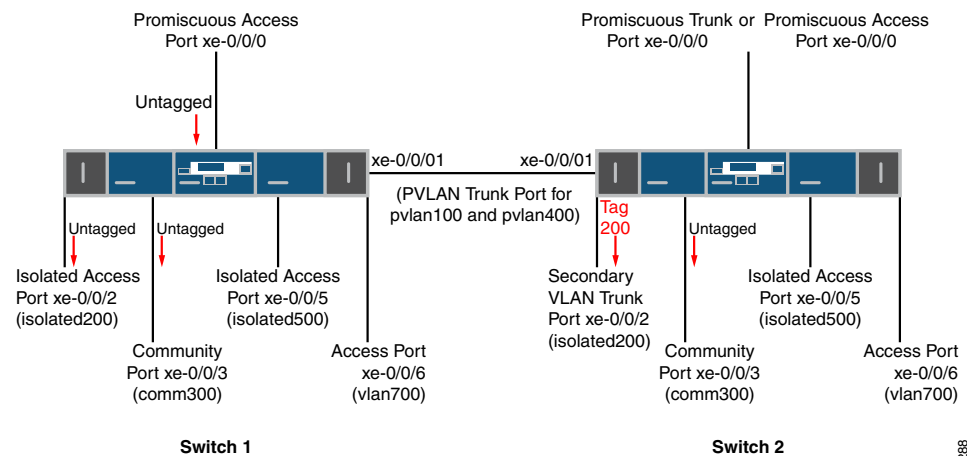
Note that traffic for VLAN isolated200 does not egress on isolated access port xe-0/0/2 on Switch 1 or secondary VLAN trunk port xe-0/0/2 on Switch 2 even though these two ports are members of the same isolated VLAN.

After traffic for vlan700 ingresses on the access port configured on xe-0/0/0 on Switch 1, it egresses on access port xe-0/0/6 because that port is a member of the same VLAN. Traffic for vlan700 is not forwarded to Switch 2 (even though xe-0/0/6 on Switch 2 is a member of vlan700) because the PVLAN trunk on xe-0/0/1 does not carry this VLAN.

Traffic Ingressing on Promiscuous Access Port

For this use case, assume you have two switches configured with the same ports and VLANs as in the previous use case except that xe-0/0/0 on Switch 1 is configured as a promiscuous access port and is a member of pvlan100. [Figure 50 on page 1895](#) shows this topology and how untagged traffic would flow after ingressing through this interface on Switch 1.

Figure 50: Traffic Ingressing on Promiscuous Access Port



g041288

As the figure shows, untagged traffic that ingresses on a promiscuous access port is forwarded to all the secondary VLAN ports that are members of the same primary VLAN that the promiscuous access port is a member of. The traffic is untagged when it egresses from access ports and tagged on egress from a trunk port (xe-0/0/2 on Switch 2).

Related Documentation

- [Understanding Private VLANs on page 1878](#)
- [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)

- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Understanding Egress Firewall Filters with PVLANS on page 1896](#)
- [Troubleshooting Private VLANs on page 2263](#)

Understanding Egress Firewall Filters with PVLANS

If you apply firewall filters to private VLANs in the output direction, the behavior of the filters might be unexpected. This topic explains how egress filters behave when applied to private VLANs.

If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

Related Documentation

- [Understanding Private VLANs on page 1878](#)
- [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)

- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Troubleshooting Private VLANs on page 2263](#)

Understanding Multiple VLAN Registration Protocol (MVRP)

Multiple VLAN Registration Protocol (MVRP) is a Layer 2 messaging protocol that automates the creation and management of virtual LANs, thereby reducing the time you have to spend on these tasks. If your QFabric system connects to servers that host many virtual machines that require their own VLANs, using MVRP can save you the time and effort that would be required to manually create and administer the VLANs on the ports that connect to the servers. For example, if a virtual machine moves between servers—and therefore connects to a different redundant server Node group interface—MVRP can configure the appropriate VLAN membership on the new server Node group interface.

When using MVRP on a QFabric system, you must manually create on the QFabric the VLANs that exist on the attached servers because the QFabric implementation of MVRP does not allow VLANs to be created dynamically. However, you do not need to manually assign VLAN membership to the QFabric ports that connect to the servers. MVRP automatically assigns VLAN membership to server-facing QFabric ports when it learns about a VLAN from an attached server.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP replace Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) and overcome GARP and GVRP limitations.



NOTE: MVRP on QFabric systems does not support private VLANs.

- [QFabric Requirements on page 1897](#)
- [MVRP Operations on page 1898](#)
- [MRP Timers Control MVRP Updates on page 1898](#)
- [MVRP Uses MRP Messages to Transmit Switch and VLAN States on page 1899](#)

QFabric Requirements

When configuring MVRP on a QFabric system, you can enable it globally or enable it only on the trunk ports that need to carry VLAN traffic from the attached servers. You also must manually create the expected VLANs, but you do not have to assign VLAN membership to the server-facing redundant server Node ports (as mentioned previously). However, you *do* have to manually assign VLAN membership to the uplink ports on the redundant server Node group and network Node group devices that will carry the VLAN traffic. [Table 183 on page 1897](#) summarizes the VLAN requirements for redundant server Node groups and network Node groups:

Table 183: MVRP VLAN Requirements for Node Devices

Node Group Type	Interface	Assign VLAN Membership to Trunk Ports?
-----------------	-----------	--

Table 183: MVRP VLAN Requirements for Node Devices (*continued*)

Redundant server Node group	Server-facing trunk	No
Redundant server Node group	Uplink trunk (to interconnect device)	Yes
Network Node groups	Uplink trunk (to interconnect device)	Yes

MVRP Operations

MVRP stays synchronized by using MVRP protocol data units (PDUs). These PDUs specify which QFabric systems and switches are members of which VLANs, and which switch interfaces are in each VLAN. The MVRP PDUs are sent to other switches in the QFabric system when an MVRP state change occurs, and the receiving switches update their MVRP states accordingly. MVRP timers dictate when PDUs can be sent and when switches receiving MVRP PDUs can update their MVRP information.

In addition to this behavior, QFX switches include a mode—called passive mode—in which an MVRP-configured interface does not announce its membership in a VLAN or send any VLAN declarations (updates) unless it receives registration for that VLAN from a peer (server) on that interface. By default MVRP-configured interfaces behave in the standard manner and automatically send PDU updates to announce any VLAN changes. (This is called active mode.)

To enable passive mode on an interface, enter and commit this statement:

```
set protocols mvrp interface interface-name passive
```

To keep VLAN membership information current, MVRP removes switches and interfaces when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants, thereby reducing network overhead.
- Limits broadcast, unknown unicast, and multicast (BUM) traffic to interested devices.

MVRP is disabled by default and is valid only for trunk interfaces.

MRP Timers Control MVRP Updates

MVRP registration and updates are controlled by timers that are part of MRP. The timers define when MVRP PDUs can be sent and when MVRP information can be updated. You configure the timers on a per-interface basis.

The following MRP timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the leave state before changing to the unregistered state.
- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.



BEST PRACTICE: Unless there is a compelling reason to change the timer settings, leave the default settings in place. Modifying timers to inappropriate values can cause an imbalance in the operation of MVRP.

MVRP Uses MRP Messages to Transmit Switch and VLAN States

MVRP uses MRP messages to register and declare MVRP states for a interface or VLAN and to inform the switching network that a interface or VLAN is leaving MVRP. These messages are communicated in the MRP PDUs sent by MVRP-enabled interfaces.

The following MRP messages are communicated for MVRP:

- Empty—MVRP information is not declared and no VLAN is registered.
- In—MVRP information is not declared but a VLAN is registered.
- JoinEmpty—MVRP information is declared but no VLAN is registered.
- JoinIn—MVRP information is declared and a VLAN is registered.
- Leave—MVRP information that was previously declared is withdrawn.
- LeaveAll—Unregister all VLANs on the switch. VLANs must re-register to participate in MVRP.
- New—The MVRP information is new and a VLAN might not be registered yet.

Related Documentation

- [Understanding VLANs on page 1873](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on page 1953](#)
- [Configuring Multiple VLAN Registration Protocol on page 2051](#)

Spanning Trees Overview

- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding MSTP on page 1901](#)
- [Understanding RSTP on page 1901](#)
- [Understanding VSTP on page 1902](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1905](#)

Overview of Spanning-Tree Protocols

QFX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default spanning-tree protocol on the QFX Series is RSTP. RSTP provides faster convergence times than STP. However, some legacy networks require the slower convergence times of basic STP.

The STP support provided for the QFX Series includes:

- IEEE 802.1d
- 802.1w RSTP
- 802.1s MSTP

If your network includes IEEE 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. See [“Configuring STP” on page 2053](#). When you explicitly configure STP, the QFX Series products use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP. If you use virtual LANs (VLANs), you should enable VSTP and use it on your network. See [“Understanding VSTP” on page 1902](#).

You can use the same operational commands (**show spanning-tree bridge** and **show spanning-tree interface**) to check the status of your spanning-tree configuration, regardless of which spanning-tree protocol has been configured.

STP uses bridge protocol data unit (BPDU) packets to exchange information with other switches. BPDUs send hello packets out at regular intervals to exchange information across bridges and detect loops in a network topology. There are two types of BPDUs:

- Configuration BPDUs—These BPDUs contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.
- Topology change notification (TCN) BPDUs—When a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.

STP uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

Understanding Spanning Tree Protocols on a QFabric System

Although there is no need to run STP in a QFabric system, you can connect a QFabric system to another Layer 2 device and use STP. STP traffic can only be processed on network Node groups. Other Node groups, such as redundant server Node groups and server Node groups, discard the STP bridge protocol data units (BPDUs) traffic and disable the

interface automatically. Server Node groups only process host-facing protocols, whereas Network Node groups process all supported protocols.

Related Documentation

- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903](#)
- [Understanding MSTP on page 1901](#)
- [Understanding RSTP on page 1901](#)
- [Understanding VSTP on page 1902](#)

Understanding MSTP

Although RSTP provides faster convergence time than STP does, it still does not solve a problem inherent in STP: all VLANs within a LAN must share the same spanning tree. To solve this problem, the QFX Series products use Multiple Spanning Tree Protocol (MSTP) to create a loop-free topology in networks with multiple spanning-tree regions.

An MSTP region allows a group of bridges to be modeled as a single bridge. An MSTP region contains multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates more efficient load sharing across redundant links.

An MSTP region can support up to 64 MSTIs, and each instance can support from 1 through 4094 VLANs.

Related Documentation

- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding RSTP on page 1901](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)

Understanding RSTP

Juniper Networks QFX Series products use Rapid Spanning Tree Protocol (RSTP) on the network side of the QFX Series to provide quicker convergence time than the base Spanning Tree Protocol (STP) does. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state, which speeds up convergence.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. The STP process to determine network state transitions is slower than the RSTP process because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state. When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

For networks with virtual LANs (VLANs), you can use VLAN Spanning Tree Protocol (VSTP), which takes the paths of each VLAN into account when calculating routes. VSTP uses RSTP by default.

An RSTP domain running from the edge outward on a QFX Series product has the following components:

- A *root port*, which is the “best path” to the root device.
- A *designated port*, which indicates that the switch is the designated bridge for the other switch connecting to this port.
- An *alternate port*, which provides an alternate root port.
- A *backup port*, which provides an alternate designated port.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once per hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines that the connection with the neighbor is lost. When a *root port* or a *designated port* fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is valid for that port and starts a *synchronizing* operation to ensure that all of its ports are in sync with the new information.

Similar sets of messages propagate through the network, restoring the connectivity very quickly after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention.

RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on the QFX Series.

Related Documentation

- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding MSTP on page 1901](#)
- [Understanding VSTP on page 1902](#)
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)

Understanding VSTP

VLAN Spanning Tree Protocol (VSTP) enables Juniper Networks switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

You can configure VSTP for a maximum of 253 VLANs.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on a switch.



NOTE: We recommend that you enable VSTP on all VLANs that could receive VSTP bridge protocol data units (BPDUs).

Related Documentation

- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding RSTP on page 1901](#)
- [Configuring VLAN Spanning Tree Protocol on page 2054](#)
- [vstp on page 2147](#)

Understanding BPDU Protection for STP, RSTP, and MSTP



NOTE: You can disable BPDU protection on interfaces by issuing the `set ethernet-switching-options bpdu-block interface-name disable` command.

A Juniper Networks QFX Series product provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Bridge protocol data unit (BPDU) protection can help prevent STP misconfigurations that can lead to network outages.

A loop-free network is supported through the exchange of a special type of frame called a BPDU. Receipt of BPDUs on certain interfaces in an STP, RSTP, VSTP, or MSTP topology, however, can lead to network outages. Enable BPDU protection on those interfaces to prevent these outages.

Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a user bridge application running on a device connected to the switch can also generate BPDUs. If these BPDUs are picked up by STP applications running on the switch, they can trigger STP miscalculations, and those miscalculations can lead to network outages.

Enable BPDU protection on switch interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If BPDUs are received on a protected interface, the interface is disabled and stops forwarding frames.

Not only can you configure BPDU protection on a switch with a spanning tree, but also on a switch without a spanning tree. This type of topology typically consists of a non-STP switch connected to an STP switch through a trunk interface.

To configure BPDU protection on a switch with a spanning tree, include the **bpdu-block-on-edge** statement at the `[edit protocols (stp | mstp | rstp)]` hierarchy level. To configure BPDU protection on a switch without a spanning tree, include the **bpdu-block** statement at the `[edit ethernet-switching-options interface interface-name]` hierarchy level.

If BPDUs are sent to an interface (indicating that the misconfiguration has been corrected), the interface can be unblocked in one of two ways:

- If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires.
- Use the operational mode command **clear ethernet-switching bpdu-error**.

Disabling the BPDU protection configuration does not unblock the interface.

Related Documentation

- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1905](#)
- [Understanding MSTP on page 1901](#)
- [Understanding RSTP on page 1901](#)
- [Understanding VSTP on page 1902](#)

Understanding Loop Protection for STP, RSTP, VSTP, and MSTP

A Juniper Networks QFX Series product provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from entering a forwarding state that would cause a loop to open in the network.

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

However, a blocking interface can mistakenly transition to the forwarding state if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor.

When loop protection is enabled, the spanning-tree topology detects root ports and blocked ports and ensures that both keep receiving BPDUs. If a loop-protection-enabled interface stops receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. It does not transition the interface to a forwarding state, but instead transitions it to a loop-inconsistent state. The interface

recovers and it transitions back to the spanning-tree blocking state as soon as it receives a BPDU.

We recommend that you enable loop protection on all switch interfaces that have a chance of becoming root or designated ports. Loop protection is most effective when enabled in the entire switched network. When you enable loop protection, you must configure at least one action (**alarm**, **block**, or both).

An interface can be configured for either loop protection or root protection, but not for both.

Related Documentation

- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1905](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903](#)
- [Understanding MSTP on page 1901](#)
- [Understanding RSTP on page 1901](#)
- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding VSTP on page 1902](#)

Understanding Root Protection for STP, RSTP, VSTP, and MSTP

A Juniper Networks QFX Series product provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). A loop-free network is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

You can also see BPDUs generated when you run a bridge application on a device attached to the switch. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

Enable root protection on interfaces that should not receive higher-priority BPDUs from the root bridge and should not be elected as the root port. These interfaces become designated ports and are typically located on an administrative boundary. If the bridge receives more STP BPDUs on a port that has root protection enabled, that port transitions to a root-prevented STP state (inconsistency state), and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. After the bridge stops receiving more STP BPDUs on the interface with root protection, the interface returns to a listening state, followed by a learning state, and ultimately back to a forwarding state. Recovery back to the forwarding state is automatic.

When root protection is enabled on an interface, it is enabled for all the STP instances on that interface. The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.

An interface can be configured for either root protection or loop protection, but not for both.

**Related
Documentation**

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1994](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Understanding MSTP on page 1901](#)
- [Understanding RSTP on page 1901](#)
- [Overview of Spanning-Tree Protocols on page 1900](#)
- [Understanding VSTP on page 1902](#)

Q-in-Q Tunneling

- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)

Understanding Q-in-Q Tunneling and VLAN Translation

Q-in-Q tunneling and VLAN translation allow service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Data centers can use Q-in-Q tunneling and VLAN translation to isolate customer traffic within a single site or to enable customer traffic flows between cloud data centers in different geographic locations.

Q-in-Q tunneling adds a service VLAN tag before the customer's 802.1Q VLAN tags. The Juniper Networks Junos operating system implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

All of the VLANs in an implementation can be service VLANs. That is, if the total number of supported VLANs is 4090, all of them can be service VLANs.

This topic describes:

- [How Q-in-Q Tunneling Works on page 1907](#)
- [How VLAN Translation Works on page 1908](#)
- [Mapping C-VLANs to S-VLANs on page 1908](#)
- [Routed VLAN Interfaces on Q-in-Q VLANs on page 1909](#)
- [Constraints for Q-in-Q Tunneling and VLAN Translation on page 1909](#)

How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's or data center VLAN (S-VLAN), another 802.1Q tag for the appropriate S-VLAN is added before the C-VLAN tag. The C-VLAN tag remains and is transmitted through the network. As the packet leaves the S-VLAN in the downstream direction, the S-VLAN 802.1Q tag is removed.

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN or multiple C-VLANs to one S-VLAN. C-VLAN and S-VLAN tags use separate name spaces, so you can have both a C-VLAN 101 and an S-VLAN 101, for example. You can limit the set of accepted customer tags to a range of tags or to discrete values.

When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network. Access interfaces are assumed to be customer-facing and accept both tagged and untagged frames. When using many-to-one bundling or mapping a specific interface, you must use the **native** option to specify an S-VLAN for untagged and priority tagged packets if you want to accept these packets. (Priority tagged packets have their VLAN ID set to 0, and their priority code point bits might be configured with a CoS value.) If you do not specify an S-VLAN for them, untagged packets are discarded. The **native** option is not available for all-in-one bundling because there is no need to specify untagged and priority tagged packets when all packets are mapped to an S-VLAN.

On QFabric systems only, you can use the **native** option to apply a specified inner tag to packets that ingress as untagged on access interfaces. This functionality is useful if your QFabric system connects to servers that host customer virtual machines that send untagged traffic and each customer's traffic requires its own VLAN while being transported through the QFabric. Instead of using individual VLANs for each customer (which can quickly lead to VLAN exhaustion), you can apply a unique inner (C-VLAN) tag to each customer's traffic and then apply a single outer tag (S-VLAN) tag for transport through the QFabric. This allows you to segregate your customers's traffic while consuming only one QFabric VLAN. Use the **inner-tag** option of the **mapping** statement to accomplish this.

Q-in-Q tunneling does not affect any class-of-service (CoS) values that are configured on a C-VLAN. These settings are retained in the C-VLAN tag and can be used after a packet leaves an S-VLAN. CoS values are not copied from C-VLAN tags to S-VLAN tags.

Depending on your interface configuration, you might need to adjust the MTU value on your trunk or access ports to accommodate the 4 bytes used for the tag added by Q-in-Q tunneling. For example, if you use the default MTU value of 1514 bytes on your access and trunk ports, you need to make one of the following adjustments:

- Reduce the MTU on the access links by at least 4 bytes so that the frames do not exceed the MTU of the trunk link when S-VLAN tags are added.
- Increase the MTU on the trunk link so that the link can handle the larger frame size.



NOTE: You can configure Q-in-Q tunneling only on access ports (not trunk ports).

How VLAN Translation Works

VLAN translation replaces an incoming C-VLAN tag with an S-VLAN tag instead of adding an additional tag. The C-VLAN tag is therefore lost, so a single-tagged packet is normally untagged when it leaves the S-VLAN (at the other end of the link). If an incoming packet has had Q-in-Q tunneling applied in advance, VLAN translation replaces the outer tag and the inner tag is retained when the packet leaves the S-VLAN at the other end of the link.

To configure VLAN translation, use the **mapping swap** statement at the **[edit vlans interface]** hierarchy level.



NOTE: You can configure VLAN translation on access ports only. You cannot configure it on trunk ports, and you cannot configure Q-in-Q tunneling on the same access port.



NOTE: VLAN translation is not supported on QFabric systems.

Mapping C-VLANs to S-VLANs

The three ways to map C-VLANs to an S-VLAN are:

- All-in-one bundling—Use the **edit vlans s-vlan-name dot1q-tunneling** statement without specifying customer VLANs. All packets received on all access interfaces (including untagged packets) are mapped to the S-VLAN.
- Many-to-one bundling—Use the **edit vlans s-vlan-name dot1q-tunneling customer-vlans** statement to specify which C-VLANs are mapped to the S-VLAN. Use this method when you want a subset of the C-VLANs to be part of the S-VLAN. If you want untagged or priority tagged packets to be mapped to the S-VLAN, use the **native** option with the **customer-vlans** statement. (Priority tagged packets have their VLAN ID set to 0, and their priority code point bits might be configured with a CoS value.)
- Mapping a specific interface—Use the **edit vlans s-vlan-name interface interface-name mapping** statement to specify a C-VLAN for a given S-VLAN. This configuration applies to only one interface—not all access interfaces as with all-in-one and many-to-one bundling. If you want untagged or priority tagged packets to be mapped to the S-VLAN, use the **native** option with the **customer-vlans** statement.

This method has two options: swap and push. With the push option, a packet retains its tag and an additional VLAN tag is added. With the swap option, the incoming tag is replaced with an S-VLAN tag. (This is VLAN translation.)

- You can configure multiple push rules for a given S-VLAN and interface. That is, you can configure an interface so that the same S-VLAN tag is added to packets arriving from multiple C-VLANs.
- You can configure only one swap rule for a given S-VLAN and interface.

This functionality is typically used to keep traffic from different customers separate or to provide individualized treatment for traffic on a certain interface.

If you configure multiple methods, the switch prioritizes the mappings in the following order:

1. Specific interface mapping
2. Many-to-one bundling
3. All-in-one bundling

You cannot configure overlapping rules for the same C-VLAN using the same mapping method. For example, you cannot use many-to-one bundling to map C-VLAN 100 to two different S-VLANs.

Routed VLAN Interfaces on Q-in-Q VLANs

Routed VLAN interfaces (RVIs) are supported on Q-in-Q VLANs. Packets arriving on an RVI that is using Q-in-Q VLANs are routed regardless of whether the packet is single-tagged or double-tagged. The original C-VLAN tag is dropped when a packet leaves the VLAN that it originated in. Outgoing routed packets retain any S-VLAN tag only when exiting a trunk interface—S-VLAN tags are dropped when traffic exits an access interface.

Constraints for Q-in-Q Tunneling and VLAN Translation

Be aware of the following constraints when configuring Q-in-Q tunneling and VLAN translation:

- Most access port security features are not supported with Q-in-Q tunneling and VLAN translation.
- Configuring Q-in-Q tunneling and VLAN translation on the same port is not supported.
- You can configure at most one VLAN translation for a given VLAN and interface. For example, you can create no more than one translation for VLAN 100 on interface xe-0/0/0.
- The combined total of VLANs and rules for Q-in-Q tunneling and VLAN translation cannot exceed 6000. For example, you can configure and commit 4000 VLANs and 2000 rules for Q-in-Q tunneling and VLAN translation. However, you cannot configure 4000 VLANs and 2500 rules for Q-in-Q tunneling and VLAN translation. If you try to commit a configuration that exceeds the limit, you see CLI and syslog errors that inform you about the problem.
- MAC addresses are learned from S-VLANs, not C-VLANs.
- Broadcast, unknown unicast, and multicast traffic is forwarded to all members in the S-VLAN.

- The following features are not supported with Q-in-Q tunneling:
 - DHCP relay
 - Fibre Channel over Ethernet
 - IP Source Guard
- The following features are not supported with VLAN translation:
 - Fibre Channel over Ethernet
 - Firewall filter applied to a port or VLAN in the output direction
 - Private VLANs
 - VLAN Spanning Tree Protocol
 - Reflective relay

Related Documentation

- [Configuring Q-in-Q Tunneling on page 2062](#)
- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 2266](#)
- [mapping on page 2166](#)
- [mtu on page 2603](#)

Layer 2 Protocol Tunneling

- [Understanding Layer 2 Protocol Tunneling on page 1910](#)

Understanding Layer 2 Protocol Tunneling

Layer 2 protocol tunneling (L2PT) allows service providers or data centers to send customer Layer 2 protocol data units (PDUs) across a cloud. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

For example, a customer might want to run STP between remote sites that are connected by a service provider network. In this case, the STP PDUs sent by the customer equipment must be tunneled through the service provider network. Otherwise, the customer STP PDUs would be processed by the STP protocol running on the service provider switches.



NOTE: Layer 2 protocol tunneling is not supported on QFabric systems.

This topic includes:

- [Layer 2 Protocols Supported by L2PT on page 1911](#)
- [How L2PT Works on page 1911](#)
- [L2PT Basics on page 1913](#)

Layer 2 Protocols Supported by L2PT

L2PT supports the following Layer 2 protocols:

- 802.1X authentication
- 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)



NOTE: If you enable L2PT for untagged OAM LFM (Operation, Administration, and Maintenance of link fault management) packets, do not configure link fault management (LFM) on the corresponding access interface.

- Cisco Discovery Protocol (CDP)
- Ethernet local management interface (E-LMI)
- MVRP VLAN Registration Protocol (MVRP)
- Link Aggregation Control Protocol (LACP)



NOTE: If you enable L2PT for untagged LACP packets, do not configure Link Aggregation Control Protocol (LACP) on the corresponding access interface.

- Link Layer Discovery Protocol (LLDP)
- Multiple MAC Registration Protocol (MMRP)
- Multiple VLAN Registration Protocol (MVRP)
- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)
- Unidirectional Link Detection (UDLD)
- VLAN Spanning Tree Protocol (VSTP)
- VLAN Trunking Protocol (VTP)



NOTE: You cannot configure all of these protocols on QFX Series devices. However, you can tunnel all of them through a QFX Series switch using L2PT.

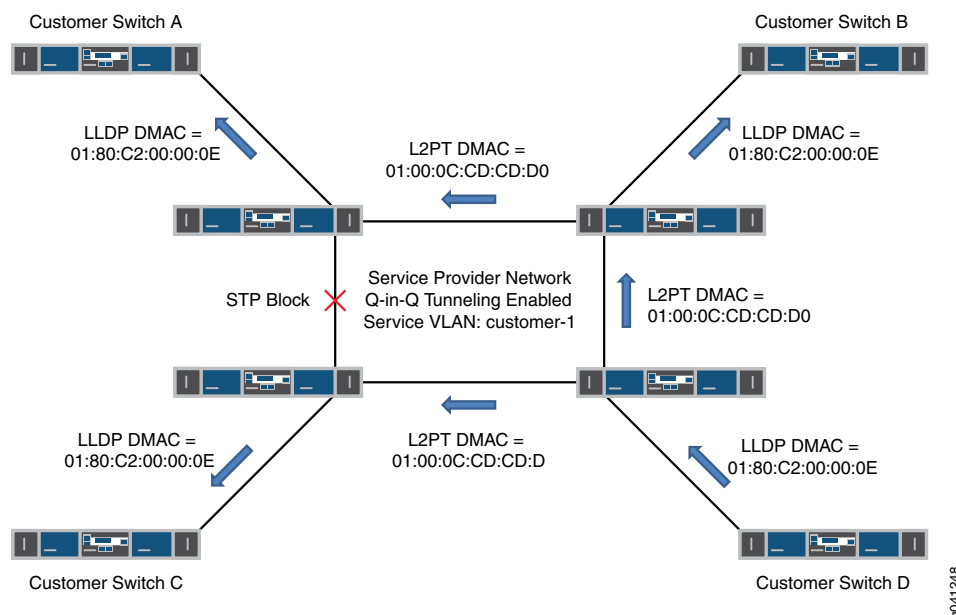
How L2PT Works

L2PT works by encapsulating Layer 2 PDUs, tunneling them across a service provider network, and de-encapsulating them for delivery to their destination switches. L2PT encapsulates Layer 2 PDUs by enabling the ingress provider edge (PE) device to rewrite the PDUs' destination media access control (DMAC) addresses before forwarding them onto the service provider network. The devices in the service provider network treat these

encapsulated PDUs as multicast Ethernet packets. Upon receipt of these PDUs, the egress PE devices de-encapsulate them by replacing the DMAC addresses with the addresses of the Layer 2 protocol that is being tunneled before forwarding the PDUs to their destination switches.

This process is illustrated in [Figure 51 on page 1912](#):

Figure 51: L2PT Example



1. Customer Switch D sends to the service provider network an LLDP PDU that is ultimately intended for the other switches in the customer network.
2. The receiving provider switch adds the L2PT DMAC and sends the frame with the encapsulated LLDP PDU to the other switches in the service provider network.
3. When the other service provider switches receive the frame, they restore the LLDP DMAC and send it to Customer Switches A, B, and C.

[Table 184 on page 1912](#) lists the destination MAC addresses of the supported Layer 2 protocols:

Table 184: Protocol Destination MAC Addresses

Protocol	Ethernet Encapsulation	MAC Address
802.1X	Ether-II	01:80:C2:00:00:03
802.3ah	Ether-II	01:80:C2:00:00:02
Cisco Discovery Protocol (CDP)	SNAP	01:00:0C:CC:CC:CC

Table 184: Protocol Destination MAC Addresses (*continued*)

Protocol	Ethernet Encapsulation	MAC Address
Ethernet local management interface (E-LMI)	Ether-II	01:80:C2:00:00:07
MVRP VLAN Registration Protocol (MVRP)	Ether-II	01:80:C2:00:00:21
Link Aggregation Control Protocol (LACP)	Ether-II	01:80:C2:00:00:02
Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)	LLC	01:80:C2:00:00:00
Link Layer Discovery Protocol (LLDP)	Ether-II	01:80:C2:00:00:0E
Multiple MAC Registration Protocol (MMRP)	Ether-II	01:80:C2:00:00:20
Unidirectional Link Detection (UDLD)	SNAP	01:00:0C:CC:CC:CC
VLAN Spanning Tree Protocol (VSTP)	SNAP	01:00:0C:CC:CC:CD
VLAN Trunking Protocol (VTP)	SNAP	01:00:0C:CC:CC:CC

When a PE device receives a Layer 2 control PDU from any of the customer PE devices, it changes the destination MAC address to 01:00:0C:CD:CD:D0. The modified packet is then sent to the provider network. All devices on the provider network treat these packets as multicast Ethernet packets and deliver them to all PE devices in the VLAN. The egress PE devices receive all the control PDUs with the same MAC address (01:00:0C:CD:CD:D0). Then they identify the packet type by doing deeper packet inspection and replace the destination MAC address 01:00:0C:CD:CD:D0 with the appropriate destination address. The modified PDUs are sent out to the customer PE devices so that the Layer 2 control PDUs are delivered, in their original state, across the provider network. The L2PT protocol is valid for all types of packets (untagged, tagged, and Q-in-Q tagged).



NOTE: VLAN translation is not compatible with L2PT. (You configure VLAN translation with the **mapping swap** statement at the **[edit vlans interface]** hierarchy level.)

L2PT Basics

L2PT is enabled on a per-VLAN basis. When you enable L2PT on a VLAN, all access interfaces are considered to be customer-facing interfaces, all trunk interfaces are considered to be service provider network-facing interfaces, and the specified Layer 2 protocol is disabled on the access interfaces. L2PT acts only on logical interfaces of the family **ethernet-switching**. L2PT PDUs are flooded to all trunk and access ports within a given service VLAN.



NOTE: Access interfaces in an L2PT-enabled VLAN should not receive L2PT-tunneled PDUs. If an access interface does receive L2PT-tunneled PDUs, it might mean that there is a loop in the network. As a result, the interface is shut down.

You configure L2PT at the `[edit vlans vlan-name dot1q-tunneling]` hierarchy level, meaning Q-in-Q tunneling is (and must be) enabled.



NOTE: If you want to tunnel untagged or priority-tagged Layer 2 control PDUs, then you must configure the switch to map untagged and priority-tagged packets to an L2PT-enabled VLAN. For more information about assigning untagged and priority-tagged packets to VLANs, see [“Understanding Q-in-Q Tunneling and VLAN Translation” on page 1906](#).

Related Documentation

- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)
- [Configuring Layer 2 Protocol Tunneling on page 2063](#)

Proxy ARP

- [Understanding Proxy ARP on page 1914](#)

Understanding Proxy ARP

You can configure proxy Address Resolution Protocol (ARP) to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

- [What Is ARP? on page 1914](#)
- [Proxy ARP Overview on page 1915](#)
- [Best Practices for Proxy ARP on page 1915](#)

What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch

maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.



NOTE: For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for a VLAN by using a routed VLAN interface (RVI).

QFX Series devices support two modes of proxy ARP, restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- **Restricted**—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.
- **Unrestricted**—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Juniper Networks Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

Best Practices for Proxy ARP

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP to restricted mode.
- Use restricted mode when configuring proxy ARP on RVIs.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

Related Documentation

- [Configuring Proxy ARP on page 2065](#)
- [proxy-arp on page 2129](#)

CHAPTER 21

Configuration

- [Configuration Examples on page 1917](#)
- [Private VLAN Configuration Examples on page 2005](#)
- [Q-in-Q Tunneling Configuration Example on page 2036](#)
- [Configuration Tasks on page 2039](#)
- [Private VLAN Configuration Tasks on page 2059](#)
- [Q-in-Q Tunneling Configuration Tasks on page 2062](#)
- [Proxy ARP Configuration Task on page 2065](#)
- [Configuration Statements on page 2066](#)
- [MVRP Configuration Statements on page 2148](#)
- [Private VLAN Configuration Statements on page 2153](#)
- [Q-in-Q Tunneling Configuration Statements on page 2159](#)

Configuration Examples

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on page 1953](#)
- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Example: Configuring Routing Between VLANs on One Switch on page 1976](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 1982](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1994](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999](#)

- [Example: Disabling MAC Learning on page 2003](#)
- [Example: Disabling MAC Learning in a VLAN on page 2003](#)

Example: Configuring Faster Convergence and Improving Network Stability with RSTP

The QFX Series products use Rapid Spanning Tree Protocol (RSTP) to provide a loop-free topology. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state. RSTP provides quicker reconvergence time than original STP because it uses protocol handshake messages rather than fixed timeouts. Eliminating the need to wait for timers to expire makes RSTP more efficient than STP.

This example describes how to configure RSTP on four QFX3500 switches:

- [Requirements on page 1918](#)
- [Overview and Topology on page 1918](#)
- [Configuring RSTP on Switch 1 on page 1920](#)
- [Configuring RSTP on Switch 2 on page 1922](#)
- [Configuring RSTP on Switch 3 on page 1924](#)
- [Configuring RSTP on Switch 4 on page 1927](#)
- [Verification on page 1929](#)

Requirements


This example uses the following hardware and software components:

- Junos OS Release 11.1 for the QFX3500 switches
- Four QFX3500 switches

Overview and Topology

In this example, QFX3500 switches are connected in the topology displayed in [Figure 52 on page 1919](#) to create a loop-free topology.

Figure 52: Network Topology for RSTP



The interfaces shown in [Table 185 on page 1919](#) will be configured for RSTP.



NOTE: You can configure RSTP on logical or physical interfaces. This example shows RSTP configured on logical interfaces.

Table 185: Topology for Configuring RSTP on the QFX Series

Components	Settings
Switch 1	<p>The following ports on Switch 1 are connected in this way:</p> <ul style="list-style-type: none">• xe-0/0/9 is connected to Switch 2• xe-0/0/13 is connected to Switch 4• xe-0/0/11 is connected to Switch 3
Switch 2	<p>The following ports on Switch 2 are connected in this way:</p> <ul style="list-style-type: none">• xe-0/0/14 is connected to Switch 1• xe-0/0/18 is connected to Switch 3
Switch 3	<p>The following ports on Switch 3 are connected in this way:</p> <ul style="list-style-type: none">• xe-0/0/26 is connected to Switch 1• xe-0/0/28 is connected to Switch 2• xe-0/0/24 is connected to Switch 4

Table 185: Topology for Configuring RSTP on the QFX Series (*continued*)

Components	Settings
Switch 4	<p>The following ports on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> • <code>xe-0/0/19</code> is connected to Switch 1 • <code>xe-0/0/23</code> is connected to Switch 3
VLAN names and tag IDs	<p><code>sales-vlan</code>, tag 10</p> <p><code>engineering-vlan</code>, tag 20</p> <p><code>publications-vlan</code>, tag 30</p> <p><code>support-vlan</code>, tag 40</p>

This configuration example creates a loop-free topology between four switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

Configuring RSTP on Switch 1

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface xe-0/0/13.0 cost 1000
set protocols rstp interface xe-0/0/13.0 mode point-to-point
set protocols rstp interface xe-0/0/9.0 cost 1000
set protocols rstp interface xe-0/0/9.0 mode point-to-point
set protocols rstp interface xe-0/0/11.0 cost 1000
set protocols rstp interface xe-0/0/11.0 mode point-to-point
```

Step-by-Step Procedure To configure interfaces and RSTP on Switch 1:

1. Configure the VLANs **sales-vlan**, **engineering-vlan** and **publications-vlan**, and **support-vlan**:


```
[edit vlans]
user@switch1# set sales-vlan description "Sales VLAN"
user@switch1# set sales-vlan vlan-id 10
user@switch1# set engineering-vlan description "Engineering VLAN"
user@switch1# set engineering-vlan vlan-id 20
user@switch1# set publications-vlan description "Publications VLAN"
user@switch1# set publications-vlan vlan-id 30
```
2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:


```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```
3. Configure the port mode for the interfaces:


```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
```
4. Configure RSTP on the switch:


```
[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface xe-0/0/13.0 cost 1000
user@switch1# rstp interface xe-0/0/13.0 mode point-to-point
user@switch1# rstp interface xe-0/0/9.0 cost 1000
user@switch1# rstp interface xe-0/0/9.0 mode point-to-point
user@switch1# rstp interface xe-0/0/11.0 cost 1000
user@switch1# rstp interface xe-0/0/11.0 mode point-to-point
```

Results Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  xe-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
```

```
    }
  }
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members [10 20 30 40];
      }
    }
  }
}
}
protocols {
  rstp {
    bridge-priority 16k;
    interface xe-0/0/13.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/9.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/11.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
}
vlands {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
  publications-vlan {
    vlan-id 30;
  }
  support-vlan {
    vlan-id 40;
  }
}
```

Configuring RSTP on Switch 2

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
```

```

set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface xe-0/0/14.0 cost 1000
set protocols rstp interface xe-0/0/14.0 mode point-to-point
set protocols rstp interface xe-0/0/18.0 cost 1000
set protocols rstp interface xe-0/0/18.0 mode point-to-point

```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 2:

1. Configure the VLANs `sales-vlan`, `engineering-vlan` and `publications-vlan`, and `support-vlan`:


```

[edit vlans]
user@switch2# set sales-vlan description "Sales VLAN"
user@switch2# set sales-vlan vlan-id 10
user@switch2# set engineering-vlan description "Engineering VLAN"
user@switch2# set engineering-vlan vlan-id 20
user@switch2# set publications-vlan description "Publications VLAN"
user@switch2# set publications-vlan vlan-id 30
user@switch2# set support-vlan vlan-description "Support VLAN"
user@switch2# set publications-vlan vlan-id 40

```
2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:


```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]

```
3. Configure the port mode for the interfaces:


```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching port-mode trunk

```
4. Configure RSTP on the switch:


```

[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface xe-0/0/14.0 cost 1000
user@switch2# rstp interface xe-0/0/14.0 mode point-to-point
user@switch2# rstp interface xe-0/0/18.0 cost 1000
user@switch2# rstp interface xe-0/0/18.0 mode point-to-point

```

Results Check the results of the configuration:

```

user@switch2> show configuration
interfaces {
  xe-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;

```

```
        vlan {
            members [10 20 30 40];
        }
    }
}
xe-0/0/18 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members [10 20 30 40];
            }
        }
    }
}
}
protocols {
    rstp {
        bridge-priority 32k;
        interface xe-0/0/14.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/18.0 {
            cost 1000;
            mode point-to-point;
        }
    }
}
vlangs {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
```

Configuring RSTP on Switch 3

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
```



```

set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 8k
set protocols rstp interface xe-0/0/26.0 cost 1000
set protocols rstp interface xe-0/0/26.0 mode point-to-point
set protocols rstp interface xe-0/0/28.0 cost 1000
set protocols rstp interface xe-0/0/28.0 mode point-to-point
set protocols rstp interface xe-0/0/24.0 cost 1000
set protocols rstp interface xe-0/0/24.0 mode point-to-point

```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 3:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch3# set sales-vlan description "Sales VLAN"
user@switch3# set sales-vlan vlan-id 10
user@switch3# set engineering-vlan description "Engineering VLAN"
user@switch3# set engineering-vlan vlan-id 20
user@switch3# set publications-vlan description "Publications VLAN"
user@switch3# set publications-vlan vlan-id 30
user@switch3# set support-vlan description "Support VLAN"
user@switch3# set publications-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:

```

[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching port-mode trunk

```

4. Configure RSTP on the switch:

```

[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface xe-0/0/26.0 cost 1000
user@switch3# rstp interface xe-0/0/26.0 mode point-to-point
user@switch3# rstp interface xe-0/0/28.0 cost 1000
user@switch3# rstp interface xe-0/0/28.0 mode point-to-point
user@switch3# rstp interface xe-0/0/24.0 cost 1000
user@switch3# rstp interface xe-0/0/24.0 mode point-to-point

```

Results Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
  xe-0/0/26 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/28 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/24 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 8k;
    interface xe-0/0/26.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/28.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/24.0 {
      cost 1000;
      mode point-to-point;
    }
  }
  bridge-priority 8k;
}
vllans {
```

```

sales-vlan {
  vlan-id 10;
}
engineering-vlan {
  vlan-id 20;
}
publications-vlan {
  vlan-id 30;
}
support-vlan {
  vlan-id 40;
}
}

```

Configuring RSTP on Switch 4

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface xe-0/0/23.0 cost 1000
set protocols rstp interface xe-0/0/23.0 mode point-to-point
set protocols rstp interface xe-0/0/19.0 cost 1000
set protocols rstp interface xe-0/0/19.0 mode point-to-point

```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 4:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch4# set sales-vlan description "Sales VLAN"
user@switch4# set sales-vlan vlan-id 10
user@switch4# set engineering-vlan description "Engineering VLAN"
user@switch4# set engineering-vlan vlan-id 20
user@switch4# set publications-vlan description "Publications VLAN"
user@switch4# set publications-vlan vlan-id 30
user@switch4# set support-vlan description "Support VLAN"
user@switch4# set publications-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch4# set xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface all cost 1000
user@switch4# rstp interface xe-0/0/23.0 cost 1000
user@switch4# rstp interface xe-0/0/23.0 mode point-to-point
user@switch4# rstp interface xe-0/0/19.0 cost 1000
user@switch4# rstp interface xe-0/0/19.0 mode point-to-point
```

Results Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  xe-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/19 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 16k;
    interface xe-0/0/23.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/19.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
vllans {
  sales-vlan {
    vlan-id 10;
  }
}
```

```

}
engineering-vlan {
  vlan-id 20;
}
publications-vlan {
  vlan-id 30;
}
support-vlan {
  vlan-id 40;
}
}

```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying RSTP Configuration on Switch 1 on page 1929](#)
- [Verifying RSTP Configuration on Switch 2 on page 1929](#)
- [Verifying RSTP Configuration on Switch 3 on page 1930](#)
- [Verifying RSTP Configuration on Switch 4 on page 1930](#)

Verifying RSTP Configuration on Switch 1

Purpose Verify that the RSTP configuration on Switch 1 is correct.

Action In operational mode, issue the **show spanning-tree interface** command:

```
user@switch1> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/13.0	128:527	128:525	16384.0019e25040e0	1000	BLK	ALT
xe-0/0/9.0	128:529	128:513	32768.0019e2503d20	1000	BLK	ALT
xe-0/0/11.0	128:531	128:513	8192.0019e25051e0	1000	FWD	ROOT

Meaning See the topology in [Figure 52 on page 1919](#). The operational mode command **show spanning-tree interface** shows that **xe-0/0/13.0** is in a forwarding state. The other interfaces on Switch 1 are blocked.

Verifying RSTP Configuration on Switch 2

Purpose Verify that the RSTP configuration on Switch 2 is correct.

Action In operational mode issue the **show spanning-tree interface** command:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/14.0	128:513	128:513	32768.0019e2503d20	1000	BLK	DESC
xe-0/0/18.0	128:519	128:515	8192.0019e25051e0	1000	FWD	ROOT

Meaning See the topology in [Figure 52 on page 1919](#). The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/18.0** is in a forwarding state and the root port. The other interface on Switch 2 is blocked.

Verifying RSTP Configuration on Switch 3

Purpose Verify that the RSTP configuration on Switch 3 is correct.

Action In operational mode, issue the **show spanning-tree interface** command:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/26.0	128:513	128:513	8192.0019e25051e0	1000	FWD	DESC
xe-0/0/28.0	128:515	128:515	8192.0019e25051e0	1000	FWD	DESC
xe-0/0/24.0	128:517	128:517	8192.0019e25051e0	1000	FWD	DESC

Meaning See the topology in [Figure 52 on page 1919](#). The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

Verifying RSTP Configuration on Switch 4

Purpose Verify the RSTP configuration on Switch 4.

Action In operational mode, issue the **show spanning-tree interface** command:

```
user@switch4> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/23.0	128:523	128:517	8192.0019e25051e0	1000	FWD	ROOT
xe-0/0/19.0	128:525	128:525	16384.0019e25040e0	1000	FWD	DESG

Meaning See the topology in [Figure 52 on page 1919](#). The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/23.0** is the root interface and is in the forwarding state.

Related Documentation

- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
- [Understanding RSTP on page 1901](#)

Example: Configuring Network Regions for VLANs with MSTP

Multiple Spanning Tree Protocol (MSTP) is used to create a loop-free topology in networks using multiple spanning-tree regions, each region containing multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates more efficient load sharing across redundant links.

You can create up to 64 MSTI instances for QFX Series products, and each MSTI supports up to 4094 VLANs.

This example describes how to configure MSTP on four QFX3500 switches:

- [Requirements on page 1931](#)
- [Overview and Topology on page 1932](#)
- [Configuring MSTP on Switch 1 on page 1934](#)
- [Configuring MSTP on Switch 2 on page 1937](#)
- [Configuring MSTP on Switch 3 on page 1939](#)
- [Configuring MSTP on Switch 4 on page 1942](#)
- [Verification on page 1945](#)

Requirements


This example uses the following hardware and software components:

- Junos OS Release 11.1 for the QFX3500 switches
- Four QFX3500 switches

Overview and Topology

When the number of VLANs grows in a network, MSTP provides a more faster way of creating a loop-free topology using MSTIs. Each MSTI in the spanning-tree domain maintains its own tree. Each tree can be mapped to different links, utilizing bandwidth that would be unavailable to a single tree. MSTIs reduce demand on system resources.

Figure 53: Network Topology for MSTP



The interfaces shown in [Table 186 on page 1933](#) will be configured for MSTP.



NOTE: You can configure MSTP on logical or physical interfaces. This example shows MSTP configured on logical interfaces.

Table 186: Topology for Configuring MSTP on the QFX Series

Components	Settings
Switch 1	<p>The following ports on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"> • xe-0/0/9 is connected to Switch 2 • xe-0/0/13 is connected to Switch 4 • xe-0/0/11 is connected to Switch 3
Switch 2	<p>The following ports on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"> • xe-0/0/14 is connected to Switch 1 • xe-0/0/18 is connected to Switch 3
Switch 3	<p>The following ports on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"> • xe-0/0/26 is connected to Switch 1 • xe-0/0/28 is connected to Switch 2 • xe-0/0/24 is connected to Switch 4
Switch 4	<p>The following ports on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> • xe-0/0/19 is connected to Switch 1 • xe-0/0/23 is connected to Switch 3
VLAN names and tag IDs	sales-vlan , tag 10 engineering-vlan , tag 20 publications-vlan , tag 30 support-vlan , tag 40
MSTIs	1 2

The topology in [Figure 53 on page 1932](#) shows a Common Internal Spanning Tree (CIST). The CIST is a single spanning tree connecting all devices in the network. The switch with the highest priority is elected as the root bridge of the CIST.

Also in an MSTP topology are ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

In this example, one MSTP region, **region1**, contains Switch 1, Switch 2, Switch 3, and Switch 4. Within the region, four VLANs are created:

- The **sales-vlan** supports sales traffic and has a VLAN tag identifier of 10.
- The **engineering-vlan** supports data traffic and has a VLAN tag identifier of 20.
- The **publications-vlan** supports publications VLAN traffic (for supplicants that fail 802.1X authentication) and has a VLAN tag identifier of 30.
- The **support-vlan** supports video traffic and has a VLAN tag identifier of 40.

The VLANs are associated with specific interfaces on each of the four switches. Two MSTIs, 1 and 2, are then associated with the VLAN tag identifiers, and some MSTP parameters, such as cost, are configured on each switch.

Configuring MSTP on Switch 1

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface xe-0/0/13.0 cost 1000
set protocols mstp interface xe-0/0/13.0 mode point-to-point
set protocols mstp interface xe-0/0/9.0 cost 1000
set protocols mstp interface xe-0/0/9.0 mode point-to-point
set protocols mstp interface xe-0/0/11.0 cost 1000
set protocols mstp interface xe-0/0/11.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 1 interface xe-0/0/11.0 cost 4000
set protocols mstp msti 2 bridge-priority 8k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 1:

1. Configure the VLANs **sales-vlan**, **engineering-vlan**, **publications-vlan**, and **support-vlan**:

```
[edit vlans]
user@switch1# set sales-vlan description "Sales VLAN"
user@switch1# set sales-vlan vlan-id 10
user@switch1# set engineering-vlan description "Engineering VLAN"
user@switch1# set engineering-vlan vlan-id 20
user@switch1# set publications-vlan description "Publications VLAN"
user@switch1# set publications-vlan vlan-id 30
user@switch1# set support-vlan description "Support VLAN"
user@switch1# set publications-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch1# mstp configuration-name region1
user@switch1# mstp bridge-priority 16k
user@switch1# mstp interface xe-0/0/13.0 cost 1000
user@switch1# mstp interface xe-0/0/13.0 mode point-to-point
user@switch1# mstp interface xe-0/0/9.0 cost 1000
user@switch1# mstp interface xe-0/0/9.0 mode point-to-point
user@switch1# mstp interface xe-0/0/11.0 cost 4000
user@switch1# mstp interface xe-0/0/11.0 mode point-to-point
user@switch1# mstp msti 1 bridge-priority 16k
user@switch1# mstp msti 1 vlan [10 20]
user@switch1# mstp msti 1 interface xe-0/0/11.0 cost 4000
user@switch1# mstp msti 2 bridge-priority 8k
user@switch1# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  xe-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  xe-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
        }
      }
    }
  }
}
```

```
        members 40;
      }
    }
  }
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members 10;
        members 20;
        members 30;
        members 40;
      }
    }
  }
}
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 16k;
    interface xe-0/0/13.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/9.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/11.0 {
      cost 4000;
      mode point-to-point;
    }
  }
  msti 1 {
    bridge-priority 16k;
    vlan [ 10 20 ];
    interface xe-0/0/11.0 {
      cost 4000;
    }
  }
  msti 2 {
    bridge-priority 8k;
    vlan [ 30 40 ];
  }
}
vlangs {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
  publications-vlan {
    vlan-id 30;
  }
}
```

```

    }
    support-vlan {
        vlan-id 40;
    }
}

```

Configuring MSTP on Switch 2

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 32k
set protocols mstp interface xe-0/0/14.0 cost 1000
set protocols mstp interface xe-0/0/14.0 mode point-to-point
set protocols mstp interface xe-0/0/18.0 cost 1000
set protocols mstp interface xe-0/0/18.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 32k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan [30 40]

```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 2:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch2# set sales-vlan description "Sales VLAN"
user@switch2# set sales-vlan vlan-id 10
user@switch2# set engineering-vlan description "Engineering VLAN"
user@switch2# set engineering-vlan vlan-id 20
user@switch2# set publications-vlan description "Publications VLAN"
user@switch2# set publications-vlan vlan-id 30
user@switch2# set support-vlan description "Support VLAN"
user@switch2# set support-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:

```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching port-mode trunk

```

- user@switch2# set xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch2# mstp configuration-name region1
user@switch2# mstp bridge-priority 32k
user@switch2# mstp interface xe-0/0/14.0 cost 1000
user@switch2# mstp interface xe-0/0/14.0 mode point-to-point
user@switch2# mstp interface xe-0/0/18.0 cost 1000
user@switch2# mstp interface xe-0/0/18.0 mode point-to-point
user@switch2# mstp interface all cost 1000
user@switch2# mstp msti 1 bridge-priority 32k
user@switch2# mstp msti 1 vlan [10 20]
user@switch2# mstp msti 2 bridge-priority 4k
user@switch2# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
  xe-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  xe-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 32k;
    interface xe-0/0/14.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/18.0 {
      cost 1000;
    }
  }
}
```

```

        mode point-to-point;
    }
    msti 1 {
        bridge-priority 32k;
        vlan [ 10 20 ];
    }
    msti 2 {
        bridge-priority 4k;
        vlan [ 30 40 ];
    }
}
vlangs {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}

```

Configuring MSTP on Switch 3

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 8k
set protocols mstp interface xe-0/0/26.0 cost 1000
set protocols mstp interface xe-0/0/26.0 mode point-to-point
set protocols mstp interface xe-0/0/28.0 cost 1000
set protocols mstp interface xe-0/0/28.0 mode point-to-point
set protocols mstp interface xe-0/0/24.0 cost 1000
set protocols mstp interface xe-0/0/24.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan [10 20]

```

- Step-by-Step Procedure**
- To configure interfaces and MSTP on Switch 3:
1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:


```
[edit vlans]
user@switch3# set sales-vlan description "Sales VLAN"
user@switch3# set sales-vlan vlan-id 10
user@switch3# set engineering-vlan description "Engineering VLAN"
user@switch3# set engineering-vlan vlan-id 20
user@switch3# set publications-vlan description "Publications VLAN"
user@switch3# set publications-vlan vlan-id 30
user@switch3# set support-vlan description "Support VLAN"
user@switch3# set publications-vlan vlan-id 40
```
 2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:


```
[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```
 3. Configure the port mode for the interfaces:


```
[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
```
 4. Configure MSTP on the switch, including the two MSTIs:


```
[edit protocols]
user@switch3# mstp configuration-name region1
user@switch3# mstp bridge-priority 8k
user@switch3# mstp interface xe-0/0/26.0 cost 1000
user@switch3# mstp interface xe-0/0/26.0 mode point-to-point
user@switch3# mstp interface xe-0/0/28.0 cost 1000
user@switch3# mstp interface xe-0/0/28.0 mode point-to-point
user@switch3# mstp interface xe-0/0/24.0 cost 1000
user@switch3# mstp interface xe-0/0/24.0 mode point-to-point
user@switch3# mstp interface all cost 1000
user@switch3# mstp msti 1 bridge-priority 4k
user@switch3# mstp msti 1 vlan [10 20]
user@switch3# mstp msti 2 bridge-priority 16k
user@switch3# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
  xe-0/0/26 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
```



```

        members 30;
        members 40;
    }
}
}
xe-0/0/28 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
xe-0/0/24 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
}
}
protocols {
    mstp {
        configuration-name region1;
        bridge-priority 8k;
        interface xe-0/0/26.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/28.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/24.0 {
            cost 1000;
            mode point-to-point;
        }
        msti 1 {
            bridge-priority 4k;
            vlan [ 10 20 ];
        }
        msti 2 {
            bridge-priority 16k;

```

```
        vlan [ 30 40 ];
    }
}
vlangs {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
```

Configuring MSTP on Switch 4

CLI Quick Configuration To quickly configure interfaces and MSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface xe-0/0/23.0 cost 1000
set protocols mstp interface xe-0/0/23.0 mode point-to-point
set protocols mstp interface xe-0/0/19.0 cost 1000
set protocols mstp interface xe-0/0/19.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 32k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure To configure interfaces and MSTP on Switch 4:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```
[edit vlans]
user@switch4# set sales-vlan description "Sales VLAN"
user@switch4# set sales-vlan vlan-id 10
user@switch4# set engineering-vlan description "Engineering VLAN"
user@switch4# set engineering-vlan vlan-id 20
user@switch4# set publications-vlan description "Publications VLAN"
user@switch4# set publications-vlan vlan-id 30
user@switch4# set support-vlan description "Support VLAN"
user@switch4# set publications-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch4# set xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch4# mstp configuration-name region1
user@switch4# mstp bridge-priority 16k
user@switch4# mstp interface all cost 1000
user@switch4# mstp interface xe-0/0/23.0 cost 1000
user@switch4# mstp interface xe-0/0/23.0 mode point-to-point
user@switch4# mstp interface xe-0/0/19.0 cost 1000
user@switch4# mstp interface xe-0/0/19.0 mode point-to-point
user@switch4# mstp msti 1 bridge-priority 16k
user@switch4# mstp msti 1 vlan [10 20]
user@switch4# mstp msti 2 bridge-priority 32k
user@switch4# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  xe-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  xe-0/0/19 {
```

```
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 16k;
    interface xe-0/0/23.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/19.0 {
      cost 1000;
      mode point-to-point;
    }
    msti 1 {
      bridge-priority 16k;
      vlan [ 10 20 ];
    }
    msti 2 {
      bridge-priority 32k;
      vlan [ 30 40 ];
    }
  }
}
vlangs {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
  publications-vlan {
    vlan-id 30;
  }
  support-vlan {
    vlan-id 40;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying MSTP Configuration on Switch 1 on page 1945](#)
- [Verifying MSTP Configuration on Switch 2 on page 1947](#)
- [Verifying MSTP Configuration on Switch 3 on page 1949](#)
- [Verifying MSTP Configuration on Switch 4 on page 1951](#)

Verifying MSTP Configuration on Switch 1

Purpose Verify the MSTP configuration on Switch 1.

Action Use the operational mode commands:

```
user@switch1> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/13.0	128:527	128:525	16384.0019e25040e0	1000	FWD	ROOT
xe-0/0/9.0	128:529	128:513	32768.0019e2503d20	1000	BLK	ALT
xe-0/0/11.0	128:531	128:513	8192.0019e25051e0	4000	BLK	ALT

```
Spanning tree interface parameters for instance 1
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/13.0	128:527	128:525	16385.0019e25040e0	1000	FWD	ROOT
xe-0/0/9.0	128:529	128:513	32769.0019e2503d20	1000	BLK	ALT
xe-0/0/11.0	128:531	128:513	4097.0019e25051e0	4000	BLK	ALT

```
Spanning tree interface parameters for instance 2
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/13.0	128:527	128:527	8194.0019e25044e0	1000	FWD	DESG
xe-0/0/9.0	128:529	128:513	4098.0019e2503d20	1000	FWD	ROOT
xe-0/0/11.0	128:531	128:531	8194.0019e25044e0	1000	FWD	DESG

```
user@switch1> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID : 0
Enabled protocol : MSTP
```

```
STP bridge parameters for CIST
```

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/13.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 2000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Message age : 0
Number of topology changes : 3
Time since last topology change : 921 seconds
Local parameters
  Bridge ID : 16384.00:19:e2:50:44:e0
  Extended system ID : 0
  Internal instance ID : 0
```

```
STP bridge parameters for MSTI 1
```

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 2000
Root port : xe-0/0/13.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Local parameters
  Bridge ID : 16385.00:19:e2:50:44:e0
```

```

Extended system ID          : 0
Internal instance ID        : 1

STP bridge parameters for MSTI 2
MSTI regional root          : 4098.00:19:e2:50:3d:20
Root cost                    : 1000
Root port                    : xe-0/0/9.0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Hop count                    : 19
Local parameters
  Bridge ID                  : 8194.00:19:e2:50:44:e0
  Extended system ID         : 0
  Internal instance ID       : 2

```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 2

Purpose Verify the MSTP configuration on Switch 2.

Action Use the operational mode commands:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/14.0	128:513	128:513	32768.0019e2503d20	1000	FWD	DESC
xe-0/0/18.0	128:519	128:515	8192.0019e25051e0	1000	FWD	ROOT

Spanning tree interface parameters for instance 1

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/14.0	128:513	128:513	32769.0019e2503d20	1000	FWD	DESC
xe-0/0/18.0	128:519	128:515	4097.0019e25051e0	1000	FWD	ROOT

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/14.0	128:513	128:513	4098.0019e2503d20	1000	FWD	DESC
xe-0/0/18.0	128:519	128:519	4098.0019e2503d20	1000	FWD	DESC

```
user@switch2> show spanning-tree bridge
```

STP bridge parameters

```
Context ID : 0
Enabled protocol : MSTP
```

STP bridge parameters for CIST

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/18.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 1
Time since last topology change : 782 seconds
Local parameters
  Bridge ID : 32768.00:19:e2:50:3d:20
  Extended system ID : 0
  Internal instance ID : 0
```

STP bridge parameters for MSTI 1

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 1000
Root port : xe-0/0/18.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Local parameters
  Bridge ID : 32769.00:19:e2:50:3d:20
```



```

Extended system ID          : 0
Internal instance ID        : 1

STP bridge parameters for MSTI 2
MSTI regional root          : 4098.00:19:e2:50:3d:20
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Local parameters
  Bridge ID                  : 4098.00:19:e2:50:3d:20
  Extended system ID         : 0
  Internal instance ID       : 2

```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 3

Purpose Verify the MSTP configuration on Switch 3.

Action Use the operational mode commands:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/26.0	128:513	128:513	8192.0019e25051e0	1000	FWD	DESC
xe-0/0/28.0	128:515	128:515	8192.0019e25051e0	1000	FWD	DESC
xe-0/0/24.0	128:517	128:517	8192.0019e25051e0	1000	FWD	DESC

Spanning tree interface parameters for instance 1

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/26.0	128:513	128:513	4097.0019e25051e0	1000	FWD	DESC
xe-0/0/28.0	128:515	128:515	4097.0019e25051e0	1000	FWD	DESC
xe-0/0/24.0	128:517	128:517	4097.0019e25051e0	1000	FWD	DESC

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/26.0	128:513	128:531	8194.0019e25044e0	1000	BLK	ALT
xe-0/0/28.0	128:515	128:519	4098.0019e2503d20	1000	FWD	ROOT
xe-0/0/24.0	128:517	128:517	16386.0019e25051e0	1000	FWD	DESC

```
user@switch3> show spanning-tree bridge
```

STP bridge parameters

```
Context ID          : 0
Enabled protocol    : MSTP
```

STP bridge parameters for CIST

```
Root ID              : 8192.00:19:e2:50:51:e0
CIST regional root   : 8192.00:19:e2:50:51:e0
CIST internal root cost : 0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay         : 15 seconds
Number of topology changes : 3
Time since last topology change : 843 seconds
Local parameters
  Bridge ID           : 8192.00:19:e2:50:51:e0
  Extended system ID   : 0
  Internal instance ID : 0
```

STP bridge parameters for MSTI 1

```
MSTI regional root   : 4097.00:19:e2:50:51:e0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay         : 15 seconds
Local parameters
  Bridge ID           : 4097.00:19:e2:50:51:e0
  Extended system ID   : 0
  Internal instance ID : 1
```

STP bridge parameters for MSTI 2

```
MSTI regional root   : 4098.00:19:e2:50:3d:20
```

```
Root cost           : 1000
Root port           : xe-0/0/28.0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay        : 15 seconds
Hop count            : 19
Local parameters
  Bridge ID          : 16386.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 2
```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 4

Purpose Verify the MSTP configuration on Switch 4.

Action Use the operational mode commands:

```
user@switch4> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/23.0	128:523	128:517	8192.0019e25051e0	1000	FWD	ROOT
xe-0/0/19.0	128:525	128:525	16384.0019e25040e0	1000	FWD	DESG

```
Spanning tree interface parameters for instance 1
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/23.0	128:523	128:517	4097.0019e25051e0	1000	FWD	ROOT
xe-0/0/19.0	128:525	128:525	16385.0019e25040e0	1000	FWD	DESG

```
Spanning tree interface parameters for instance 2
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/23.0	128:523	128:517	16386.0019e25051e0	1000	BLK	ALT
xe-0/0/19.0	128:525	128:527	8194.0019e25044e0	1000	FWD	ROOT

```
user@switch4> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID : 0
Enabled protocol : MSTP
```

```
STP bridge parameters for CIST
```

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/23.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 4
Time since last topology change : 887 seconds
Local parameters
  Bridge ID : 16384.00:19:e2:50:40:e0
  Extended system ID : 0
  Internal instance ID : 0
```

```
STP bridge parameters for MSTI 1
```

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 1000
Root port : xe-0/0/23.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Local parameters
  Bridge ID : 16385.00:19:e2:50:40:e0
  Extended system ID : 0
```

```

Internal instance ID          : 1

STP bridge parameters for MSTI 2
MSTI regional root           : 4098.00:19:e2:50:3d:20
Root cost                     : 2000
Root port                    : xe-0/0/19.0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                 : 15 seconds
Hop count                    : 18
Local parameters
  Bridge ID                   : 32770.00:19:e2:50:40:e0
  Extended system ID         : 0
  Internal instance ID       : 2

```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
 - [Understanding MSTP on page 1901](#)

Example: Configuring Automatic VLAN Administration Using MVRP

As the numbers of servers and VLANs attached to a QFabric systems increase, VLAN administration becomes complex and the task of efficiently configuring VLANs on multiple redundant server Node group devices becomes increasingly difficult. To partially automate VLAN administration, you can enable Multiple VLAN Registration Protocol (MVRP) on your QFabric system. If your QFabric system connects to servers that host many virtual machines that require their own VLANs, using MVRP can save you the time and effort that would be required to manually configure and administer the VLANs on the interfaces that connect to the servers. For example, if a virtual machine moves between servers—and therefore connects to a different redundant server Node group interface—MVRP can configure the appropriate VLAN membership on the new server Node group interface.



NOTE: Only trunk interfaces can be enabled for MVRP.

This example describes how to configure MVRP on a QFabric system.

- [Requirements on page 1954](#)
- [Overview and Topology on page 1954](#)
- [Configuring VLANs and Network Node Group Interfaces on page 1955](#)
- [Configuring the Redundant Server Node Group on page 1956](#)
- [Verification on page 1957](#)

Requirements

This example uses the following hardware and software components:

- One QFabric system
- Junos OS Release 13.1 for the QFX Series

Overview and Topology

MVRP ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs in a static or dynamic VLAN creation setup.

You do not need to explicitly bind a VLAN to the trunk interface. When MVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. An MVRP-enabled trunk interface does not advertise VLANs that have been configured on the switch but that are not currently bound to an access interface. Thus, MVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

When VLAN access interfaces become active or inactive, MVRP ensures that the updated information is advertised on the trunk interface. Thus, in this example, distribution Switch C does not forward traffic to inactive VLANs.

A redundant server Node group device is connected to a server that hosts virtual machines for three customers, each of which requires its own VLAN.

- **customer-1**: VLAN ID 100
- **customer-2**: VLAN ID 200
- **customer-3**: VLAN ID 300

Table 187 on page 1954 explains the components of the example topology.

Table 187: Components of the Example Topology

Settings	Settings
Hardware	<ul style="list-style-type: none"> • Redundant server Node group device • Network Node group device
VLAN names and IDs	<ul style="list-style-type: none"> • customer-1, VLAN ID (tag)100 • customer-2, VLAN ID (tag)200 • customer-3, VLAN ID (tag)300
Interfaces	<p>Redundant server Node group device interfaces:</p> <ul style="list-style-type: none"> • RSNG:xe-0/1/1—Uplink to interconnect device • RSNG:xe-0/0/1—Server-facing interface <p>Network Node group device interface:</p> <ul style="list-style-type: none"> • NNG:xe-0/0/1—Uplink to interconnect device

Configuring VLANs and Network Node Group Interfaces

To configure VLANs, bind the VLANs to the server-facing trunk interface, and enable MVRP on the trunk interface of the network Node group device, perform these tasks:

CLI Quick Configuration To quickly configure VLANs on the QFabric system, assign VLAN membership to the uplink port on the network Node group device, and configure the uplink port to be trunk:

```
[edit]
set vlans customer-1 vlan-id 100
set vlans customer-2 vlan-id 200
set vlans customer-3 vlan-id 300
set interfaces NNG:xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces NNG:xe-0/0/1 unit 0 family ethernet-switching vlan members [customer-1
customer-2 customer-3]
```



NOTE: As recommended as a best practice, default MVRP timers are used in this example, so they are not configured. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Step-by-Step Procedure To create the VLANs and configure the network Node group device for MVRP, follow these steps. Note that you are creating VLANs for the entire QFabric system, so you do not need to create them on specific QFabric devices.

1. Configure the VLAN for customer 1:


```
[edit]
user@qfabric# set vlans customer-1 vlan-id 100
```
2. Configure the VLAN for customer 2:


```
[edit]
user@qfabric# set vlans customer-2 vlan-id 200
```
3. Configure the VLAN for customer 3:


```
[edit]
user@qfabric# set vlans customer-3 vlan-id 300
```
4. Configure an uplink interface (one that connects to an interconnect device) to be a trunk:


```
[edit]
user@qfabric# set interfaces NNG:xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
```
5. Configure the uplink interface to be a member of all three VLANs:


```
[edit]
user@qfabric# set interfaces NNG:xe-0/0/1 unit 1 family ethernet-switching vlan members [customer-1 customer-2 customer-3]
```



NOTE: If you want the uplink interface to be a member of all the VLANs in the QFabric system, you can enter `all` instead of specifying the individual VLANs.

Results Check the results of the configuration on the network Node group device:

```
[edit]
user@qfabric# show interfaces NNG:xe-0/0/1.0
family ethernet-switching {
  port-mode trunk;
  vlan {
    members customer-1 customer-2 customer-3;
  }
}

[edit]
user@qfabric# show vlans
customer-1 {
  vlan-id 100;
}
customer-2 {
  vlan-id 200;
}
customer-3 {
  vlan-id 300;
}
```

Configuring the Redundant Server Node Group

CLI Quick Configuration

To quickly configure the redundant server Node group device for MVRP:

```
[edit]
set interfaces RSNG:xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set interfaces RSNG:xe-0/1/1 unit 0 family ethernet-switching vlan members [customer-1
customer-2 customer-3]
set interfaces RSNG:xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface RSNG:xe-0/0/1.0 passive
```

Step-by-Step Procedure

To configure the redundant server Node group device, follow these steps. Note that you do not need to configure the VLANs on the server-facing interface (RSNG:xe-0/0/1), but you do need to configure the VLANs on the uplink interface. Also notice that in this example you configure the server-facing interface to be passive, which means that it will not announce its membership in a VLAN or send any VLAN declarations (updates) unless it receives registration for that VLAN from the server.

1. Configure an uplink interface (one that connects to the interconnect device) to be a trunk:

```
[edit]
user@qfabric# set interfaces RSNG:xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
```

2. Configure the uplink interface to be a member of all three VLANs:

```
[edit]
user@qfabric# set interfaces NNG:xe-0/1/1 unit 0 family ethernet-switching vlan members [customer-1 customer-2 customer-3]
```

3. Configure an interface that connects to the server that hosts multiple virtual machines to be a trunk:

```
[edit]
```


- ```

user@qfabric# set interfaces RSNG:xe-0/0/1 unit 0 family ethernet-switching port-mode
trunk

```
4. Enable MVRP on the server-facing trunk interface and configure it to be passive:
- ```

[edit]
user@qfabric# set protocols mvrp interface RSNG:xe-0/0/1.0 passive

```

Results Check the results of the configuration for the redundant server Node group:

```

[edit]
user@qfabric# show interfaces RSNG:xe-0/0/1.0
family ethernet-switching {
  port-mode trunk;
}

[edit]
user@qfabric# show interfaces RSNG:xe-0/1/1.0
family ethernet-switching {
  port-mode trunk;
}
passive
}

[edit]
user@qfabric# show protocols mvrp
interface RSNG:xe-0/0/1.0;

```

Verification

To confirm that the configuration is updating VLAN membership, perform these tasks:

- [Verifying That MVRP Is Enabled On The QFabric System on page 1957](#)

Verifying That MVRP Is Enabled On The QFabric System

Purpose Verify that MVRP is enabled on the appropriate interfaces

Action Show the MVRP configuration:

```
user@qfabric> show mvrp
```

```

MVRP configuration
MVRP status           : Enabled

```

```

MVRP timers (ms):
Interface      Join    Leave    LeaveAll
-----
NNG:xe-0/0/1.0  200     1000     10000
RSNG:xe-0/0/1.0  200     1000     10000
RSNG:xe-0/1/1.0  200     1000     10000

```

```

Interface      Status    Registration Mode
-----
NNG:xe-0/0/1.0  Enabled   Normal
RSNG:xe-0/1/1.0  Enabled   Normal
RSNG:xe-0/0/1.0  Enabled   Passive

```

Meaning The results show that MVRP is enabled on the appropriate network Node group and redundant server Node group interfaces and that the default timers are used.

Related Documentation

- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on page 1897](#)

Example: Configuring Reflective Relay for Use with VEPA Technology

Reflective relay must be configured on a switch that receives virtual machine aggregated packets, such as Virtual Ethernet Port Aggregator (VEPA) packets, because some of these packets might be sent back to the server destined for another virtual machine on the same server. Reflective relay returns those packets to the original device using the same downstream port that delivered the packets to the switch.

This example shows how to configure a switch port interface to return packets sent by VEPA on the downstream interface back to the server using the same downstream interface:

- [Requirements on page 1958](#)
- [Overview and Topology on page 1958](#)
- [Configuration on page 1960](#)
- [Verification on page 1961](#)

Requirements

This example uses the following hardware and software components:

- One QFX3500 switch
- One server
- Junos OS Release 12.1 or later for the QFX Series

Before you configure reflective relay on a switch port, be sure you have:

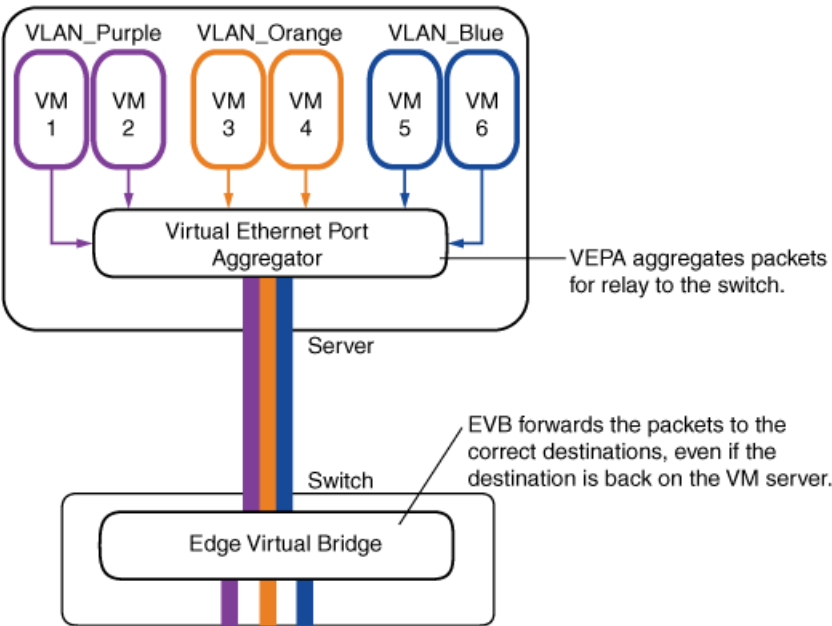
- Configured a server with six virtual machines, VM 1 through VM 6.
- Configured the server with three VLANs named VLAN_Purple, VLAN_Orange, and VLAN_Blue and added two virtual machines to each VLAN.
- Configured the same three VLANs named VLAN_Purple, VLAN_Orange, and VLAN_Blue on one interface.
- Installed and configured VEPA to aggregate the virtual machine packets.

Overview and Topology

In this example, illustrated in [Figure 54 on page 1959](#), a switch is connected to one server that is hosting six virtual machines and is configured with a VEPA for aggregating packets. The server's six virtual machines are VM1 through VM 6, and each virtual machine belongs to one of the three server VLANs, VLAN_Purple, VLAN_Orange, or VLAN_Blue. Instead of the server directly passing packets between virtual machines, packets from any of the three VLANs that are destined for another one of the three VLANs are aggregated using

VEPA technology and passed to the switch for processing. You must configure the switch port to accept these aggregated packets on the downstream interface and to return appropriate packets to the server on the same downstream interface after they are processed. [Figure 54 on page 1959](#) shows the topology for this example.

Figure 54: Reflective Relay Topology



g020996

In this example, you configure the physical Ethernet switch port interface for tagged-access port mode and reflective relay. Configuring tagged-access port mode allows the interface to accept VLAN tagged packets. Configuring reflective relay allows the downstream port to return those packets on the same interface. [Table 188 on page 1959](#) shows the components used in this example.

Table 188: Components of the Topology for Configuring Reflective Relay

Component	Description
QFX3500 switch	Switch that supports reflective relay. For a list of switches that support this feature, see “QFX Series Software Features Overview” on page 15 .
xe-0/0/2	Switch interface to the server.
Server	Server with virtual machines and VEPA technology.
Virtual machines	Six virtual machines located on the server: V1, V2, V3, V4, V5, and V6.
VLANs	Three VLANs: VLAN_Purple, VLAN_Orange, and VLAN_Blue. Each VLAN has two virtual machine members.
VEPA	Virtual Ethernet port aggregator that aggregates virtual machine packets on the server before the resulting single stream is transmitted to the switch.

Configuration

To configure reflective relay, perform these tasks:

- [Configuring Reflective Relay on the Port on page 1960](#)

Configuring Reflective Relay on the Port

CLI Quick Configuration

To quickly configure reflective relay, copy the following commands and paste them into the switch window:

```
[edit]
set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members [VLAN_Blue VLAN_Orange VLAN_Purple]
```

Step-by-Step Procedure

To configure reflective relay:

1. Configure the tagged-access port mode on the interface:



NOTE: Configure the port mode as tagged-access otherwise you will receive an error when you commit the configuration.

- ```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode tagged-access
```
2. Configure reflective relay on the interface to allow it to both accept and send packets:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the three VLANs on the server:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members [VLAN_Purple VLAN_Orange VLAN_Blue]
```

#### Results

Check the results of the configuration:

```
[edit interfaces xe-0/0/2]
user@switch# show
unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 reflective-relay;
 vlan {
 members [VLAN_Purple VLAN_Orange VLAN_Blue];
 }
 }
}
```

## Verification

To confirm that reflective relay is enabled and working correctly, perform these tasks:

- [Verifying That Reflective Relay Is Enabled and Working Correctly on page 1961](#)

### *Verifying That Reflective Relay Is Enabled and Working Correctly*

**Purpose** Verify that reflective relay is enabled and working correctly.

**Action** Use the `show ethernet-switching interfaces detail` command to display the reflective relay status:

```
user@switch> show ethernet-switching interfaces xe-0/0/2 detail
Interface: xe-0/0/2, Index: 66, State: down, Port mode: Tagged-access
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
 VLAN_Purple, 802.1Q Tag: 450, tagged, unblocked
 VLAN_Orange, 802.1Q Tag: 460, tagged, unblocked
 VLAN_Blue, 802.1Q Tag: 470, tagged, unblocked
Number of MACs learned on IFL: 0
```

Confirm that reflective relay is working by sending a Layer 2 broadcast message from one virtual machine to another virtual machine located on the same VLAN. Check the switch to verify that the switch sends the packets back on the same interface on which they were received. One way to check this is to set up port mirroring on the switch interface, connect a traffic generator to the mirrored interface, and use the traffic generator to examine packets.

Alternatively, if you do not have a traffic generator available, you can send traffic between two virtual machines with FTP, Telnet, or SSH, while running the `tcpdump` utility on the receiver virtual machine port to capture reflected packets.

**Meaning** The reflective relay status is **Enabled**, meaning that interface **xe-0/0/2** is configured for the tagged-access port mode, which accepts VLAN-tagged packets, and for reflective relay, which accepts and returns packets on the same interface.

When the traffic generator shows packets arriving at the switch and returning to the server on the same interface, reflective relay is working.

**Related Documentation**

- [Configuring Reflective Relay on page 2044](#)
- [Configuring Port Mirroring on page 4904](#)
- [Configuring VLANs on page 2048](#)
- [port-mode on page 2112](#)
- [reflective-relay on page 2130](#)

Example: Setting Up Basic Bridging and a VLAN on the QFX Series

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices—storage devices, file servers, and other LAN components—in a LAN and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure basic bridging and VLANs for the QFX Series:

- [Requirements on page 1962](#)
- [Overview and Topology on page 1962](#)
- [Configuration on page 1963](#)
- [Verification on page 1969](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX Series
- A configured and provisioned QFX Series product

Overview and Topology

To use a switch to connect network devices on a LAN, you must at a minimum configure bridging and VLANs. By default, bridging is enabled on all switch interfaces, all interfaces are in access mode, and all interfaces belong to a VLAN called **blue**, which is automatically configured. When you plug in access devices—such as desktop computers, file servers, printers, and wireless access points—they are joined immediately into the **blue** VLAN, and the LAN is up and running.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.) You use the ports to connect devices that have their own power sources. Table 1 details the topology used in this configuration example.

Table 189: Components of the Basic Bridging Configuration Topology

| Property        | Settings                                        |
|-----------------|-------------------------------------------------|
| Switch hardware | QFX 3500 switch, with 48 10-Gbps Ethernet ports |
| VLAN name       | <b>blue</b>                                     |

Table 189: Components of the Basic Bridging Configuration Topology (*continued*)

| Property                    | Settings                                                    |
|-----------------------------|-------------------------------------------------------------|
| Connections to file servers | xe-0/0/17 and xe-0/0/18                                     |
| Unused ports                | xe-0/0/0 through xe-0/0/16, and xe-0/0/19 through xe-0/0/47 |

### Configuration

#### CLI Quick Configuration

By default, after you perform the initial configuration on the switch:

- Switching is enabled on all interfaces.
- A VLAN named **blue** is created.
- All interfaces are placed into this VLAN.

You do not need to perform any other configuration on the switch to set up bridging and VLANs. To use the switch, simply connect the servers to the ports **xe-0/0/17** and **xe-0/0/18**.

#### Step-by-Step Procedure

To configure bridging and VLANs:

1. Make sure the switch is powered on.
2. Connect the two file servers to ports **xe-0/0/17** and **xe-0/0/18**.

#### Results

Check the results of the configuration:

```
[edit]
user@switch> show configuration
Last commit: 2010-12-06 00:11:22 UTC by triumph
version 11.1;
system {
 root-authentication {
 encrypted-password "1urmA7AFM$x5SaGEUOdSI3u1K/iITGh1"; ## SECRET-DATA
 }
 syslog {
 user * {
 any emergency;
 }
 file messages {
 any notice;
 authorization info;
 }
 file interactive-commands {
 interactive-commands any;
 }
 }
 commit {
 factory-settings {
 reset-chassis-lcd-menu;
 reset-virtual-chassis-configuration;
```

```
 }
 }
}
interfaces {
 xe-0/0/0 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/1 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/2 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/3 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/4 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/5 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/6 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/7 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/8 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/9 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/10 {
 unit 0 {
```



```
 family ethernet-switching;
 }
}
xe-0/0/11 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/12 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/13 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/14 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/15 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/16 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/17 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/18 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/19 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/20 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/21 {
 unit 0 {
 family ethernet-switching;
```

```
 }
 }
 xe-0/0/22 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/23 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/24 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/25 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/26 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/27 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/28 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/29 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/30 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/31 {
 unit 0 {
 family ethernet-switching;
 }
 }
 xe-0/0/32 {
 unit 0 {
 family ethernet-switching;
 }
 }
```

```
}
xe-0/0/33 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/34 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/35 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/36 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/37 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/38 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/39 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/40 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/41 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/42 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/43 {
 unit 0 {
 family ethernet-switching;
 }
}
```

```
xe-0/0/44 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/45 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/46 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/0/47 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/0 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/0 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/1 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/1 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/2 {
 unit 0 {
 family ethernet-switching;
 }
}
xe-0/1/3 {
 unit 0 {
 family ethernet-switching;
 }
}
}
protocols {
 lldp {
 interface all;
 }
}
rstp;
```

```
}
```

## Verification

To verify that switching is operational and that a VLAN has been created, perform these tasks:

- [Verifying That the VLAN Has Been Created on page 1969](#)
- [Verifying That Interfaces Are Associated with the Proper VLANs on page 1969](#)

### *Verifying That the VLAN Has Been Created*

**Purpose** Verify that the VLAN named **blue** has been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
```

| Name | Tag | Interfaces                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| blue |     | xe-0/0/0.0*, xe-0/0/1.0, xe-0/0/2.0, xe-0/0/3.0,<br>xe-0/0/4.0, xe-0/0/5.0, xe-0/0/6.0, xe-0/0/7.0,<br>xe-0/0/8.0*, xe-0/0/9.0, xe-0/0/10.0, xe-0/0/11.0*,<br>xe-0/0/12.0, xe-0/0/13.0, xe-0/0/14.0, xe-0/0/15.0,<br>xe-0/0/16.0, xe-0/0/17.0, xe-0/0/18.0, xe-0/0/19.0*,<br>xe-0/0/20.0, xe-0/0/21.0, xe-0/0/22.0, xe-0/0/23.0,<br>xe-0/0/24.0, xe-0/0/25.0, xe-0/0/26.0, xe-0/0/27.0,<br>xe-0/0/28.0, xe-0/0/29.0, xe-0/0/30.0, xe-0/0/31.0,<br>xe-0/0/32.0, xe-0/0/33.0, xe-0/0/34.0, xe-0/0/35.0,<br>xe-0/0/36.0, xe-0/0/37.0, xe-0/0/38.0, xe-0/0/39.0,<br>xe-0/0/40.0, xe-0/0/41.0, xe-0/0/42.0, xe-0/0/43.0,<br>xe-0/0/44.0, xe-0/0/45.0, xe-0/0/46.0, xe-0/0/47.0,<br>xe-0/1/0.0*, xe-0/1/1.0*, xe-0/1/2.0*, xe-0/1/3.0* |
| mgmt |     | me0.0*                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Meaning** The **show vlans** command lists the VLANs configured on the switch. This output shows that the VLAN **blue** has been created.

### *Verifying That Interfaces Are Associated with the Proper VLANs*

**Purpose** Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.

**Action** List all interfaces on which switching is enabled:

```
user@switch> show ethernet-switching interfaces
```

| Interface   | State | VLAN members | Blocking                     |
|-------------|-------|--------------|------------------------------|
| xe-0/0/0.0  | up    | blue         | unblocked                    |
| xe-0/0/1.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/2.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/3.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/4.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/5.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/6.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/7.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/8.0  | up    | blue         | unblocked                    |
| xe-0/0/9.0  | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/10.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/11.0 | up    | blue         | unblocked                    |
| xe-0/0/12.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/13.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/14.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/15.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/16.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/17.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/18.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/19.0 | up    | blue         | unblocked                    |
| xe-0/0/20.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/21.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/22.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/23.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/24.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/25.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/26.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/27.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/28.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/29.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/30.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/31.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/32.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/33.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/34.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/35.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/36.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/37.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/38.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/39.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/40.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/41.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/42.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/43.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/44.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/45.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/46.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/0/47.0 | down  | blue         | blocked - blocked by STP/RTG |
| xe-0/1/0.0  | up    | blue         | unblocked                    |
| xe-0/1/1.0  | up    | blue         | unblocked                    |
| xe-0/1/2.0  | up    | blue         | unblocked                    |
| xe-0/1/3.0  | up    | blue         | unblocked                    |
| me0.0       | up    | mgmt         | unblocked                    |

**Meaning** The `show ethernet-switching interfaces` command lists all interfaces on which switching

is enabled (in the Interfaces column), along with the VLANs that are active on the interfaces (in the VLAN members column). Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows **xe-0/0/0.0** instead of **xe-0/0/0**. This is because Junos OS creates VLANs on logical interfaces, not directly on physical interfaces.

- Related Documentation**
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
  - [Understanding Bridging on page 1869](#)

## Example: Setting Up Bridging with Multiple VLANs

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices in a LAN—storage devices, file servers, and other network components—and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure bridging for the QFX Series and how to create two VLANs to segment the LAN:

- [Requirements on page 1971](#)
- [Overview and Topology on page 1971](#)
- [Configuration on page 1972](#)
- [Verification on page 1974](#)

### Requirements

This example uses the following hardware and software components:

- A configured and provisioned QFX3500 switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview and Topology

Switches connect all devices in an office or data center into a single LAN to provide sharing of common resources such as file servers. The default configuration creates a single VLAN, and all traffic on the switch is part of that broadcast domain. Creating separate network segments reduces the span of the broadcast domain and enables you to group related users and network resources without being limited by physical cabling or by the location of a network device in the building or on the LAN.

This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch. One VLAN, called **sales**, is for the sales and marketing group,

and a second, called **support**, is for the customer support team. The sales and support groups each have their own dedicated file servers and other resources. For the switch ports to be segmented across the two VLANs, each VLAN must have its own broadcast domain, identified by a unique name and tag (VLAN ID). In addition, each VLAN must be on its own distinct IP subnet.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.)

**Table 190: Components of the Multiple VLAN Topology**

| Property                          | Settings                                                                                                                                             |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware                   | QFX3500 switch configured with 48 10-Gbps Ethernet ports (xe-0/0/0 through xe-0/0/47)                                                                |
| VLAN names and tag IDs            | <b>sales</b> , tag 100<br><b>support</b> , tag 200                                                                                                   |
| VLAN subnets                      | <b>sales</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126)<br><b>support</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254) |
| Interfaces in VLAN <b>sales</b>   | File servers: xe-0/0/20 and xe-0/0/21                                                                                                                |
| Interfaces in VLAN <b>support</b> | File servers: xe-0/0/46 and xe-0/0/47                                                                                                                |
| Unused interfaces                 | xe-0/0/2 and xe-0/0/25                                                                                                                               |

This configuration example creates two IP subnets, one for the sales VLAN and the second for the support VLAN. The switch bridges traffic within a VLAN. For traffic passing between two VLANs, the switch routes the traffic using a Layer 3 routing interface on which you have configured the address of the IP subnet.

To keep the example simple, the configuration steps show only a few devices in each of the VLANs. Use the same configuration procedure to add more LAN devices.

### Configuration

#### CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**sales** and **support**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/20 unit 0 description "Sales file server port"
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/46 unit 0 description "Support file server port"
```



```

set interfaces xe-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces vlan unit 0 family inet address 192.0.2.0/25
set interfaces vlan unit 1 family inet address 192.0.2.128/25
set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans support vlan-id 200
set vlans support l3-interface vlan.1

```

**Step-by-Step Procedure** Configure the switch interfaces and the VLANs to which they belong. By default, all interfaces are in access mode, so you do not have to configure the port mode.

1. Configure the interface for the file server in the **sales** VLAN:  

```

[edit interfaces xe-0/0/20 unit 0]
user@switch# set description "Sales file server port"
user@switch# set family ethernet-switching vlan members sales

```
2. Configure the interface for the file server in the **support** VLAN:  

```

[edit interfaces xe-0/0/46 unit 0]
user@switch# set description "Support file server port"
user@switch# set family ethernet-switching vlan members support

```
3. Create the subnet for the **sales** broadcast domain:  

```

[edit interfaces]
user@switch# set vlan unit 0 family inet address 192.0.2.1/25

```
4. Create the subnet for the **support** broadcast domain:  

```

[edit interfaces]
user@switch# set vlan unit 1 family inet address 192.0.2.129/25

```
5. Configure the VLAN tag IDs for the **sales** and **support** VLANs:  

```

[edit vlans]
user@switch# set sales vlan-id 100
user@switch# set support vlan-id 200

```
6. To route traffic between the **sales** and **support** VLANs, define the interfaces that are members of each VLAN and associate a Layer 3 interface:  

```

[edit vlans]
user@switch# set sales l3-interface vlan.0
user@switch# set support l3-interface vlan.1

```

Display the results of the configuration:

```

user@switch> show configuration
interfaces {
 xe-0/0/20 {
 unit 0 {
 description "Sales file server port";
 family ethernet-switching {
 vlan members sales;
 }
 }
 }
 xe-0/0/46 {
 unit 0 {
 description "Support file server port";
 family ethernet-switching {
 vlan members support;
 }
 }
 }
}

```

```
}
vlands {
 unit 0 {
 family inet address 192.0.2.1/25;
 }
 unit 1 {
 family inet address 192.0.2.129/25;
 }
}
}
}
vlands {
 sales {
 vlan-id 100;
 interface xe-0/0/0.0;
 interface xe-0/0/3.0;
 interface xe-0/0/20.0;
 interface xe-0/0/22.0;
 l3-interface vlan 0;
 }
 support {
 vlan-id 200;
 interface xe-0/0/24.0;
 interface xe-0/0/26.0;
 interface xe-0/0/44.0;
 interface xe-0/0/46.0;
 l3-interface vlan 1;
 }
}
```



**TIP:** To quickly configure the sales and support VLAN interfaces, issue the `load merge terminal` command. Then copy the hierarchy and paste it into the switch terminal window.

---

## Verification

Verify that the **sales** and **support** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 1974](#)
- [Verifying That Traffic Is Being Routed Between the Two VLANs on page 1975](#)
- [Verifying That Traffic Is Being Switched Between the Two VLANs on page 1975](#)

### ***Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces***

**Purpose** Verify that the **sales** and **support** VLANs have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** To list all VLANs configured on the switch, use the **show vlans** command:

```
user@switch> show vlans
Name Tag Interfaces
default
xe-0/0/1.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/5.0,
xe-0/0/6.0, xe-0/0/7.0, xe-0/0/8.0, xe-0/0/9.0,
xe-0/0/10.0*, xe-0/0/11.0, xe-0/0/12.0, xe-0/0/13.0*,
xe-0/0/14.0, xe-0/0/15.0, xe-0/0/16.0, xe-0/0/17.0,
xe-0/0/18.0, xe-0/0/19.0, xe-0/0/21.0, xe-0/0/23.0*,
xe-0/0/25.0, xe-0/0/27.0, xe-0/0/28.0, xe-0/0/29.0,
xe-0/0/30.0, xe-0/0/31.0, xe-0/0/32.0, xe-0/0/33.0,
xe-0/0/34.0, xe-0/0/35.0, xe-0/0/36.0, xe-0/0/37.0,
xe-0/0/38.0, xe-0/0/39.0, xe-0/0/40.0, xe-0/0/41.0,
xe-0/0/42.0, xe-0/0/43.0, xe-0/0/45.0, xe-0/0/47.0,
xe-0/1/0.0*, xe-0/1/1.0*, xe-0/1/2.0*, xe-0/1/3.0*

sales 100
xe-0/0/0.0*, xe-0/0/3.0, xe-0/0/20.0, xe-0/0/22.0

support 200
xe-0/0/0.24, xe-0/0/26.0, xe-0/0/44.0, xe-0/0/46.0*

mgmt
me0.0*
```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **sales** and **support** VLANs have been created. The **sales** VLAN has a tag ID of 100 and is associated with interfaces **xe-0/0/0.0**, **xe-0/0/3.0**, **xe-0/0/20.0**, and **xe-0/0/22.0**. VLAN **support** has a tag ID of 200 and is associated with interfaces **xe-0/0/24.0**, **xe-0/0/26.0**, **xe-0/0/44.0**, and **xe-0/0/46.0**.

#### *Verifying That Traffic Is Being Routed Between the Two VLANs*

**Purpose** Verify routing between the two VLANs.

**Action** List the Layer 3 routes in the switch Address Resolution Protocol (ARP) table:

```
user@switch> show arp
MAC Address Address Name Flags
00:00:0c:06:2c:0d 192.0.2.3 vlan.0 None
00:13:e2:50:62:e0 192.0.2.11 vlan.1 None
```

**Meaning** Sending IP packets on a multiaccess network requires mapping from an IP address to a MAC address (the physical or hardware address). The ARP table displays the mapping between the IP address and MAC address for both **vlan.0** (associated with **sales**) and **vlan.1** (associated with **support**). These VLANs can route traffic to each other.

#### *Verifying That Traffic Is Being Switched Between the Two VLANs*

**Purpose** Verify that learned entries are being added to the Ethernet switching table.

**Action** List the contents of the Ethernet switching table:

```
user@switch> show ethernet-switching table
```

Ethernet-switching table: 8 entries, 5 learned

| VLAN    | MAC address       | Type  | Age | Interfaces    |
|---------|-------------------|-------|-----|---------------|
| default | *                 | Flood |     | - All-members |
| default | 00:00:05:00:00:01 | Learn |     | - xe-0/0/10.0 |
| default | 00:00:5e:00:01:09 | Learn |     | - xe-0/0/13.0 |
| default | 00:19:e2:50:63:e0 | Learn |     | - xe-0/0/23.0 |
| sales   | *                 | Flood |     | - All-members |
| sales   | 00:00:5e:00:07:09 | Learn |     | - xe-0/0/0.0  |
| support | *                 | Flood |     | - All-members |
| support | 00:00:5e:00:01:01 | Learn |     | - xe-0/0/46.0 |

**Meaning** The output shows that learned entries for the **sales** and **support** VLANs have been added to the Ethernet switching table, and are associated with interfaces **xe-0/0/0.0** and **xe-0/0/46.0**. Even though the VLANs were associated with more than one interface in the configuration, these interfaces are the only ones that are currently operating.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Understanding Bridging on page 1869](#)

## Example: Configuring Routing Between VLANs on One Switch

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs you normally you need a router that connects the VLANs. However, you can accomplish this on a Juniper Networks switch without using a router by configuring a routed VLAN interface (RVI). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

- [Requirements on page 1976](#)
- [Overview and Topology on page 1977](#)
- [Configure Layer 2 switching for two VLANs on page 1978](#)
- [Verification on page 1980](#)

---

### Requirements


This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later

Overview and Topology

This example uses an RVI to route traffic between two VLANs on the same switch. The topology is shown in [Figure 55 on page 1977](#).

Figure 55: RVI with One Switch



This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch and configuring an RVI to enable routing between the VLANs. One VLAN, called **blue**, is for the sales and marketing group, and a second, called **red**, is for the customer support team. The sales and support groups each have their own file servers and wireless access points. Each VLAN must have a unique name, tag (VLAN ID), and distinct IP subnet. [Table 191 on page 1977](#) lists the components of the sample topology.

Table 191: Components of the Multiple VLAN Topology

| Property                       | Settings                                                                                                                                        |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN names and tag IDs         | <b>blue</b> , ID 100<br><b>red</b> , ID 200                                                                                                     |
| Subnets associated with VLANs  | <b>blue</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126)<br><b>red</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254) |
| Interfaces in VLAN <b>blue</b> | Sales server port: <b>xe-0/0/4</b><br>Sales wireless access points: <b>xe-0/0/6</b>                                                             |
| Interfaces in VLAN <b>red</b>  | Support server port: <b>xe-0/0/0</b><br>Support wireless access points: <b>xe-0/0/2</b>                                                         |
| RVI name                       | interface <b>vlan</b>                                                                                                                           |
| RVI units and addresses        | logical unit 100: <b>192.0.2.1/25</b><br><br>logical unit 200: <b>192.0.2.129/25</b>                                                            |

This configuration example creates two IP subnets, one for the blue VLAN and the second for the red VLAN. The switch bridges traffic within the VLANs. For traffic passing between two VLANs, the switch routes the traffic using an RVI on which you have configured addresses in each IP subnet.

To keep the example simple, the configuration steps show only a few interfaces and VLANs. Use the same configuration procedure to add more interfaces and VLANs. By default, all interfaces are in access mode, so you do not have to configure the port mode.

## Configure Layer 2 switching for two VLANs

**CLI Quick Configuration** To quickly configure Layer 2 switching for the two VLANs (**blue** and **red**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces xe-0/0/4 unit 0 description "Sales server port"
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/6 unit 0 description "Sales wireless access point port"
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/0 unit 0 description "Support servers"
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members red
set interfaces xe-0/0/2 unit 0 description "Support wireless access point port"
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members red
set interfaces vlan unit 100 family inet address 192.0.2.1/25
set interfaces vlan unit 200 family inet address 192.0.2.129/25
set vlans blue l3-interface vlan.100
set vlans blue vlan-id 100
set vlans red vlan-id 200
set vlans red l3-interface vlan.200
```

**Step-by-Step Procedure** To configure the switch interfaces and the VLANs to which they belong:

1. Configure the interface for the sales server in the blue VLAN:  

```
[edit interfaces xe-0/0/4 unit 0]
user@switch# set description "Sales server port"
user@switch# set family ethernet-switching vlan members blue
```
2. Configure the interface for the wireless access point in the blue VLAN:  

```
[edit interfaces xe-0/0/6 unit 0]
user@switch# set description "Sales wireless access point port"
user@switch# set family ethernet-switching vlan members blue
```
3. Configure the interface for the support server in the red VLAN:  

```
[edit interfaces xe-0/0/0 unit 0]
user@switch# set description "Support server port"
user@switch# set family ethernet-switching vlan members red
```
4. Configure the interface for the wireless access point in the red VLAN:  

```
[edit interfaces xe-0/0/2 unit 0]
user@switch# set description "Support wireless access point port"
user@switch# set family ethernet-switching vlan members red
```

**Step-by-Step Procedure** Now create the VLANs and the RVI. The RVI will have logical units in the broadcast domains of both VLANs.

1. Create the red and blue VLANs by configuring the VLAN IDs for them:  

```
[edit vlans]
user@switch# set blue vlan-id 100
user@switch# set red vlan-id 200
```
2. Create the interface named **vlan** with a logical unit in the sales broadcast domain (blue VLAN):  

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 192.0.2.1/25
```

The unit number is arbitrary and does not have to match the VLAN tag ID. However, configuring the unit number to match the VLAN ID can help avoid confusion.

3. Add a logical unit in the support broadcast domain (red VLAN) to the **vlan** interface:

```
[edit interfaces]
```

```
user@switch# set vlan unit 200 family inet address 192.0.2.129/25
```

4. Complete the RVI configuration by binding the red and blue VLANs (Layer 2) with the appropriate logical units of the **vlan** interface (Layer 3):

```
[edit vlans]
```

```
user@switch# set blue l3-interface vlan.100
```

```
user@switch# set red l3-interface vlan.200
```

Display the results of the configuration:

```
user@switch> show configuration
```

```
interfaces {
 xe-0/0/4 {
 unit 0 {
 description "Sales server port";
 family ethernet-switching {
 vlan members blue;
 }
 }
 }
 xe-0/0/6 {
 unit 0 {
 description "Sales wireless access point port";
 family ethernet-switching {
 vlan members blue;
 }
 }
 }
 xe-0/0/0 {
 unit 0 {
 description "Support server port";
 family ethernet-switching {
 vlan members red;
 }
 }
 }
 xe-0/0/2 {
 unit 0 {
 description "Support wireless access point port";
 family ethernet-switching {
 vlan members red;
 }
 }
 }
}
vlan {
 unit 100 {
 family inet address 192.0.2.1/25;
 }
 unit 200 {
 family inet address 192.0.2.129/25;
 }
}
```

```

 }
 }
}
vllans {
 blue {
 vlan-id 100;
 interface xe-0/0/4.0;
 interface xe-0/0/6.0;
 l3-interface vlan 100;
 }
 red {
 vlan-id 200;
 interface xe-0/0/0.0;
 interface xe-0/0/2.0;
 l3-interface vlan 200;
 }
}

```



**TIP:** To quickly configure the blue and red VLAN interfaces, issue the **load merge terminal** command, copy the hierarchy, and paste it into the switch terminal window.

## Verification

To verify that the **blue** and **red** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 1980](#)
- [Verifying That Traffic Can Be Routed Between the Two VLANs on page 1981](#)

### *Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces*

**Purpose** Verify that the VLANs **blue** and **red** have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** List all VLANs configured on the switch:

```

user@switch> show vlans
Name Tag Interfaces
default 0 xe-0/0/0.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/6.0,
blue 100 xe-0/0/4.0, xe-0/0/6,
red 200 xe-0/0/0.0, xe-0/0/2.0, *
mgmt 0 me0.0*

```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **blue** and **red** VLANs have been created. The **blue** VLAN has a tag ID of 100 and is associated with interfaces



**xe-0/0/4.0** and **xe-0/0/6.0**. VLAN **red** has a tag ID of 200 and is associated with interfaces **xe-0/0/0.0** and **xe-0/0/2.0**.

### *Verifying That Traffic Can Be Routed Between the Two VLANs*

**Purpose** Verify routing between the two VLANs.

**Action** Verify that the RVI logical units are up:

```
user@switch> show interfaces terse
vlan.100 up up inet 192.0.2.1/25
vlan.200 up up inet 192.0.2.129/25
```



**NOTE:** At least one port (access or trunk) with an appropriate VLAN assigned to it must be up for the vlan interface to be up.

Verify that switch has created routes that use the RVI logical units:

```
user@switch> show route
192.0.2.0/25 *[Direct/0] 1d 03:26:45
 > via vlan.100
192.0.2.1/32 *[Local/0] 1d 03:26:45
 Local via vlan.100
192.0.2.128/25 *[Direct/0] 1d 03:26:45
 > via vlan.200
192.0.2.129/32 *[Local/0] 1d 03:26:45
 Local via vlan.200
```

List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```
user@switch> show arp
```

| MAC Address       | Address     | Name     | Flags |
|-------------------|-------------|----------|-------|
| 00:00:0c:06:2c:0d | 192.0.2.7   | vlan.100 | None  |
| 00:13:e2:50:62:e0 | 192.0.2.132 | vlan.200 | None  |

**Meaning** The output of the **show interfaces** and **show route** commands show that the Layer 3 RVI logical units are working and the switch has used them to create direct routes that it will use to forward traffic between the VLAN subnets. The **show arp** command displays the mappings between the IP addresses and MAC addresses for devices on both **vlan.100** (associated with VLAN **blue**) and **vlan.200** (associated with VLAN **red**). These two devices can communicate.

**Related Documentation**

- [Understanding Routed VLAN Interfaces on page 1875](#)
- [I3-interface on page 2098](#)

## Example: Connecting an Access Switch to a Distribution Switch

In large local area networks (LANs), you commonly need to aggregate traffic from a number of access switches into a distribution switch.

This example describes how to connect an access switch to a distribution switch:

- [Requirements on page 1982](#)
- [Overview and Topology on page 1982](#)
- [Configuring the Access Switch on page 1983](#)
- [Configuring the Distribution Switch on page 1987](#)
- [Verification on page 1989](#)

### Requirements

This example uses the following hardware and software components:

- For the distribution switch, one EX 4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an EX-UM-2XFP uplink module with two 10-Gigabit Ethernet XFP ports.
- For the access switch, one EX 3200-24P, which has twenty-four 1-Gigabit Ethernet ports, all of which support Power over Ethernet (PoE), and an uplink module with four 1-Gigabit Ethernet ports.
- Junos OS Release 11.1 or later for the QFX Series

### Overview and Topology

In a large office that is spread across several floors or buildings, or in a data center, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single access switch to a distribution switch.

In the topology, the LAN is segmented into two VLANs, one for the sales department and the second for the support team. One 1-Gigabit Ethernet port on the access switch's uplink module connects to the distribution switch, to one 1-Gigabit Ethernet port on the distribution switch.

[Table 192 on page 1982](#) explains the components of the example topology. The example shows how to configure one of the three access switches. The other access switches could be configured in the same manner.

**Table 192: Components of the Topology for Connecting an Access Switch to a Distribution Switch**

| Property               | Settings                                                                                                                                                         |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Access switch hardware | EX 3200-24P, 24 1-Gigabit Ethernet ports, all PoE-enabled ( <b>ge-0/0/0</b> through <b>ge-0/0/23</b> ); one 4-port 1-Gigabit Ethernet uplink module (EX-UM-4SFP) |

**Table 192: Components of the Topology for Connecting an Access Switch to a Distribution Switch (*continued*)**

|                                                                  |                                                                                                                                                                                                                               |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Distribution switch hardware                                     | EX 4200-24F, 24 1-Gigabit Ethernet fiber SPF ports ( <b>ge-0/0/0</b> through <b>ge-0/0/23</b> ); one 2-port 10-Gigabit Ethernet XFP uplink module (EX-UM-4SFP)                                                                |
| VLAN names and tag IDs                                           | <b>sales</b> , tag 100 <b>support</b> , tag 200                                                                                                                                                                               |
| VLAN subnets                                                     | <b>sales</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>support</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)                                                                             |
| Trunk port interfaces                                            | On the access switch: <b>ge-0/1/0</b> On the distribution switch: <b>ge-0/0/0</b>                                                                                                                                             |
| Access port interfaces in VLAN <b>sales</b> (on access switch)   | Avaya IP telephones: <b>ge-0/0/3</b> through <b>ge-0/0/19</b> Wireless access points: <b>ge-0/0/0</b> and <b>ge-0/0/1</b> Printers: <b>ge-0/0/22</b> and <b>ge-0/0/23</b> File servers: <b>ge-0/0/20</b> and <b>ge-0/0/21</b> |
| Access port interfaces in VLAN <b>support</b> (on access switch) | Avaya IP telephones: <b>ge-0/0/25</b> through <b>ge-0/0/43</b> Wireless access points: <b>ge-0/0/24</b> Printers: <b>ge-0/0/44</b> and <b>ge-0/0/45</b> File servers: <b>ge-0/0/46</b> and <b>ge-0/0/47</b>                   |
| Unused interfaces on access switch                               | <b>ge-0/0/2</b> and <b>ge-0/0/25</b>                                                                                                                                                                                          |

### Configuring the Access Switch

To configure the access switch:

#### CLI Quick Configuration

To quickly configure the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 description "Sales Wireless access point port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description "Sales phone port"
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description "Sales printer port"
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description "Sales file server port"
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/24 unit 0 description "Support wireless access point port"
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/26 unit 0 description "Support phone port"
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description "Support printer port"
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description "Support file server port"
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/1/0 unit 0 description "Uplink module port connection to distribution switch"
set interfaces ge-0/1/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/1/0 unit 0 family ethernet-switching native-vlan-id 1
set interfaces ge-0/1/0 unit 0 family ethernet-switching vlan members [sales support]
set interfaces vlan unit 0 family inet address 192.0.2.1/25
set interfaces vlan unit 1 family inet address 192.0.2.129/25
set vlans sales interface ge-0/0/0.0
set vlans sales interface ge-0/0/3.0
set vlans sales interface ge-0/0/22.0
set vlans sales interface ge-0/0/20.0
```

```

set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans sales vlan-description "Sales VLAN"
set vlans support interface ge-0/0/24.0
set vlans support interface ge-0/0/26.0
set vlans support interface ge-0/0/44.0
set vlans support interface ge-0/0/46.0
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
set vlans support vlan-description "Support VLAN"

```

### Step-by-Step Procedure

To configure the access switch:

1. Configure the 1-Gigabit Ethernet interface on the uplink module to be the trunk port that connects to the distribution switch:

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set description "Uplink
module port connection to distribution switch"user@access-switch# set ethernet-switching
port-mode trunk
```

2. Specify the VLANs to be aggregated on the trunk port:

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set ethernet-switching
vlan-members [sales support]
```

3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set ethernet-switching
native-vlan-id 1
```

4. Configure the sales VLAN:

```
[edit vlans sales]user@access-switch# set vlan-description "Sales
VLAN"user@access-switch# set vlan-id (VLANs) 100user@access-switch# set l3-interface
(VLAN) vlan.0
```

5. Configure the support VLAN:

```
[edit vlans support]user@access-switch# set vlan-description "Support
VLAN"user@access-switch# set vlan-id (VLANs) 200user@access-switch# set
l3-interface (VLAN) vlan.1
```

6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]user@access-switch# set vlan unit 0 family inet address 192.0.2.1/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces]user@access-switch# set vlan unit 1 family inet address 192.0.2.129/25
```

8. Configure the interfaces in the sales VLAN:

```
[edit interfaces]user@access-switch# set ge-0/0/0 unit 0 description "Sales wireless
access point port"user@access-switch# set ge-0/0/0 unit 0 family ethernet-switching
vlan members salesuser@access-switch# set ge-0/0/3 unit 0 description "Sales phone
port"user@access-switch# set ge-0/0/3 unit 0 family ethernet-switching vlan members
salesuser@access-switch# set ge-0/0/20 unit 0 description "Sales file server
port"user@access-switch# set ge-0/0/20 unit 0 family ethernet-switching vlan members
salesuser@access-switch# set ge-0/0/22 unit 0 description "Sales printer
port"user@access-switch# set ge-0/0/22 unit 0 family ethernet-switching vlan members
sales
```

9. Configure the interfaces in the support VLAN:

```
[edit interfaces]user@access-switch# set ge-0/0/24 unit 0 description "Support
wireless access point port"user@access-switch# set ge-0/0/24 unit 0 family
ethernet-switching vlan members supportuser@access-switch# set ge-0/0/26 unit 0
description "Support phone port"user@access-switch# set ge-0/0/26 unit 0 family
```

```

ethernet-switching vlan members support
user@access-switch# set ge-0/0/44 unit 0
description "Support printer port"
user@access-switch# set ge-0/0/44 unit 0 family
ethernet-switching vlan members support
user@access-switch# set ge-0/0/46 unit 0
description "Support file server port"
user@access-switch# set ge-0/0/46 unit 0 family
ethernet-switching vlan members support

```

10. Configure descriptions and VLAN tag IDs for the sales and support VLANs:

```

[edit vlans]user@access-switch# set sales vlan-description "Sales
VLAN"
user@access-switch# set sales vlan-id 100
user@access-switch# set support
vlan-description "Support VLAN"
user@access-switch# set support vlan-id 200

```

11. To route traffic between the sales and support VLANs and associate a Layer 3 interface with each VLAN:

```

[edit vlans]user@access-switch# set sales l3-interface vlan.0
user@access-switch# set support l3-interface vlan.1

```

**Results** Display the results of the configuration:

```

user@access-switch> show
interfaces {
 ge-0/0/0 {
 unit 0 {
 description "Sales wireless access point port";
 family ethernet-switching {
 vlan members sales;
 }
 }
 }
 ge-0/0/3 {
 unit 0 {
 description "Sales phone port";
 family ethernet-switching {
 vlan members sales;
 }
 }
 }
 ge-0/0/20 {
 unit 0 {
 description "Sales file server port";
 family ethernet-switching {
 vlan members sales;
 }
 }
 }
 ge-0/0/22 {
 unit 0 {
 description "Sales printer port";
 family ethernet-switching {
 vlan members sales;
 }
 }
 }
 ge-0/0/24 {
 unit 0 {
 description "Support wireless access point port";
 family ethernet-switching {
 vlan members support;
 }
 }
 }
}

```

```
 }
 }
}
ge-0/0/26 {
 unit 0 {
 description "Support phone port";
 family ethernet-switching {
 vlan members support;
 }
 }
}
ge-0/0/44 {
 unit 0 {
 description "Support printer port";
 family ethernet-switching {
 vlan members sales;
 }
 }
}
ge-0/0/46 {
 unit 0 {
 description "Support file server port";
 family ethernet-switching {
 vlan members support;
 }
 }
}
ge-0/1/0 {
 unit 0 {
 description "Uplink module port connection to distribution switch";
 family ethernet-switching {
 port-mode trunk;
 vlan members [sales support];
 native-vlan-id 1;
 }
 }
}
vlan {
 unit 0 {
 family inet address 192.0.2.1/25;
 }
 unit 1 {
 family inet address 192.0.2.129/25;
 }
}
vlands {
 sales {
 vlan-id 100;
 vlan-description "Sales VLAN";
 l3-interface vlan.0;
 }
 support {
 vlan-id 200;
 vlan-description "Support VLAN";
 l3-interface vlan.1;
 }
}
```

```
}
}
```



**TIP:** To quickly configure the distribution switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

## Configuring the Distribution Switch

To configure the distribution switch:

### CLI Quick Configuration

To quickly configure the distribution switch, copy the following commands and paste them into the switch terminal window:

```
set interfaces ge-0/0/0 description "Connection to access switch"
set interfaces ge-0/0/0 ethernet-switching port-mode trunk
set interfaces ge-0/0/0 ethernet-switching vlan members [sales support]
set interfaces ge-0/0/0 ethernet-switching native-vlan-id 1
set interfaces vlan unit 0 family inet address 192.0.2.2/25
set interfaces vlan unit 1 family inet address 192.0.2.130/25
set vlans sales vlan-description "Sales VLAN"
set vlans sales vlan-id 100
set vlans sales l3-interface vlan.0
set vlans support vlan-description "Support VLAN"
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
```

### Step-by-Step Procedure

To configure the distribution switch:

1. Configure the interface on the switch to be the trunk port that connects to the access switch:  

```
[edit interfaces ge-0/0/0 unit 0]user@distribution-switch# set description "Connection to access switch"
user@distribution-switch# set ethernet-switching port-mode trunk
```
2. Specify the VLANs to be aggregated on the trunk port:  

```
[edit interfaces ge-0/0/0 unit 0]user@distribution-switch# set ethernet-switching vlanmembers [sales support]
```
3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):  

```
[edit interfaces]user@distribution-switch# set ge-0/0/0 ethernet-switching native-vlan-id 1
```
4. Configure the sales VLAN:  

```
[edit vlans sales]user@distribution-switch# set vlan-description "Sales VLAN"
user@distribution-switch# set vlan-id (VLANs) 100
user@distribution-switch# set l3-interface (VLAN) vlan.0
```
5. Configure the support VLAN:  

```
[edit vlans support]user@distribution-switch# set vlan-description "Support VLAN"
user@distribution-switch# set vlan-id (VLANs) 200
user@distribution-switch# set l3-interface (VLAN) vlan.1
```
6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]user@distribution-switch# set vlan unit 0 family inet address
192.0.2.2/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces] user@distribution-switch# set vlan unit 1 family inet address
192.0.2.130/25
```

**Results** Display the results of the configuration:

```
user@distribution-switch> show
interfaces {
 ge-0/0/0 {
 description "Connection to access switch";
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan members [sales support];
 native-vlan-id 1;
 }
 }
 }
}
vlan {
 unit 0 {
 family inet address 192.0.2.2/25;
 }
 unit 1 {
 family inet address 192.0.2.130/25;
 }
}
vpls {
 sales {
 vlan-id 100;
 vlan-description "Sales VLAN";
 l3-interface vlan.0;
 }
 support {
 vlan-id 200;
 vlan-description "Support VLAN";
 l3-interface vlan.1;
 }
}
```



**TIP:** To quickly configure the distribution switch, issue the `load merge terminal` command, then copy the hierarchy and paste it into the switch terminal window.



## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the VLAN Members and Interfaces on the Access Switch on page 1989](#)
- [Verifying the VLAN Members and Interfaces on the Distribution Switch on page 1989](#)

### *Verifying the VLAN Members and Interfaces on the Access Switch*

**Purpose** Verify that the **sales** and **support** have been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
```

| Name    | Tag | Interfaces                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| default |     | ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0,<br>ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0,<br><br>ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0,<br>ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0,<br>ge-0/0/18.0, ge-0/0/19.0*, ge-0/0/21.0, ge-0/0/23.0,<br>ge-0/0/25.0, ge-0/0/27.0*, ge-0/0/28.0, ge-0/0/29.0,<br>ge-0/0/30.0, ge-0/0/31.0*, ge-0/0/32.0, ge-0/0/33.0,<br>ge-0/0/34.0, ge-0/0/35.0*, ge-0/0/36.0, ge-0/0/37.0,<br>ge-0/0/38.0, ge-0/0/39.0*, ge-0/0/40.0, ge-0/0/41.0,<br>ge-0/0/42.0, ge-0/0/43.0*, ge-0/0/45.0, ge-0/0/47.0,<br>ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0* |
| sales   | 100 | ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0,<br>ge-0/1/0.0*,                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| support | 200 | ge-0/0/24.0*, ge-0/0/26.0, ge-0/0/44.0, ge-0/0/46.0,                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| mgmt    |     | me0.0*                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Meaning** The output shows the **sales** and **support** VLANs and the interfaces associated with them.

### *Verifying the VLAN Members and Interfaces on the Distribution Switch*

**Purpose** Verify that the **sales** and **support** have been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
```

| Name    | Tag | Interfaces                                                                                                                                                                                                                                                                                                                                                        |
|---------|-----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| default |     | ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,<br>ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0*, ge-0/0/8.0,<br><br>ge-0/0/9.0, ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0,<br>ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0,<br>ge-0/0/17.0, ge-0/0/18.0*, ge-0/0/19.0, ge-0/0/20.0,<br>ge-0/0/21.0, ge-0/0/22.0*, ge-0/0/23.0, ge-0/1/1.0*,<br>ge-0/1/2.0*, ge-0/1/3.0* |
| sales   | 100 | ge-0/0/0.0*                                                                                                                                                                                                                                                                                                                                                       |
| support | 200 | ge-0/0/0.0*                                                                                                                                                                                                                                                                                                                                                       |
| mgmt    |     | me0.0*                                                                                                                                                                                                                                                                                                                                                            |

**Meaning** The output shows the **sales** and **support** VLANs associated to interface **ge-0/0/0.0**. Interface **ge-0/0/0.0** is the trunk interface connected to the access switch.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Understanding Bridging on page 1869](#)

## Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations

The QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Configure BPDU protection on interfaces to prevent them from receiving BPDUs that could result in STP misconfigurations, which could lead to network outages.

This example describes how to configure BPDU protection on access interfaces in QFX Series products in an RSTP topology:

- [Requirements on page 1990](#)
- [Overview and Topology on page 1991](#)
- [Configuration on page 1992](#)
- [Verification on page 1992](#)

---

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series

- Two edged-linked switches in an RSTP topology



**NOTE:** By default, RSTP is enabled on the QFX Series.

### Overview and Topology


A loop-free network is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). However, receipt of BPDUs on certain interfaces in an STP, RSTP, or MSTP topology. It can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on those interfaces that should not receive BPDUs.

Enable BPDU protection on switch interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If a BPDU is received on a BPDU-protected interface, the interface is disabled and stops forwarding frames.

Two switches are displayed in [Figure 56 on page 1991](#). In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are access ports.

This example shows you how to configure interface **xe-0/0/5** and interface **xe-0/0/6** as edge ports and how to configure BPDU protection. When BPDU protection is enabled, the interfaces transition to a blocking state when they receive BPDUs.

**Figure 56: BPDU Protection Topology**



g020153

[Table 193 on page 1992](#) shows the components that will be configured for BPDU protection.

**Table 193: Components of the Topology for Configuring BPDU Protection on the QFX Series**

| Component                     | Settings                                                                                                                                       |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 (Distribution Layer) | Switch 1 is connected to Switch 2 on a trunk interface.                                                                                        |
| Switch 2 (Access Layer)       | Switch 2 has these access ports that require BPDU protection: <ul style="list-style-type: none"> <li>• xe-0/0/5</li> <li>• xe-0/0/6</li> </ul> |

This configuration example uses an RSTP topology. You also can configure BPDU protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

### Configuration

#### CLI Quick Configuration

To quickly configure BPDU protection on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp interface xe-0/0/5 edge
set protocols rstp interface xe-0/0/6 edge
set protocols rstp bpdu-block-on-edge
```

#### Step-by-Step Procedure

To configure BPDU protection:

1. Configure interface `xe-0/0/5` and interface `xe-0/0/6` on Switch 2 as edge ports:

```
[edit protocols rstp]
user@switch# set interface xe-0/0/5 edge
user@switch# set interface xe-0/0/6 edge
```

2. Configure BPDU protection on all edge ports:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/5.0 {
 edge;
}
interface xe-0/0/6.0 {
 edge;
}
bpdu-block-on-edge;
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before BPDU Protection Is Triggered on page 1993](#)
- [Verifying That BPDU Protection Is Working Correctly on page 1993](#)

*Displaying the Interface State Before BPDU Protection Is Triggered*

**Purpose** Before BPDUs are being received from the devices connected to interface **xe-0/0/5** and interface **xe-0/0/6**, confirm the interface state.

**Action** You can verify the interface state using the **show spanning-tree interface** command:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/6.0 | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output shows that interface **xe-0/0/5.0** and interface **xe-0/0/6.0** are designated ports in a forwarding state.

*Verifying That BPDU Protection Is Working Correctly*

**Purpose** In this example, the devices connected to Switch 2 start sending BPDUs to interface **xe-0/0/5.0** and interface **xe-0/0/6.0**. Verify that BPDU protection is configured on the interfaces.

**Action** You can verify that BPDU protection is configured on the interfaces by using the **show spanning-tree interface** command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface    | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|--------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0   | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0   | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0   | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0   | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0   | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0   | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| (Bpdu-Incon) |         |                       |                         |              |       |      |
| xe-0/0/6.0   | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| (Bpdu-Incon) |         |                       |                         |              |       |      |
| xe-0/0/7.0   | 128:520 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| xe-0/0/8.0   | 128:521 | 128:521               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** When BPDUs are sent from the devices to interface **xe-0/0/5.0** and interface **xe-0/0/6.0** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state blocks the interfaces and prevents them from forwarding traffic.

Disabling the BPDU protection configuration on an interface does not unblock the interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. Otherwise, use the operational mode command **clear ethernet-switching bpd-error** to unblock the interface.

If the devices connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state. In such cases, you need to find and repair the misconfiguration on the devices that is triggering the sending of BPDUs to Switch 2.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
  - [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999](#)
  - [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1994](#)
  - [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903](#)

## Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees

QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol

(MSTP). Root protection increases the efficiency of STP, RSTP, and MSTP by allowing network administrators to enforce the root bridge placement in the network manually.

This example describes how to configure root protection on an interface for the QFX Series.

- [Requirements on page 1995](#)
- [Overview and Topology on page 1995](#)
- [Configuration on page 1997](#)
- [Verification on page 1997](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Four switches in an RSTP topology

Before you configure the interface for root protection, be sure you have:

- RSTP operating on the switches.



**NOTE:** By default, RSTP is enabled on the QFX Series.

## Overview and Topology

Peer STP applications running on switch interfaces exchange a special type of frame called a bridge protocol data unit (BPDU). Switches communicate interface information using BPDUs to create a loop-free topology that ultimately determines the root bridge and which interfaces block or forward traffic in the spanning tree.

You can also see BPDUs generated when you run a bridge application on a device attached to the switch. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

To prevent this from happening, enable root protection on interfaces that should not receive more BPDUs from the root bridge and should not be elected as the root port. These interfaces are typically located on an administrative boundary and are designated ports.

When root protection is enabled on an interface:

- The interface is blocked from becoming the root port.
- Root protection is enabled for all STP instances on that interface.
- The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.




**NOTE:** An interface can be configured for either root protection or loop protection, but not for both.

Four switches are displayed in [Figure 57 on page 1996](#). In this example, they are configured for RSTP and create a loop-free topology. Interface `xe-0/0/7` on Switch 1 is a designated port on an administrative boundary. It connects to Switch 4. Switch 3 is the root bridge. Interface `xe-0/0/6` on Switch 1 is the root port.

This example shows how to configure root protection on interface `xe-0/0/7` to prevent it from transitioning to become the root port.

Figure 57: Network Topology for Root Protection



[Table 194 on page 1996](#) shows the components that will be configured for root protection.

Table 194: Topology for Configuring Root Protection on the QFX Series

| Component | Settings                                                                    |
|-----------|-----------------------------------------------------------------------------|
| Switch 1  | Switch 1 is connected to Switch 4 through interface <code>xe-0/0/7</code> . |



Table 194: Topology for Configuring Root Protection on the QFX Series (*continued*)

| Component | Settings                                                                                                                                                                                    |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 2  | Switch 2 is connected to Switch 1 and Switch 3. Interface <b>xe-0/0/4</b> is the alternate port in the RSTP topology.                                                                       |
| Switch 3  | Switch 3 is the root bridge and is connected to Switch 1 and Switch 2.                                                                                                                      |
| Switch 4  | Switch 4 is connected to Switch 1. After loop protection is configured on interface <b>xe-0/0/7</b> , Switch 4 sends more BPDUs that trigger loop protection on interface <b>xe-0/0/7</b> . |

A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you can also configure root protection for STP or MSTP topologies at the **[edit protocols (mstp | stp)]** hierarchy level.

### Configuration

#### CLI Quick Configuration

To quickly configure root protection on interface **xe-0/0/7**, copy the following command and paste it into the switch terminal window:

```
[edit]
set protocols rstp interface xe-0/0/7 no-root-port
```

#### Step-by-Step Procedure

To configure root protection:

1. Configure interface **xe-0/0/7**:  

```
[edit protocols rstp]
user@switch#
set interface xe-0/0/7 no-root-port
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/7.0 {
 no-root-port;
}
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Root Protection Is Triggered on page 1998](#)
- [Verifying That Root Protection Is Working on the Interface on page 1998](#)

**Displaying the Interface State Before Root Protection Is Triggered**

**Purpose** Before root protection is triggered on interface **xe-0/0/7**, confirm the interface state.

**Action** Confirm the state of the interfaces before root protection is configured:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| xe-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| xe-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/7.0** is a designated port in a forwarding state.

**Verifying That Root Protection Is Working on the Interface**

**Purpose** A configuration change takes place on Switch 4. A lower bridge priority on Switch 4 causes it to send more BPDUs to interface **xe-0/0/7**. Receipt of more BPDUs on interface **xe-0/0/7** triggers root protection. Verify that root protection is operating on interface **xe-0/0/7**.

**Action** Verify that root protection has been configured and is operating correctly:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| xe-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| xe-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Root-Incon)  
[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/7.0** has transitioned to a loop inconsistent state. The loop inconsistent state blocks the interface and prevents it from becoming a candidate for the root port.

When the root bridge no longer receives more STP BPDUs from the interface, the interface recovers and transitions back to a forwarding state. Recovery is automatic.

**Related Documentation**

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1905](#)

## Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree

The QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing interfaces from moving into a forwarding state that would create a loop in the network.

This example describes how to configure loop protection for an interface for the QFX Series in an RSTP topology:

- [Requirements on page 1999](#)
- [Overview and Topology on page 1999](#)
- [Configuration on page 2001](#)
- [Verification on page 2001](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Three switches in an RSTP topology



**NOTE:** By default, RSTP is enabled for the QFX Series.

### Overview and Topology

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

A blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can

occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor. When this happens, a loop appears in the spanning tree. Loops in a Layer 2 topology cause broadcast, unicast, and multicast frames to continuously circle the looped network. As a switch processes a flood of frames in a looped network, its resources become depleted, and the ultimate result is a network outage.




**NOTE:** An interface can be configured for either loop protection or root protection, but not for both.

Three switches are displayed in [Figure 58 on page 2000](#). In this example, they are configured for RSTP and create a loop-free topology. Interface **xe-0/0/6** is blocking traffic between Switch 3 and Switch 1; thus, traffic is forwarded through interface **xe-0/0/7** on Switch 2. BPDUs are being sent from the root bridge on Switch 1 to both of these interfaces.

This example shows how to configure loop protection on interface **xe-0/0/6** to prevent it from transitioning from a blocking state to a forwarding state and creating a loop in the spanning-tree topology.

**Figure 58: Network Topology for Loop Protection**



[Table 195 on page 2000](#) shows the components that will be configured for loop protection.

**Table 195: Topology for Configuring Loop Protection on the QFX Series**

| Components | Settings                                     |
|------------|----------------------------------------------|
| Switch 1   | Switch 1 is the root bridge.                 |
| Switch 2   | Switch 2 has the root port <b>xe-0/0/7</b> . |

Table 195: Topology for Configuring Loop Protection on the QFX Series (*continued*)

| Components                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Settings                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Switch 3 is connected to Switch 1 through interface <b>xe-0/0/6</b> .                                                                                                                                                                      |
| <p>A spanning-tree topology contains ports that have specific roles:</p> <ul style="list-style-type: none"> <li>• The <i>root port</i> is responsible for forwarding data to the root bridge.</li> <li>• The <i>alternate port</i> is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.</li> <li>• The <i>designated port</i> forwards data to the downstream network segment or device.</li> </ul> <p>This configuration example uses an RSTP topology. However, you can also configure loop protection for STP or MSTP topologies at the <code>[edit protocols (mstp   stp)]</code> hierarchy level.</p> |                                                                                                                                                                                                                                            |
| <b>Configuration</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                            |
| <b>CLI Quick Configuration</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>To quickly configure loop protection on interface <b>xe-0/0/6</b>:</p> <pre>[edit] set protocols rstp interface xe-0/0/6 bpdu-timeout-action block</pre>                                                                                |
| <b>Step-by-Step Procedure</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>To configure loop protection:</p> <ol style="list-style-type: none"> <li>1. Configure interface <b>xe-0/0/6</b> on Switch 3: <pre>[edit protocols rstp] user@switch# set interface xe-0/0/6 bpdu-timeout-action block</pre> </li> </ol> |
| <b>Results</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>Check the results of the configuration:</p> <pre>user@switch&gt; show configuration protocols rstp interface xe-0/0/6.0 {   bpdu-timeout-action {     block;   } }</pre>                                                                |
| <b>Verification</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                            |
| <p>To confirm that the configuration is working properly, perform these tasks:</p> <ul style="list-style-type: none"> <li>• <a href="#">Displaying the Interface State Before Loop Protection Is Triggered on page 2001</a></li> <li>• <a href="#">Verifying That Loop Protection Is Working on an Interface on page 2002</a></li> </ul> <p><b>Displaying the Interface State Before Loop Protection Is Triggered</b></p>                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                            |
| <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Before loop protection is triggered on interface <b>xe-0/0/6</b> , confirm that the interface is blocked.                                                                                                                                  |

**Action** Display the interface state and role before applying root protection:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/6.0 | 128:519 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/6.0** is the alternate port and is blocked.

#### *Verifying That Loop Protection Is Working on an Interface*

**Purpose** Verify that the loop protection configuration on interface **xe-0/0/6**. RSTP has been disabled on interface **xe-0/0/4** on Switch 1. This stops BPDUs from being sent to interface **xe-0/0/6** and triggering loop protection on that interface.

**Action** Display the interface state and role after applying root protection:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| xe-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| xe-0/0/6.0 | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Loop-Incon)

[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/6.0** has detected that BPDUs are no longer being forwarded to it and has moved into a loop-inconsistent state. The loop-inconsistent state prevents the interface from transitioning to a forwarding state. The interface recovers and transitions back to its original state as soon as it receives BPDUs.

**Related Documentation**

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1994](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904](#)

## Example: Disabling MAC Learning

By default, MAC learning is enabled on the QFX Series. This topic provides examples for disabling, enabling, and verifying the operation of MAC learning on the QFX Series. These examples require that you be logged in as the root user to the switch on which you wish to modify MAC learning.

- To disable MAC learning in a VLAN:

```
[edit]
user@switch# edit ethernet-switching-options interfaces xe-0/0/0.0
[edit ethernet-switching-options interfaces xe-0/0/0.0]
user@switch# set no-mac-learning
```

- To reenable MAC learning:

```
[edit]
user@switch# edit ethernet-switching-options interfaces xe-0/0/0.0
[edit ethernet-switching-options interfaces xe-0/0/0.0]
user@switch# delete no-mac-learning
```

- To verify the status of MAC learning on the QFX Series:

```
user@switch> show ethernet-switching table
Learning stats: 10 learn msg rcvd, 2 error, 0 forced update
Interface Local pkts Transit pkts Error
xe-0/0/0.0 0 6 1
xe-0/0/22.0 0 0 0
xe-0/0/1.0 0 4 1
xe-0/0/2.0 0 0 0
xe-0/0/3.0 0 0 0
xe-0/0/4.0 0 0 0
xe-0/0/19.0 0 0 0
xe-0/0/18.0 0 0 0
xe-0/0/9.0 0 0 0
```

### Related Documentation

- [Understanding MAC Learning on page 1876](#)
- [Disabling MAC Learning on page 2055](#)
- [no-mac-learning \(Per VLAN\) on page 2109](#)

## Example: Disabling MAC Learning in a VLAN

When MAC learning is enabled, a MAC address is learned dynamically from a packet's source MAC address. By default, MAC learning is enabled on a VLAN. This topic provides examples for disabling, enabling, and verifying the operation of MAC learning in a VLAN. Disabling dynamic MAC learning in a VLAN on a QFX Series product prevents a node from

learning source and destination MAC addresses. These examples require that you be logged in as the root user to the switch on which you wish to modify MAC learning. This example uses a VLAN named *blue*.

- To disable MAC learning in a VLAN:

```
[edit vlans vlan-name]
user@switch# set no-mac-learning
```

For example:

```
[edit vlans blue]
user@switch# set no-mac-learning
```

- To verify that you have disabled MAC learning, issue the **show ethernet-switching table** command:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 2 learned
```

| VLAN    | MAC address       | Type   | Age  | Interfaces    |
|---------|-------------------|--------|------|---------------|
| blue    | *                 | Flood  |      | - All-members |
| blue    | 00:1f:12:39:90:80 | Static |      | - Router      |
| default | *                 | Flood  |      | - All-members |
| default | 00:1f:12:39:90:89 | Learn  | 3:15 | ge-0/0/1.0    |
| default | 00:1f:12:39:a3:81 | Learn  | 0    | ge-0/0/1.0    |

The CLI output shows that the VLAN named *blue* is not configured for MAC learning. The **Type** column includes only **static** (MAC address that are manually created) and **flood** (MAC addresses that are unknown and flooded to all members of the VLAN) entries.

- To reenable MAC learning in a VLAN, issue either of the following two commands::

```
[edit vlans vlan-name]
user@switch delete no-mac-learning
user@switch# deactivate no-mac-learning
```

For example:

```
[edit vlans blue]
user@switch delete no-mac-learning
user@switch# deactivate no-mac-learning
```

- To verify that you have enabled MAC learning, issue the **show ethernet-switching table** command:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 6 entries, 3 learned
```

| VLAN    | MAC address       | Type   | Age | Interfaces    |
|---------|-------------------|--------|-----|---------------|
| blue    | *                 | Flood  |     | - All-members |
| blue    | 00:1f:12:39:90:80 | Static |     | - Router      |
| blue    | 00:1f:12:39:a3:80 | Learn  | 0   | ge-0/0/9.0    |
| default | *                 | Flood  |     | - All-members |
| default | 00:1f:12:39:90:89 | Learn  | 0   | ge-0/0/1.0    |
| default | 00:1f:12:39:a3:81 | Learn  | 0   | ge-0/0/1.0    |

The CLI output shows that the VLAN named *blue* is configured for MAC learning. The **Type** column includes **static** (MAC address that are manually created), **flood** (MAC addresses that are unknown and flooded to all members of the VLAN), and **.Learn** (MAC addresses that are earned dynamically from a packet's source MAC address) entries.



- Related Documentation**
- [Understanding MAC Learning on page 1876](#)
  - [Disabling MAC Learning in a VLAN on page 2056](#)
  - [no-mac-learning \(Per VLAN\) on page 2109](#)
  - [show ethernet-switching table on page 2211](#)

## Private VLAN Configuration Examples

---

- [Example: Configuring a Private VLAN on a Single Switch on page 2005](#)
- [Example: Configuring a Private VLAN Spanning Multiple Switches on page 2010](#)
- [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024](#)

### Example: Configuring a Private VLAN on a Single Switch

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and even to limit the communication between known hosts. The private VLAN (PVLAN) feature allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

This example describes how to create a PVLAN on a single switch:

- [Requirements on page 2005](#)
- [Overview and Topology on page 2005](#)
- [Configuration on page 2006](#)
- [Verification on page 2009](#)

#### Requirements

---

This example uses the following hardware and software components:

- One QFX3500 device
- Junos OS Release 12.1 or later for the QFX Series

Before you begin configuring a PVLAN, make sure you have created and configured the necessary VLANs. See [“Configuring VLANs” on page 2048](#).

#### Overview and Topology

---

In a large office with multiple buildings and VLANs, you might need to isolate some workgroups or other endpoints for security reasons or to partition the broadcast domain. This configuration example shows a simple topology to illustrate how to create a PVLAN with one primary VLAN and two community VLANs, one for HR and one for finance, as well as two isolated ports—one for the mail server and the other for the backup server.

[Table 196 on page 2006](#) lists the settings for the sample topology.

Table 196: Components of the Topology for Configuring a PVLAN

| Interface   | Description                                       |
|-------------|---------------------------------------------------|
| ge-0/0/0.0  | Primary VLAN ( <b>pvlan100</b> ) trunk interface  |
| ge-0/0/11.0 | User 1, HR Community ( <b>hr-comm</b> )           |
| ge-0/0/12.0 | User 2, HR Community ( <b>hr-comm</b> )           |
| ge-0/0/13.0 | User 3, Finance Community ( <b>finance-comm</b> ) |
| ge-0/0/14.0 | User 4, Finance Community ( <b>finance-comm</b> ) |
| ge-0/0/15.0 | Mail server, Isolated ( <b>isolated</b> )         |
| ge-0/0/16.0 | Backup server, Isolated ( <b>isolated</b> )       |
| ge-1/0/0.0  | Primary VLAN ( <b>pvlan100</b> ) trunk interface  |

### Configuration

#### CLI Quick Configuration

To quickly create and configure a PVLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans pvlan100 vlan-id 100
set interfaces ge-0/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-1/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-1/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/15 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/16 unit 0 family ethernet-switching port-mode access
set vlans pvlan100 pvlan
set vlans pvlan100 interface ge-0/0/0.0
set vlans pvlan100 interface ge-1/0/0.0
set vlans hr-comm interface ge-0/0/11.0
set vlans hr-comm interface ge-0/0/12.0
set vlans finance-comm interface ge-0/0/13.0
set vlans finance-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan100
set vlans finance-comm primary-vlan pvlan100
set pvlan100 interface ge-0/0/15.0 isolated
set pvlan100 interface ge-0/0/16.0 isolated
```

#### Step-by-Step Procedure

To configure the PVLAN:

- Set the VLAN ID for the primary VLAN:  

```
[edit vlans]
user@switch# set pvlan vlan-id 100
```
- Set the interfaces and port modes:

**[edit interfaces]**

```

user@switch# set ge-0/0/0 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members pvlan
user@switch# set ge-1/0/0 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-1/0/0 unit 0 family ethernet-switching vlan members pvlan
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/15 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/16 unit 0 family ethernet-switching port-mode access

```

3. Set the primary VLAN to have no local switching:



**NOTE:** The primary VLAN must be a tagged VLAN.

**[edit vlans]**

```
user@switch# set pvlan100 pvlan
```

4. Add the trunk interfaces to the primary VLAN:

**[edit vlans]**

```

user@switch# set pvlan100 interface ge-0/0/0.0
user@switch# set pvlan100 interface ge-1/0/0.0

```

5. For each secondary VLAN, configure access interfaces:



**NOTE:** We recommend that the secondary VLANs be untagged VLANs. It does not impair functioning if you tag the secondary VLANs. However, the tags are not used when a secondary VLAN is configured on a single switch.

**[edit vlans]**

```

user@switch# set hr-comm interface ge-0/0/11.0
user@switch# set hr-comm interface ge-0/0/12.0
user@switch# set finance-comm interface ge-0/0/13.0
user@switch# set finance-comm interface ge-0/0/14.0

```

6. For each community VLAN, set the primary VLAN:

**[edit vlans]**

```

user@switch# set hr-comm primary-vlan pvlan100
user@switch# set finance-comm primary-vlan pvlan100

```

7. Configure the isolated interfaces in the primary VLAN:

**[edit vlans]**

```

user@switch# set pvlan100 interface ge-0/0/15.0 isolated
user@switch# set pvlan100 interface ge-0/0/16.0 isolated

```

**Results**

Check the results of the configuration:

**[edit]**

```

user@switch# show
interfaces {
 ge-0/0/0 {

```

```
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members pvlan100;
 }
 }
 }
 }
 ge-1/0/0 {
 unit 0 {
 family ethernet-switching;
 }
 }
 ge-0/0/11 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
 }
 ge-0/0/12 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
 }
 ge-0/0/13 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
 }
 ge-0/0/14 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
 }
 vlans {
 finance-comm {
 interface {
 ge-0/0/13.0;
 ge-0/0/14.0;
 }
 primary-vlan pvlan100;
 }
 hr-comm {
 interface {
 ge-0/0/11.0;
 ge-0/0/12.0;
 }
 primary-vlan pvlan100;
 }
 }
```

```

}
pvlan100 {
 vlan-id 100;
 interface {
 ge-0/0/15.0;
 ge-0/0/16.0;
 ge-0/0/0.0;
 ge-1/0/0.0;
 }
 pvlan;
}
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Private VLAN and Secondary VLANs Were Created on page 2009](#)

#### *Verifying That the Private VLAN and Secondary VLANs Were Created*

**Purpose** Verify that the primary VLAN and secondary VLANs were properly created on the switch.

**Action** Use the `show vlans` command:

```

user@switch> show vlans pvlan100 extensive
VLAN: pvlan100, Created at: Tue Sep 16 17:59:47 2008
802.1Q Tag: 100, Internal index: 18, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 6 (Active = 0)
 ge-0/0/0.0, tagged, trunk
 ge-0/0/11.0, untagged, access
 ge-0/0/12.0, untagged, access
 ge-0/0/13.0, untagged, access
 ge-0/0/14.0, untagged, access
 ge-0/0/15.0, untagged, access
 ge-0/0/16.0, untagged, access
 ge-1/0/0.0, tagged, trunk
Secondary VLANs: Isolated 2, Community 2
 Isolated VLANs :
 __pvlan_pvlan_ge-0/0/15.0__
 __pvlan_pvlan_ge-0/0/16.0__
 Community VLANs :
 finance-comm
 hr-comm

user@switch> show vlans hr-comm extensive
VLAN: hr-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 22, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 2 (Active = 0)
 ge-0/0/0.0, tagged, trunk
 ge-0/0/11.0, untagged, access
 ge-0/0/12.0, untagged, access
 ge-1/0/0.0, tagged, trunk

user@switch> show vlans finance-comm extensive

```

```
VLAN: finance-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 21, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 2 (Active = 0)
 ge-0/0/0.0, tagged, trunk
 ge-0/0/13.0, untagged, access
 ge-0/0/14.0, untagged, access
 ge-1/0/0.0, tagged, trunk

user@switch> show vlans __pvlan_pvlan_ge-0/0/15.0__ extensive
VLAN: __pvlan_pvlan_ge-0/0/15.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 19, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 1 (Active = 0)
 ge-0/0/0.0, tagged, trunk
 ge-0/0/15.0, untagged, access
 ge-1/0/0.0, tagged, trunk

user@switch> show vlans __pvlan_pvlan_ge-0/0/16.0__ extensive
VLAN: __pvlan_pvlan_ge-0/0/16.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 20, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 1 (Active = 0)
 ge-0/0/0.0, tagged, trunk
 ge-0/0/16.0, untagged, access
 ge-1/0/0.0, tagged, trunk
```

**Meaning** The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.

- Related Documentation**
- [Understanding Private VLANs on page 1878](#)
  - [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
  - [Creating a Private VLAN on a Single Switch on page 2057](#)

## Example: Configuring a Private VLAN Spanning Multiple Switches

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and even to limit the communication between known hosts. The private VLAN (PVLAN) feature allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN. A PVLAN can span multiple switches.

This example describes how to create a PVLAN spanning multiple switches. The example creates one primary PVLAN, containing multiple secondary VLANs:

- [Requirements on page 2011](#)
- [Overview and Topology on page 2011](#)
- [Configuring a PVLAN on Switch 1 on page 2014](#)
- [Configuring a PVLAN on Switch 2 on page 2016](#)

- [Configuring a PVLAN on Switch 3 on page 2019](#)
- [Verification on page 2020](#)

---

## Requirements

This example uses the following hardware and software components:

- Three QFX3500 devices
- Junos OS Release 12.1 or later for the QFX Series

Before you begin configuring a PVLAN, make sure you have created and configured the necessary VLANs. See [“Configuring VLANs” on page 2048](#).

---

## Overview and Topology

In a large office with multiple buildings and VLANs, you might need to isolate some workgroups or other endpoints for security reasons or to partition the broadcast domain. This configuration example shows how to create a PVLAN spanning multiple QFX devices, with one primary VLAN containing two community VLANs (one for HR and one for Finance), and an interswitch isolated VLAN (for the mail server, the backup server, and the CVS server). The PVLAN comprises three switches, two access switches and one distribution switch. The PVLAN is connected to a router through a promiscuous port, which is configured on the distribution switch.



.....

**NOTE:** The isolated ports on Switch 1 and on Switch 2 do not have Layer 2 connectivity with one another even though they are included within the same domain. See [“Understanding Private VLANs” on page 1878](#).

.....

[Figure 59 on page 2012](#) shows the topology for this example—two access switches connecting to a distribution switch, which has a connection (through a promiscuous port) to the router.

Figure 59: PVLAN Topology Spanning Multiple Switches




Table 197 on page 2013, Table 198 on page 2013, and Table 199 on page 2014 list the settings for the example topology.



Table 197: Components of Switch 1 in the Topology for Configuring a PVLAN Spanning Multiple Devices

| Property                              | Settings                                                                                            |
|---------------------------------------|-----------------------------------------------------------------------------------------------------|
| VLAN names and tag IDs                | primary-vlan, tag 100<br><br>isolation-vlan-id, tag 50<br>finance-comm, tag 300<br>hr-comm, tag 400 |
| PVLAN trunk interfaces                | ge-0/0/0.0, connects Switch 1 to Switch 3<br><br>ge-0/0/5.0, connects Switch 1 to Switch 2          |
| Isolated Interfaces in primary VLAN   | ge-0/0/15.0, mail server<br><br>ge-0/0/16.0, backup server                                          |
| Interfaces in VLAN <b>finance-com</b> | ge-0/0/11.0<br><br>ge-0/0/12.0                                                                      |
| Interfaces in VLAN <b>hr-comm</b>     | ge-0/0/13.0<br><br>ge-0/0/14.0                                                                      |

Table 198: Components of Switch 2 in the Topology for Configuring a PVLAN Spanning Multiple Devices

| Property                              | Settings                                                                                            |
|---------------------------------------|-----------------------------------------------------------------------------------------------------|
| VLAN names and tag IDs                | primary-vlan, tag 100<br><br>isolation-vlan-id, tag 50<br>finance-comm, tag 300<br>hr-comm, tag 400 |
| PVLAN trunk interfaces                | ge-0/0/0.0, connects Switch 2 to Switch 3<br><br>ge-0/0/5.0, connects Switch 2 to Switch 1          |
| Isolated Interface in primary VLAN    | ge-0/0/17.0, CVS server                                                                             |
| Interfaces in VLAN <b>finance-com</b> | ge-0/0/11.0<br><br>ge-0/0/12.0                                                                      |
| Interfaces in VLAN <b>hr-comm</b>     | ge-0/0/13.0<br><br>ge-0/0/14.0                                                                      |

**Table 199: Components of Switch 3 in the Topology for Configuring a PVLAN Spanning Multiple Devices**

| Property               | Settings                                                                                                                                                                                                                                                       |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN names and tag IDs | <b>primary-vlan</b> , tag 100<br><br><b>isolation-vlan-id</b> , tag 50<br><b>finance-comm</b> , tag 300<br><b>hr-comm</b> , tag 400                                                                                                                            |
| PVLAN trunk interfaces | <b>ge-0/0/0.0</b> , connects Switch 3 to Switch 1<br><br><b>ge-0/0/1.0</b> , connects Switch 3 to Switch 2                                                                                                                                                     |
| Promiscuous port       | <b>ge-0/0/2</b> , connects the PVLAN to the router<br><br><b>NOTE:</b> You must configure the trunk port that connects the PVLAN to another switch or router outside the PVLAN as a member of the PVLAN, which implicitly configures it as a promiscuous port. |

### Configuring a PVLAN on Switch 1

When configuring a PVLAN on multiple switches, these rules apply:

- The primary VLAN must be a tagged VLAN. We recommend that you configure the primary VLAN first.
- If you are going to configure a community VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port. You must also configure the primary VLAN to be private using the **pvlan** statement.
- If you are going to configure an isolation VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port.

#### CLI Quick Configuration

To quickly create and configure a PVLAN spanning multiple switches, copy the following commands and paste them into the terminal window of Switch 1:

```
[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm interface ge-0/0/11.0
set vlans finance-comm interface ge-0/0/12.0
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm interface ge-0/0/13.0
set vlans hr-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/15.0
set vlans pvlan100 interface ge-0/0/16.0
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/5.0 pvlan-trunk
set vlans pvlan100 pvlan
set vlans pvlan100 pvlan isolation-vlan-id 50
set pvlan100 interface ge-0/0/15.0 isolated
set pvlan100 interface ge-0/0/16.0 isolated
```

**Step-by-Step  
Procedure**

1. Set the VLAN ID for the primary VLAN:  

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
```
2. Set the PVLAN trunk interfaces to connect this VLAN across neighboring switches:  

```
[edit vlans]
user@switch# set pvlan100 interface ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk
```
3. Set the primary VLAN to be private and have no local switching:  

```
[edit vlans]
user@switch# set pvlan100 pvlan
```
4. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:  

```
[edit vlans]
user@switch# set finance-comm vlan-id 300
```
5. Configure access interfaces for the **finance-comm** VLAN:  

```
[edit vlans]
user@switch# set finance-comm interface ge-0/0/11.0
user@switch# set finance-comm interface ge-0/0/12.0
```
6. Set the primary VLAN of this secondary community VLAN, **finance-comm** :  

```
[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100
```
7. Set the VLAN ID for the HR community VLAN that spans the switches.  

```
[edit vlans]
user@switch# set hr-comm vlan-id 400
```
8. Configure access interfaces for the **hr-comm** VLAN:  

```
[edit vlans]
user@switch# set hr-comm interface ge-0/0/13.0
user@switch# set hr-comm interface ge-0/0/14.0
```
9. Set the primary VLAN of this secondary community VLAN, **hr-comm**:  

```
[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100
```
10. Set the interswitch isolated ID to create an interswitch isolated domain that spans the switches:  

```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 50
```
11. Configure the isolated interfaces in the primary VLAN:  

```
[edit vlans]
user@switch# set pvlan100 interface ge-0/0/15.0 isolated
user@switch# set pvlan100 interface ge-0/0/16.0 isolated
```



**NOTE:** When you configure an isolated port, include it as a member of the primary VLAN, but do not configure it as a member of any community VLAN.

**Results**

Check the results of the configuration:

```
[edit]
user@switch# show
vlangs {
 finance-comm {
 vlan-id 300;
 interface {
 ge-0/0/11.0;
 ge-0/0/12.0;
 }
 primary-vlan pvlan100;
 }
 hr-comm {
 vlan-id 400;
 interface {
 ge-0/0/13.0;
 ge-0/0/14.0;
 }
 primary-vlan pvlan100;
 }
 pvlan100 {
 vlan-id 100;
 interface {
 ge-0/0/15.0;
 ge-0/0/16.0;
 ge-0/0/0.0 {
 pvlan-trunk;
 }
 ge-0/0/5.0 {
 pvlan-trunk;
 }
 }
 pvlan;
 isolation-vlan-id 50;
 }
}
```

---

### Configuring a PVLAN on Switch 2

#### CLI Quick Configuration

To quickly create and configure a private VLAN spanning multiple switches, copy the following commands and paste them into the terminal window of Switch 2:



**NOTE:** The configuration of Switch 2 is the same as the configuration of Switch 1 except for the interface in the interswitch isolated domain. For Switch 2, the interface is ge-0/0/17.0.

```
[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm interface ge-0/0/11.0
set vlans finance-comm interface ge-0/0/12.0
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm interface ge-0/0/13.0
```

```

set vlans hr-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/17.0
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/5.0 pvlan-trunk
set vlans pvlan100 pvlan
set vlans pvlan100 pvlan isolation-vlan-id 50
set pvlan100 interface ge-0/0/17.0 isolated

```

### Step-by-Step Procedure

To configure a PVLAN on Switch 2 that will span multiple switches:

1. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:  

```

[edit vlans]
user@switch# set finance-comm vlan-id 300

```
2. Configure access interfaces for the **finance-comm** VLAN:  

```

[edit vlans]
user@switch# set finance-comm interface ge-0/0/11.0
user@switch# set finance-comm interface ge-0/0/12.0

```
3. Set the primary VLAN of this secondary community VLAN, **finance-comm**:  

```

[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100

```
4. Set the VLAN ID for the HR community VLAN that spans the switches.  

```

[edit vlans]
user@switch# set hr-comm vlan-id 400

```
5. Configure access interfaces for the **hr-comm** VLAN:  

```

[edit vlans]
user@switch# set hr-comm interface ge-0/0/13.0
user@switch# set hr-comm interface ge-0/0/14.0

```
6. Set the primary VLAN of this secondary community VLAN, **hr-comm**:  

```

[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100

```
7. Set the VLAN ID for the primary VLAN:  

```

[edit vlans]
user@switch# set pvlan100 vlan-id 100

```
8. Set the PVLAN trunk interfaces that will connect this VLAN across neighboring switches:  

```

[edit vlans]
user@switch# set pvlan100 interface ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk

```
9. Set the primary VLAN to be private and have no local switching:  

```

[edit vlans]
user@switch# set pvlan100 pvlan

```
10. Set the interswitch isolated ID to create an interswitch isolated domain that spans the switches:  

```

[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 50

```



**NOTE:** To configure an isolated port, include it as one of the members of the primary VLAN, but do not configure it as belonging to one of the community VLANs.

11. Configure the isolated interface in the primary VLAN:

```
[edit vlans]
user@switch# set pvlan100 interface ge-0/0/17.0 isolated
```

### Results

Check the results of the configuration:

```
[edit]
user@switch# show
vlans {
 finance-comm {
 vlan-id 300;
 interface {
 ge-0/0/11.0;
 ge-0/0/12.0;
 }
 primary-vlan pvlan100;
 }
 hr-comm {
 vlan-id 400;
 interface {
 ge-0/0/13.0;
 ge-0/0/14.0;
 }
 primary-vlan pvlan100;
 }
 pvlan100 {
 vlan-id 100;
 interface {
 ge-0/0/15.0;
 ge-0/0/16.0;
 ge-0/0/0.0 {
 pvlan-trunk;
 }
 ge-0/0/5.0 {
 pvlan-trunk;
 }
 ge-0/0/17.0;
 }
 pvlan;
 isolation-vlan-id 50;
 }
}
```

### Configuring a PVLAN on Switch 3

**CLI Quick Configuration** To quickly configure Switch 3 to function as the distribution switch of this PVLAN, copy the following commands and paste them into the terminal window of Switch 3:



**NOTE:** Interface ge-0/0/2.0 is a trunk port connecting the PVLAN to a router.

```
[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/1.0 pvlan-trunk
set vlans pvlan100 pvlan
set vlans pvlan100 pvlan isolation-vlan-id 50
```

**Step-by-Step Procedure** To configure Switch 3 to function as the distribution switch for this PVLAN, use the following procedure:

1. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:  

```
[edit vlans]
user@switch# set vlans finance-comm vlan-id 300
```
2. Set the primary VLAN of this secondary community VLAN, **finance-comm**:  

```
[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100
```
3. Set the VLAN ID for the HR community VLAN that spans the switches:  

```
[edit vlans]
user@switch# set vlans hr-comm vlan-id 400
```
4. Set the primary VLAN of this secondary community VLAN, **hr-comm**:  

```
[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100
```
5. Set the VLAN ID for the primary VLAN:  

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
```
6. Set the PVLAN trunk interfaces that will connect this VLAN across neighboring switches:  

```
[edit vlans]
user@switch# set pvlan100 interface ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk
```
7. Set the primary VLAN to be private and have no local switching:  

```
[edit vlans]
user@switch# set pvlan100 pvlan
```
8. Set the interswitch isolated ID to create an interswitch isolated domain that spans the switches:  

```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 50
```



**NOTE:** To configure an isolated port, include it as one of the members of the primary VLAN, but do not configure it as belonging to one of the community VLANs.

### Results

Check the results of the configuration:

```
[edit]
user@switch# show
vlands {
 finance-comm {
 vlan-id 300;
 primary-vlan pvlan100;
 }
 hr-comm {
 vlan-id 400;
 primary-vlan pvlan100;
 }
 pvlan100 {
 vlan-id 100;
 interface {
 ge-0/0/0.0 {
 pvlan-trunk;
 }
 ge-0/0/1.0 {
 pvlan-trunk;
 }
 ge-0/0/2.0;
 }
 pvlan;
 isolation-vlan-id 50;
 }
}
```

---

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 1 on page 2020](#)
- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 2 on page 2022](#)
- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 3 on page 2023](#)

#### ***Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 1***

**Purpose** Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 1:



**Action** Use the **show vlans extensive** command:

```

user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_ge-0/0/15.0__, Created at: Thu Sep 16 23:15:27 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/15.0*, untagged, access

VLAN: __pvlan_pvlan100_ge-0/0/16.0__, Created at: Thu Sep 16 23:15:27 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/16.0*, untagged, access

VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk

VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)

VLAN: finance-comm, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 300, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/11.0*, untagged, access
 ge-0/0/12.0*, untagged, access

VLAN: hr-comm, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 400, Internal index: 9, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/13.0*, untagged, access
 ge-0/0/14.0*, untagged, access

VLAN: pvlan100, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 6 (Active = 6)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk

```

```

ge-0/0/11.0*, untagged, access
ge-0/0/12.0*, untagged, access
ge-0/0/13.0*, untagged, access
ge-0/0/14.0*, untagged, access
ge-0/0/15.0*, untagged, access
ge-0/0/16.0*, untagged, access
Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
__pvlan_pvlan100_ge-0/0/15.0__
__pvlan_pvlan100_ge-0/0/16.0__
Community VLANs :
finance-comm
hr-comm
Inter-switch-isolated VLAN :
__pvlan_pvlan100_isiv__

```

**Meaning** The output shows that a PVLAN was created on Switch 1 and shows that it includes two isolated VLANs, two community VLANs, and an interswitch isolated VLAN. The presence of the pvlan-trunk and Inter-switch-isolated fields indicates that this PVLAN is spanning more than one switch.

#### *Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 2*

**Purpose** Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 2:

**Action** Use the `show vlans extensive` command:

```

user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_ge-0/0/17.0__, Created at: Thu Sep 16 23:19:22 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
ge-0/0/0.0*, tagged, trunk, pvlan-trunk
ge-0/0/5.0*, tagged, trunk, pvlan-trunk
ge-0/0/17.0*, untagged, access

VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 50, Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
ge-0/0/0.0*, tagged, trunk, pvlan-trunk
ge-0/0/5.0*, tagged, trunk, pvlan-trunk

VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)

VLAN: finance-comm, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 300, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
ge-0/0/0.0*, tagged, trunk, pvlan-trunk
ge-0/0/5.0*, tagged, trunk, pvlan-trunk
ge-0/0/11.0*, untagged, access

```

```
ge-0/0/12.0*, untagged, access
```

```
VLAN: hr-comm, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 400, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/13.0*, untagged, access
 ge-0/0/14.0*, untagged, access
```

```
VLAN: pvlan100, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 5 (Active = 5)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/5.0*, tagged, trunk, pvlan-trunk
 ge-0/0/11.0*, untagged, access
 ge-0/0/12.0*, untagged, access
 ge-0/0/13.0*, untagged, access
 ge-0/0/14.0*, untagged, access
 ge-0/0/17.0*, untagged, access
Secondary VLANs: Isolated 1, Community 2, Inter-switch-isolated 1
Isolated VLANs :
 __pvlan_pvlan100_ge-0/0/17.0__
Community VLANs :
 finance-comm
 hr-comm
Inter-switch-isolated VLAN :
 __pvlan_pvlan100_isiv__
```

**Meaning** The output shows that a PVLAN was created on Switch 2 and shows that it includes one isolated VLAN, two community VLANs, and an interswitch isolated VLAN. The presence of the pvlan-trunk and Inter-switch-isolated fields indicates that this PVLAN is spanning more than one switch. When you compare this output to the output of Switch 1, you can see that both switches belong to the same PVLAN (**pvlan100**).

### *Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 3*

**Purpose** Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 3:

**Action** Use the **show vlans extensive** command:

```
user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 50, Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
```

```
VLAN: finance-comm, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 300, Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: hr-comm, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 400, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: pvlan100, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
 ge-0/0/0.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, tagged, trunk, pvlan-trunk
Secondary VLANs: Isolated 0, Community 2, Inter-switch-isolated 1
Community VLANs :
 finance-comm
 hr-comm
Inter-switch-isolated VLAN :
 __pvlan_pvlan100_isiv__
```

**Meaning** The output shows that the PVLAN (**pvlan100**) is configured on Switch 3 and that it includes no isolated VLANs, two community VLANs, and an interswitch isolated VLAN. But Switch 3 is functioning as a distribution switch, so the output does not include access interfaces within the PVLAN. It shows only the **pvlan-trunk** interfaces that connect **pvlan100** from Switch 3 to the other switches (Switch 1 and Switch 2) in the same PVLAN.

**Related Documentation**

- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Example: Configuring a Private VLAN on a Single Switch on page 2005](#)

## Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports

This example shows how to configure secondary VLAN trunk ports and promiscuous access ports as part of a private VLAN configuration. Secondary VLAN trunk ports carry secondary VLAN traffic.

For a given private VLAN, a secondary VLAN trunk port can carry traffic for only one secondary VLAN. However, a secondary VLAN trunk port can carry traffic for multiple secondary VLANs as long as each secondary VLAN is a member of a different private (primary) VLAN. For example, a secondary VLAN trunk port can carry traffic for a community VLAN that is part of primary VLAN **pvlan100** and also carry traffic for an isolated VLAN that is part of primary VLAN **pvlan400**.

To configure a trunk port to carry secondary VLAN traffic, use the **isolated** and **interface** statements, as shown in steps 12 and 13 of the example configuration for Switch 1.



**NOTE:** When traffic egresses from a secondary VLAN trunk port, it normally carries the tag of the primary VLAN that the secondary port is a member of. If you want traffic that egresses from a secondary VLAN trunk port to retain its secondary VLAN tag, use the **extend-secondary-vlan-id** statement.

A promiscuous access port carries untagged traffic and can be a member of only one primary VLAN. Traffic that ingresses on a promiscuous access port is forwarded to the ports of the secondary VLANs that are members of the primary VLAN that the promiscuous access port is a member of. This traffic carries the appropriate secondary VLAN tags when it egresses from the secondary VLAN ports if the secondary VLAN port is a trunk port.

To configure an access port to be promiscuous, use the **promiscuous** statement, as shown in step Figure 50 on page 1895 of the example configuration for Switch 2.

If traffic ingresses on a secondary VLAN port and egresses on a promiscuous access port, the traffic is untagged on egress. If tagged traffic ingresses on a promiscuous access port, the traffic is discarded.

- [Requirements on page 2025](#)
- [Overview and Topology on page 2025](#)
- [Configuring the PVLANS on Switch 1 on page 2027](#)
- [Configuring the PVLANS on Switch 2 on page 2031](#)
- [Verification on page 2035](#)

## Requirements

This example uses the following hardware and software components:

- Two QFX devices
- Junos OS Release 12.2 or later for the QFX Series

## Overview and Topology

Figure 60 on page 2026 shows the topology used in this example. Switch 1 includes several primary and secondary private VLANs and also includes two secondary VLAN trunk ports configured to carry secondary VLANs that are members of primary VLANs pvlan100 and pvlan400.

Switch 2 includes the same private VLANs. The figure shows xe-0/0/0 on Switch 2 as configured with promiscuous access ports or promiscuous trunk ports. The example configuration included here configures this port as a promiscuous access port.

The figure also shows how traffic would flow after ingressing on the secondary VLAN trunk ports on Switch 1.

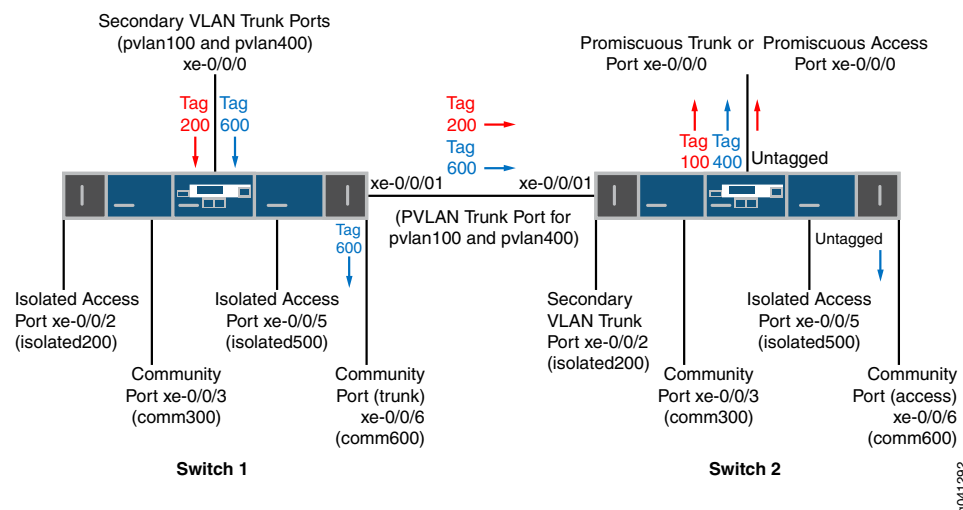
**Figure 60: PVLAN Topology with Secondary VLAN Trunk Ports and Promiscuous Access Port**

Table 200 on page 2026 and Table 201 on page 2027 list the settings for the example topology on both switches.

**Table 200: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 1**

| Component             | Description                                                       |
|-----------------------|-------------------------------------------------------------------|
| pvlan100, ID 100      | Primary VLAN                                                      |
| pvlan400, ID 400      | Primary VLAN                                                      |
| comm300, ID 300       | Community VLAN, member of pvlan100                                |
| comm600, ID 600       | Community VLAN, member of pvlan400                                |
| isolation-vlan-id 200 | VLAN ID for isolated VLAN, member of pvlan100                     |
| isolation-vlan-id 500 | VLAN ID for isolated VLAN, member of pvlan400                     |
| xe-0/0/0.0            | Secondary VLAN trunk port for primary VLANs pvlan100 and pvlan400 |
| xe-0/0/1.0            | PVLAN trunk port for primary VLANs pvlan100 and pvlan400          |
| xe-0/0/2.0            | Isolated access port for pvlan100                                 |
| xe-0/0/3.0            | Community access port for comm300                                 |
| xe-0/0/5.0            | Isolated access port for pvlan400                                 |
| xe-0/0/6.0            | Community trunk port for comm600                                  |

Table 201: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 2

| Component             | Description                                                |
|-----------------------|------------------------------------------------------------|
| pvlan100, ID 100      | Primary VLAN                                               |
| pvlan400, ID 400      | Primary VLAN                                               |
| comm300, ID 300       | Community VLAN, member of pvlan100                         |
| comm600, ID 600       | Community VLAN, member of pvlan400                         |
| isolation-vlan-id 200 | VLAN ID for isolated VLAN, member of pvlan100              |
| isolation-vlan-id 500 | VLAN ID for isolated VLAN, member of pvlan400              |
| xe-0/0/0.0            | Promiscuous access port for primary VLANs pvlan100         |
| xe-0/0/1.0            | PVLAN trunk port for primary VLANs pvlan100 and pvlan400   |
| xe-0/0/2.0            | Secondary trunk port for isolated VLAN, member of pvlan100 |
| xe-0/0/3.0            | Community access port for comm300                          |
| xe-0/0/5.0            | Isolated access port for pvlan400                          |
| xe-0/0/6.0            | Community access port for comm600                          |

### Configuring the PVLANS on Switch 1

**CLI Quick Configuration** To quickly create and configure the PVLANS on Switch 1, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode trunk
set vlans pvlan100 vlan-id 100
set vlans pvlan400 vlan-id 400
set vlans pvlan100 pvlan
set vlans pvlan400 pvlan
set vlans pvlan100 interface xe-0/0/1.0 pvlan-trunk
set vlans pvlan400 interface xe-0/0/1.0 pvlan-trunk
set vlans comm300 vlan-id 300
set vlans comm300 primary-vlan pvlan100
set vlans comm300 interface xe-0/0/3.0
set vlans comm600 vlan-id 600
set vlans comm600 primary-vlan pvlan400
set vlans comm600 interface xe-0/0/6.0
set vlans pvlan100 pvlan isolation-vlan-id 200
```

```

set vlans pvlan400 pvlan isolation-vlan-id 500
set vlans pvlan100 interface xe-0/0/0.0 isolated
set vlans pvlan400 interface xe-0/0/0.0 isolated
set vlans comm600 interface xe-0/0/0.0
set vlans pvlan100 interface xe-0/0/2.0 isolated
set vlans pvlan400 interface xe-0/0/5.0 isolated

```

### Step-by-Step Procedure

To configure the private VLANs and secondary VLAN trunk ports:

1. Configure the interfaces and port modes:

[edit interfaces]

```

user@switch# set xe-0/0/0 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
user@switch# set xe-0/0/2 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/3 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/5 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode access

```

2. Create the primary VLANs:

[edit vlans]

```

user@switch# set pvlan100 vlan-id 100
user@switch# set pvlan400 vlan-id 400

```



**NOTE:** Primary VLANs must always be tagged VLANs, even if they exist on only one device.

3. Configure the primary VLANs to be private:

[edit vlans]

```

user@switch# set pvlan100 pvlan
user@switch# set pvlan400 pvlan

```

4. Configure the PVLAN trunk port to carry the private VLAN traffic between the switches:

[edit vlans]

```

user@switch# set pvlan100 interface xe-0/0/1.0 pvlan-trunk
user@switch# set pvlan400 interface xe-0/0/1.0 pvlan-trunk

```

5. Create secondary VLAN comm300 with VLAN ID 300:

[edit vlans]

```

user@switch# set comm300 vlan-id 300

```

6. Configure the primary VLAN for comm300:

[edit vlans]

```

user@switch# set comm300 primary-vlan pvlan100

```

7. Configure the interface for comm300:

[edit vlans]

```

user@switch# set comm300 interface xe-0/0/3.0

```

8. Create secondary VLAN comm600 with VLAN ID 600:

[edit vlans]

```

user@switch# set comm600 vlan-id 600

```

9. Configure the primary VLAN for comm600:



- ```
[edit vlans]
user@switch# set comm600 primary-vlan pvlan400
```
10. Configure the interface for comm600:
- ```
[edit vlans]
user@switch# set comm600 interface xe-0/0/6.0
```
11. Configure the interswitch isolated VLANs:
- ```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 200
user@switch# set pvlan400 pvlan isolation-vlan-id 500
```



NOTE: When you configure a secondary VLAN trunk port to carry an isolated VLAN, you must also configure an **isolation-vlan-id**. This is true even if the isolated VLAN exists only on one switch.

12. Enable trunk port xe-0/0/0 to carry secondary VLANs for the primary VLANs:
- ```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/0.0 isolated
user@switch# set pvlan400 interface xe-0/0/0.0 isolated
```
13. Configure trunk port xe-0/0/0 to carry comm600 (member of pvlan400):
- ```
[edit vlans]
user@switch# set comm600 interface xe-0/0/0.0
```



NOTE: You do not need to explicitly configure xe-0/0/0 to carry the isolated VLAN traffic (tags 200 and 500) because all the isolated ports in pvlan100 and pvlan400—including xe-0/0/0.0—are automatically included in the isolated VLANs created when you configured **isolation-vlan-id 200** and **isolation-vlan-id 500**.

14. Configure xe-0/0/2 and xe-0/0/6 to be isolated:
- ```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/2.0 isolated
user@switch# set pvlan400 interface xe-0/0/5.0 isolated
```

## Results

Check the results of the configuration on Switch 1:

```
[edit]
user@switch# show
interfaces {
 xe-0/0/0 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members pvlan100;
 members pvlan400;
 }
 }
 }
 }
}
```

```
 }
 }
}
xe-0/0/1 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members pvlan100;
 members pvlan400;
 }
 }
 }
}
xe-0/0/2 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
}
xe-0/0/3 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
}
xe-0/0/5 {
 unit 0 {
 family ethernet-switching {
 port-mode access;
 }
 }
}
xe-0/0/6 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 }
 }
}
}
vlands {
 comm300 {
 vlan-id 300;
 interface {
 xe-0/0/3.0;
 }
 primary-vlan pvlan100;
 }
 comm600 {
 vlan-id 600;
 interface {
 xe-0/0/6.0;
 }
 }
}
```

```

 primary-vlan pvlan400;
 }
 pvlan100 {
 vlan-id 100;
 interface {
 xe-0/0/0.0;
 xe-0/0/2.0;
 xe-0/0/3.0;
 xe-0/0/1.0 {
 pvlan-trunk;
 }
 }
 no-local-switching;
 isolation-id 200;
 }
 pvlan400 {
 vlan-id 400;
 interface {
 xe-0/0/0.0;
 xe-0/0/5.0;
 xe-0/0/6.0;
 xe-0/0/1.0 {
 pvlan-trunk;
 }
 }
 no-local-switching;
 isolation-id 500;
 }
}

```

### Configuring the PVLANS on Switch 2

The configuration for Switch 2 is almost identical to the configuration for Switch 1. The most significant difference is that xe-0/0/0 on Switch 2 is configured as a promiscuous trunk port or a promiscuous access port, as [Figure 60 on page 2026](#) shows. In the following configuration, xe-0/0/0 is configured as a promiscuous access port for primary VLAN pvlan100.

If traffic ingresses on VLAN-enabled port and egresses on a promiscuous access port, the VLAN tags are dropped on egress and the traffic is untagged at that point. For example, traffic for comm600 ingresses on the secondary VLAN trunk port configured on xe-0/0/0.0 on Switch 1 and carries tag 600 as it is forwarded through the secondary VLAN. When it egresses from xe-0/0/0.0 on Switch 2, it will be untagged if you configure xe-0/0/0.0 as a promiscuous access port as shown in this example. If you instead configure xe-0/0/0.0 as a promiscuous trunk port (port-mode trunk), the traffic for comm600 carries its primary VLAN tag (400) when it egresses.

#### CLI Quick Configuration

To quickly create and configure the PVLANS on Switch 2, copy the following commands and paste them into a switch terminal window:

```

[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400

```

```
set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/3 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode access
set vlans pvlan100 vlan-id 100
set vlans pvlan400 vlan-id 400
set vlans pvlan100 pvlan
set vlans pvlan400 pvlan
set vlans pvlan100 interface xe-0/0/1.0 pvlan-trunk
set vlans pvlan400 interface xe-0/0/1.0 pvlan-trunk
set vlans comm300 vlan-id 300
set vlans comm300 primary-vlan pvlan100
set vlans comm300 interface xe-0/0/3.0
set vlans comm600 vlan-id 600
set vlans comm600 primary-vlan pvlan400
set vlans comm600 interface xe-0/0/6.0
set vlans pvlan100 pvlan isolation-vlan-id 200
set vlans pvlan400 pvlan isolation-vlan-id 500
set vlans pvlan100 interface xe-0/0/0.0 promiscuous
set vlans pvlan100 interface xe-0/0/2.0 isolated
set vlans pvlan400 interface xe-0/0/5.0 isolated
```

#### Step-by-Step Procedure

To configure the private VLANs and secondary VLAN trunk ports:

1. Configure the interfaces and port modes:  
  
[edit interfaces]  
user@switch# set xe-0/0/0 unit 0 family ethernet-switching port-mode access  
user@switch# set xe-0/0/1 unit 0 family ethernet-switching port-mode trunk  
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100  
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400  
user@switch# set xe-0/0/2 unit 0 family ethernet-switching port-mode trunk  
user@switch# set xe-0/0/3 unit 0 family ethernet-switching port-mode access  
user@switch# set xe-0/0/5 unit 0 family ethernet-switching port-mode access  
user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode access
2. Create the primary VLANs:  
  
[edit vlans]  
user@switch# set pvlan100 vlan-id 100  
user@switch# set pvlan400 vlan-id 400
3. Configure the primary VLANs to be private:  
  
[edit vlans]  
user@switch# set pvlan100 pvlan  
user@switch# set pvlan400 pvlan
4. Configure the PVLAN trunk port to carry the private VLAN traffic between the switches:  
  
[edit vlans]  
user@switch# set pvlan100 interface xe-0/0/1.0 pvlan-trunk  
user@switch# set pvlan400 interface xe-0/0/1.0 pvlan-trunk
5. Create secondary VLAN comm300 with VLAN ID 300:  
  
[edit vlans]  
user@switch# set comm300 vlan-id 300
6. Configure the primary VLAN for comm300:  
  
[edit vlans]  
user@switch# set comm300 primary-vlan pvlan100
7. Configure the interface for comm300:

- ```
[edit vlans]
user@switch# set comm300 interface xe-0/0/3.0
```
8. Create secondary VLAN comm600 with VLAN ID 600:


```
[edit vlans]
user@switch# set comm600 vlan-id 600
```
 9. Configure the primary VLAN for comm600:


```
[edit vlans]
user@switch# set comm600 primary-vlan pvlan400
```
 10. Configure the interface for comm600:


```
[edit vlans]
user@switch# set comm600 interface xe-0/0/6.0
```
 11. Configure the interswitch isolated VLANs:


```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 200
user@switch# set pvlan400 pvlan isolation-vlan-id 500
```
 12. Configure access port xe-0/0/0 to be promiscuous for pvlan100:


```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/0.0 promiscuous
```



NOTE: A promiscuous access port can be a member of only one primary VLAN.

13. Configure xe-0/0/2 and xe-0/0/6 to be isolated:


```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/2.0 isolated
user@switch# set pvlan400 interface xe-0/0/5.0 isolated
```

Results

Check the results of the configuration on Switch 2:

```
[edit]
user@switch# show
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
          members pvlan100;
        }
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members pvlan100;
        }
      }
    }
  }
}
```

```
        members pvlan400;
    }
}
}
xe-0/0/2 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
        }
    }
}
xe-0/0/3 {
    unit 0 {
        family ethernet-switching {
            port-mode access;
        }
    }
}
xe-0/0/5 {
    unit 0 {
        family ethernet-switching {
            port-mode access;
        }
    }
}
xe-0/0/6 {
    unit 0 {
        family ethernet-switching {
            port-mode access;
        }
    }
}
vlands {
    comm300 {
        vlan-id 300;
        interface {
            xe-0/0/3.0;
        }
        primary-vlan pvlan100;
    }
    comm600 {
        vlan-id 600;
        interface {
            xe-0/0/6.0;
        }
        primary-vlan pvlan400;
    }
    pvlan100 {
        vlan-id 100;
        interface {
            xe-0/0/0.0;
            xe-0/0/2.0;
            xe-0/0/3.0;
            xe-0/0/1.0 {
                pvlan-trunk;
            }
        }
    }
}
```

```

    }
  }
  no-local-switching;
  isolation-id 200;
}
pvlan400 {
  vlan-id 400;
  interface {
    xe-0/0/5.0;
    xe-0/0/6.0;
    xe-0/0/1.0 {
      pvlan-trunk;
    }
  }
  no-local-switching;
  isolation-id 500;
}
}

```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Private VLAN and Secondary VLANs Were Created on page 2035](#)
- [Verifying The Ethernet Switching Table Entries on page 2035](#)

Verifying That the Private VLAN and Secondary VLANs Were Created

Purpose Verify that the primary VLAN and secondary VLANs were properly created on Switch 1.

Action Use the `show vlans` command:

```
user@switch> show vlans private-vlan
```

Name	Role	Tag	Interfaces
pvlan100	Primary	100	xe-0/0/0.0, xe-0/0/1.0, xe-0/0/2.0,
xe-0/0/3.0			
__iso_pvlan100__	Isolated	200	xe-0/0/2.0
comm300	Community	300	xe-0/0/3.0
pvlan400	Primary	400	xe-0/0/0.0, xe-0/0/1.0, xe-0/0/5.0,
xe-0/0/6.0			
__iso_pvlan400__	Isolated	500	xe-0/0/5.0
comm600	Community	600	xe-0/0/6.0

Meaning The output shows that the private VLANs were created and identifies the interfaces and secondary VLANs associated with them.

Verifying The Ethernet Switching Table Entries

Purpose Verify that the Ethernet switching table entries were created for primary VLAN pvlan100.

Action Show the Ethernet switching table entries for pvlan100.

```
user@switch> show ethernet-switching table vlan pvlan100 private-vlan
```

```
Ethernet-switching table: 0 unicast entries
pvlan100          *          Flood          - All-members
pvlan100          00:10:94:00:00:02 Learn      xe-0/0/2.0
__iso_pvlan100__  *          Flood          - All-members
__iso_pvlan100__  00:10:94:00:00:02 Replicated - xe-0/0/2.0
```

- Related Documentation**
- [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887](#)
 - [Understanding Private VLANs on page 1878](#)
 - [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
 - [Creating a Private VLAN on a Single Switch on page 2057](#)
 - [Understanding Egress Firewall Filters with PVLANS on page 1896](#)
 - [Troubleshooting Private VLANs on page 2263](#)

Q-in-Q Tunneling Configuration Example

- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)

Example: Setting Up Q-in-Q Tunneling

Service providers can use Q-in-Q tunneling to transparently pass Layer 2 VLAN traffic between customer sites without removing or changing the customer VLAN tags or class-of-service (CoS) settings. Data centers can use Q-in-Q tunneling to isolate customer traffic within a single site or when customer traffic flows between cloud data centers in different geographic locations.

This example describes how to set up Q-in-Q tunneling:

- [Requirements on page 2036](#)
- [Overview and Topology on page 2036](#)
- [Configuration on page 2037](#)
- [Verification on page 2038](#)

Requirements

This example requires one QFX Series device with Junos OS Release 12.1 or later.

Before you begin setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs on the neighboring switches. See [“Configuring VLANs” on page 2048](#).

Overview and Topology

In this service provider network, there are multiple customer VLANs mapped to one service VLAN.

[Table 202 on page 2037](#) lists the settings for the sample topology.

Table 202: Components of the Topology for Setting Up Q-in-Q Tunneling

Interface	Description
xe-0/0/11.0	Tagged S-VLAN trunk port
xe-0/0/12.0	Untagged customer-facing access port
xe-0/0/13.0	Untagged customer-facing access port
xe-0/0/14.0	Tagged S-VLAN trunk port

Configuration

CLI Quick Configuration

To quickly create and configure Q-in-Q tunneling, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans service-vlan vlan-id 1000
set vlans service-vlan dot1q-tunneling customer-vlans 1-100
set vlans service-vlan dot1q-tunneling customer-vlans 201-300
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members 1000
set interfaces xe-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/12 unit 0 family ethernet-switching vlan members 1000
set interfaces xe-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members 1000
set interfaces xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members 1000
set ethernet-switching-options dot1q-tunneling ether-type 0x9100
```

Step-by-Step Procedure

To configure Q-in-Q tunneling:

1. Set the VLAN ID for the S-VLAN:

```
[edit vlans]
user@switch# set service-vlan vlan-id 1000
```

2. Enable Q-in-Q tunneling and specify the customer VLAN ranges:

```
[edit vlans]
user@switch# set service-vlan dot1q-tunneling customer-vlans 1-100
user@switch# set service-vlan dot1q-tunneling customer-vlans 201-300
```

3. Set the port mode and VLAN information for the interfaces:

```
[edit interfaces]
user@switch# set xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/11 unit 0 family ethernet-switching vlan members 1000
user@switch# set xe-0/0/12 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/12 unit 0 family ethernet-switching vlan members 1000
user@switch# set xe-0/0/13 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/13 unit 0 family ethernet-switching vlan members 1000
user@switch# set xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/14 unit 0 family ethernet-switching vlan members 1000
```

4. Set the Q-in-Q Ethertype value (optional):

```
[edit]
user@switch# set ethernet-switching-options dot1q-tunneling ether-type 0x9100
```

Results

Check the results of the configuration:

```
user@switch> show configuration vlans service-vlan
vlan-id 1000 {
  dot1q-tunneling {
    customer-vlans [ 1-100 201-300 ];
  }
}
user@switch> show configuration interfaces
xe-0/0/11 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan members 1000;
    }
  }
}
xe-0/0/12 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
      vlan members 1000;
    }
  }
}
xe-0/0/13 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
      vlan members 1000;
    }
  }
}
xe-0/0/14 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan members 1000;
    }
  }
}
user@switch> show ethernet-switching-options
dot1q-tunneling {
  ether-type 0x9100;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That Q-in-Q Tunneling Was Enabled

Purpose Verify that Q-in-Q tunneling was properly enabled.

Action Use the **show vlans** command:

```
user@switch> show vlans service-vlan extensive
VLAN: service-vlan, Created at: Wed Mar 14 07:17:53 2012
802.1Q Tag: 1000, Internal index: 18, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    1-100
    201-300
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 2 (Active = 0)
    xe-0/0/11.0, tagged, trunk
    xe-0/0/14.0, tagged, trunk
    xe-0/0/12.0, untagged, access
    xe-0/0/13.0, untagged, access
```

Meaning The output indicates that Q-in-Q tunneling is enabled and that the VLAN is tagged and shows the associated customer VLANs.

Related Documentation

- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
- [Configuring Q-in-Q Tunneling on page 2062](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 2266](#)

Configuration Tasks

- [Configuring the Interface Address on page 2040](#)
- [Configuring MAC Notification on page 2042](#)
- [Configuring MAC Table Aging on page 2044](#)
- [Configuring Reflective Relay on page 2044](#)
- [Configuring Static ARP Entries on page 2045](#)
- [Configuring Reflective Relay on page 2045](#)
- [Configuring Routed VLAN Interfaces on page 2046](#)
- [Configuring the Native VLAN Identifier on page 2047](#)
- [Configuring VLANs on page 2048](#)
- [Creating a Series of Tagged VLANs on page 2050](#)
- [Configuring Multiple VLAN Registration Protocol on page 2051](#)
- [Configuring STP on page 2053](#)
- [Unblocking an Interface That Receives BPDUs in Error on page 2054](#)
- [Configuring VLAN Spanning Tree Protocol on page 2054](#)
- [Disabling MAC Learning on page 2055](#)
- [Disabling MAC Learning in a VLAN on page 2056](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Configuring the Forwarding Mode on page 2059](#)

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.



NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, include the **address** statement:

```
address address {  
    broadcast address;  
    destination address;  
    destination-profile name;  
    eui-64;  
    preferred;  
    primary;  
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the [edit access group-profile *name* **ppp**] hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**)

because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.

- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.

- [Configuring Interface IPv4 Addresses on page 2041](#)
- [Configuring Interface IPv6 Addresses on page 2041](#)

Configuring Interface IPv4 Addresses

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

To configure an IPv4 address on routers and switches running Junos OS, use the **edit interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.



NOTE: Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.

Configuring Interface IPv6 Addresses



NOTE: IPv6 is not currently supported for the QFX Series.

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```



NOTE: You cannot configure a subnet zero IPv6 address because RFC 2461 reserves the subnet-zero address for anycast addresses, and Junos OS complies with the RFC.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {  
  unit 0 {  
    family inet6 {  
      address fec0:1:1::2/64;  
    }  
  }  
}
```



NOTE: You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

**Related
Documentation**

- *Configuring IPCP Options*
- *Configuring Default, Primary, and Preferred Addresses and Interfaces*

Configuring MAC Notification

When a MAC address is learned or unlearned, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as *MAC notification*.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- [Enabling MAC Notification on page 2043](#)
- [Disabling MAC Notification on page 2043](#)
- [Setting the MAC Notification Interval on page 2043](#)

Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit ethernet-switching-options]  
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, 60 seconds):

```
[edit ethernet-switching-options]  
user@switch# set mac-notification notification-interval 60
```

Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit ethernet-switching-options]  
user@switch# delete mac-notification
```

Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit ethernet-switching-options]  
user@switch# set mac-notification notification-interval 5
```

Related Documentation

- [Verifying That MAC Notification Is Working Properly on page 2173](#)

Configuring MAC Table Aging

MAC table aging ensures that a switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

To manage MAC entries more efficiently, you can configure an entry's aging time, which is the maximum time that an entry can remain in the MAC address table before it is deleted because it has reached its maximum age.

You can use the **set-mac-table-aging-time** command to configure how long entries remain in the Ethernet switching table before expiring. Here the VLAN is **employee-vlan**:

```
[edit vlans employee-vlan]
user@switch# set mac-table-aging-time 200
```

Related Documentation

- [Understanding Bridging on page 1869](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 1982](#)

Configuring Reflective Relay

Configure reflective relay when a switch port must return packets on a downstream port. For example, configure reflective relay when a switch port receives aggregated virtual machine packets from a technology such as virtual Ethernet port aggregator (VEPA). When these packets are passed through the switch, reflective relay allows the switch to send those packets back on the same interface that was used for delivery.

Before you begin configuring reflective relay, ensure that you have:

- Configured packet aggregation on the server connected to the port. See your server documentation.
- Configured the port for all VLANs that could be included in aggregated packets..

To configure reflective relay:

1. Configure an Ethernet interface with a port mode of **tagged-access**:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type port-mode
tagged-access
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure the interface for reflective relay:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type reflective-relay
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```


3. Configure the interface for the VLANs that exist on the VM server:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type vlan members
vlan-names
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members
[VLAN_Purple VLAN_Orange VLAN_Blue]
```

Related Documentation

- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958](#)
- [Understanding Reflective Relay for Use with VEPA Technology on page 1877](#)

Configuring Static ARP Entries

You can create static ARP table entries, which are explicit mappings between IP addresses and MAC addresses.

- To configure a static ARP entry:

```
[edit interfaces interface-name unit logical-unit-number family inet address
address]
user@switch# set arp ip-address (mac | multicast-mac) mac-address
```

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, use the **multicast-mac** statement.

Specify the MAC address as 6 hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

Related Documentation

- [Understanding Static ARP Entries on page 4693](#)
- [arp on page 2072](#)

Configuring Reflective Relay

Configure reflective relay when a switch port must return packets on a downstream port. For example, configure reflective relay when a switch port receives aggregated virtual machine packets from a technology such as virtual Ethernet port aggregator (VEPA). When these packets are passed through the switch, reflective relay allows the switch to send those packets back on the same interface that was used for delivery.

Before you begin configuring reflective relay, ensure that you have:

- Configured packet aggregation on the server connected to the port. See your server documentation.
- Configured the port for all VLANs that could be included in aggregated packets..

To configure reflective relay:

1. Configure an Ethernet interface with a port mode of **tagged-access**:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type port-mode
tagged-access
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure the interface for reflective relay:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type reflective-relay
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the VLANs that exist on the VM server:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type vlan members
vlan-names
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members
[VLAN_Purple VLAN_Orange VLAN_Blue]
```

Related Documentation

- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958](#)
- [Understanding Reflective Relay for Use with VEPA Technology on page 1877](#)

Configuring Routed VLAN Interfaces

Routed VLAN interfaces (RVIs) enable a switch to recognize which packets are being sent to local addresses so that they are bridged whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address look-ups.

To configure the routed VLAN interface:

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

```
[edit]
user@switch# set vlans support vlan-id 111
```

2. Assign an interface to the VLAN by specifying the logical interface (with the **unit** statement) and specifying the VLAN name as the member:

```
[edit]
user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members
support
```

3. Create the subnet for the VLAN's broadcast domain:

```
[edit]
```

```
user@switch# set interfaces vlan unit 111 family inet address 111.111.111.1/24
```

4. Bind a Layer 3 interface with the VLAN:

```
[edit]
```

```
user@switch# set vlans support l3-interface vlan.111
```



NOTE: Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces vlan terse
```

Interface	Admin	Link	Proto	Local	Remote
vlan	up	up			
vlan.111	up	up	inet	111.111.111.1/24	

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		None
employee-vlan	20	ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
marketing	40	ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
support	111	ge-0/0/18.0
mgmt		bme0.32769, bme0.32771*

```
user@switch> show ethernet-switching table
```

Ethernet-switching table: 1 entries, 0 learned

VLAN	MAC address	Type	Age	Interfaces
support	00:19:e2:50:95:a0	Static		- Router

Related Documentation

Configuring the Native VLAN Identifier

Switches support receiving and forwarding routed or bridged Ethernet frames with 802.1Q VLAN tags. The logical interface on which untagged packets are received must have the same native VLAN ID as that on the physical interface.

To configure the native VLAN ID using the CLI:

1. Configure the port mode as **trunk** so that the interface is on multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

```
[edit interfaces xe-0/0/3 unit 0 family ethernet-switching]
```

```
user@switch# set port-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces xe-0/0/3 unit 0 family ethernet-switching]
```

```
user@switch# set native-vlan-id 3
```

Related Documentation

- [Understanding Bridging on page 1869](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)

Configuring VLANs

Switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Specify the description of the VLAN:

```
[edit interfaces xe-chassis/pic/port unit 0]
user@switch# set description vlan-description
```

2. Specify the unique name of the VLAN:



NOTE: In a QFabric system, do not configure “default” as the name of a VLAN. Though the QFabric system will allow you to configure and commit a VLAN with the name “default” in the current software with no commit errors, it will not work. Junos OS 12.2 and onwards will not allow you to commit a VLAN with the name “default.”

```
[edit interfaces xe-chassis/pic/port unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

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

4. Configure the VLAN tag ID or VLAN ID range for the VLAN:

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

or
```

```
[edit vlans]
user@switch# set vlan-name vlan-range vlan-id-low-vlan-id-high
```

5. Specify the maximum time that an entry can remain in the forwarding table before it ages out:

```
[edit vlans]
user@switch# set vlan-name mac-table-aging-time time
```

6. Specify a VLAN firewall filter to be applied to incoming or outgoing packets:

```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Configuring Routed VLAN Interfaces on page 2046](#)
- [Creating a Series of Tagged VLANs on page 2050](#)

- [Understanding Bridging on page 1869](#)

Creating a Series of Tagged VLANs

When you divide an Ethernet LAN into multiple VLANs, each VLAN is assigned a unique IEEE 802.1Q tag. This tag is associated with each frame in the VLAN, and the network nodes receiving the traffic can use the tag to identify which VLAN a frame is associated with.

Instead of configuring VLANs and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames can detect which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN **employee** and specify a tag range of **10 through 12**. This creates the following VLANs and tags:

- VLAN **employee-10**, tag 10
- VLAN **employee-11**, tag 11
- VLAN **employee-12**, tag 12

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.

To configure a series of tagged VLANs using the CLI (here, the VLAN is **employee**):

1. Configure the series (here, a VLAN series from 120 through 130):

```
[edit]
user@switch# set vlans employee vlan-range 120-130
```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlanmembers employee
```

- Include the VLAN range:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlan members 120-130
```

Associating a series of tagged VLANs to an interface by name or by VLAN range has the same result: VLANs **__employee_120__** through **__employee_130__** are created.



NOTE: When a series of VLANs is created using the `vlan-range` command, the VLAN names are preceded and followed by a double underscore.

Related Documentation

- [Verifying That a Series of Tagged VLANs Has Been Created on page 2173](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Understanding Bridging on page 1869](#)

Configuring Multiple VLAN Registration Protocol

Multiple VLAN Registration Protocol (MVRP) automates the creation and management of VLANs. When using MVRP on a QFabric system, you must manually create on the QFabric the VLANs that exist on the attached servers because the QFabric implementation of MVRP does not allow VLANs to be created dynamically. However, you do not need to manually assign VLAN membership to the QFabric ports that connect to the servers. MVRP automatically assigns VLAN membership to server-facing QFabric ports when it learns about a VLAN from an attached server.

MVRP is disabled by default. To enable MVRP or set MVRP options, follow these instructions:

- [Enabling MVRP on page 2051](#)
- [Disabling MVRP on page 2051](#)
- [Configuring Timer Values on page 2052](#)
- [Configuring Passive Mode on page 2052](#)

Enabling MVRP

MVRP can be enabled only on trunk interfaces. To enable MVRP on a trunk interface:

```
[edit protocols mvrp]
user@qfabric# set interface interface-name
```



NOTE: On QFX Series switches, you must configure specific interfaces—you cannot specify `interface all`. You can enable MVRP on an interface range.

Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

To disable MVRP on the entire QFabric system:

```
[edit protocols mvrp]
user@qfabric# set disable
```

To disable MVRP on a specific trunk interface:

```
[edit protocols mvrp]
user@qfabric# set disable interface interface-name
```

Configuring Timer Values

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the switch after receiving an MVRP PDU. The join timer controls the amount of time the switch waits to accept a registration request, the leave timer controls the period of time that the switch waits in the Leave state before changing to the unregistered state, and the leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer.



BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

To set the join timer for an interface:

```
[edit protocols mvrp]
user@qfabric# set interface interface-name 300
```

To set the leave timer for an interface:

```
[edit protocols mvrp]
user@qfabric# set interface interface-name leave-timer 1200
```

To set the leaveall timer for an interface:

```
[edit protocols mvrp]
user@qfabric# set interface interface-name leaveall-timer 12000
```

Configuring Passive Mode

QFX switches include a mode—called passive mode—in which an MVRP-configured interface does not announce its membership in a VLAN or send any VLAN declarations (updates) unless it receives registration for that VLAN from a peer (server).

To configure an interface to operate in passive mode:

```
[edit protocols mvrp]
user@qfabric# set interface interface-name passive
```

Related Documentation

- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on page 1897](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on page 1953](#)
- [Verifying That MVRP Is Working Correctly on page 2182](#)

Configuring STP

The default spanning-tree protocol on the QFX Series is Rapid Spanning Tree Protocol (RSTP). RSTP provides faster convergence times than Spanning Tree Protocol (STP) does. However, some legacy networks require the slower convergence times of basic STP.

If your network includes 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. When you explicitly configure STP, the QFX Series products use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP.

To configure STP using the CLI:

1. Delete the RSTP configuration on the interface (here, the interface is **xe-0/0/5**):

```
[edit]
user@switch# delete protocols rstp interface xe-0/0/5
```

2. Configure STP on the interface:

```
[edit]
user@switch# set protocols stp interface xe-0/0/5
```

3. Commit the configuration:

```
[edit]
user@switch# commit
```

Related Documentation

- [show spanning-tree bridge on page 2235](#)
- [show spanning-tree interface on page 2240](#)
- [Overview of Spanning-Tree Protocols on page 1900](#)

Unblocking an Interface That Receives BPDUs in Error



NOTE: BPDU block protection is disabled on Node devices.

QFX Series switches use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface transitions to a blocking state and stops forwarding frames.

After you fix the misconfiguration that triggered the sending of BPDUs to an interface, you can unblock the interface and return it to service.

To unblock an interface after fixing the misconfiguration that triggered the BPDUs and return it to service using the CLI:

- Automatically unblock an interface by configuring a timer that expires (here, the interface is **xe-0/0/6**):

```
[edit ethernet-switching-options]
user@switch# set bpd-blockdisable-timeout 30 interface xe-0/0/6
```

- Manually unblock an interface using the operational mode command:

```
user@switch> clear ethernet-switching bpd-error interface xe-0/0/6
```

Related Documentation

- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903](#)

Configuring VLAN Spanning Tree Protocol

VLAN Spanning Tree Protocol (VSTP) enables the QFX Series to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining the best paths within the VLANs instead of within the entire network.

To configure VSTP:

- (Optional) Enable Rapid Spanning Tree Protocol (RSTP):

```
[edit protocols]
user@switch# set rstp
```

VSTP can run on a maximum of 253 VLANs; RSTP runs on the remaining VLANs if configured. Enabling RSTP ensures that a spanning-tree protocol runs on all VLANs.

- Enable VSTP.

- To enable VSTP on multiple VLANs using a VLAN group:

```
[edit protocols]
user@switch# set vstp vlan-group group group-name vlan vlan-id-range
```

- To enable VSTP on all VLANs:

```
[edit protocols]
user@switch# set vstp vlan all
```



NOTE: You must enable RSTP if you used the `set vstp vlan all` statement to enable VSTP and if the switch has more than 253 VLANs. If you use the `set vstp vlan all` statement to enable VSTP on a switch with more than 253 VLANs, the configuration cannot be committed.

- To enable VSTP on a VLAN using a single VLAN ID:

```
[edit protocols]
user@switch# set vstp vlan vlan-id
```

- To enable VSTP on a VLAN using a single VLAN name:

```
[edit protocols]
user@switch# set vstp vlan vlan-name
```

Related Documentation

- [Understanding VSTP on page 1902](#)

Disabling MAC Learning

By default, MAC learning is globally enabled on all nodes in a QFX Series product. This topic describes how to disable MAC learning, as well as how to reenable and verify that MAC learning has been enabled or disabled.

Disabling dynamic MAC learning on the QFX Series prevents a node from learning source and destination MAC addresses.

- To disable MAC learning on the QFX Series:

```
[edit ethernet-switching-options interfaces interface]
user@switch# set no-mac-learning
```

- To enable MAC learning on the QFX Series:

```
[edit ethernet-switching-options interfaces interface]
user@switch# delete no-mac-learning
user@switch# deactivate no-mac-learning
```

- To verify the status of MAC learning on the QFX Series, view the Ethernet MAC learning statistics in operational mode.

```
user@switch> show ethernet-switching statistics mac-learning
Ethernet-switching table: 2 entries, 1 learned
  VLAN      MAC address      Type      Age  Interfaces
  default   *                Flood     -    All-members
  default   00:1f:12:39:90:80 Learn     29   xe-/0/0.0
```

Related Documentation

- [Understanding MAC Learning on page 1876](#)
- [Example: Disabling MAC Learning on page 2003](#)
- [no-mac-learning \(Global\) on page 2109](#)

Disabling MAC Learning in a VLAN

By default, MAC learning is enabled on a VLAN. This topic describes how to disable MAC learning in a VLAN, as well as how to reenable and verify that MAC learning has been enabled or disabled.

Disabling dynamic MAC learning in a VLAN on a QFX Series product prevents a node from learning source and destination MAC addresses.

- To disable MAC learning in a VLAN:

```
[edit vlans vlan-name]  
user@switch# set no-mac-learning
```

- To reenable MAC learning in a VLAN, use either of the following two commands:

```
[edit vlans vlan-name]  
user@switch# delete no-mac-learning  
user@switch# deactivate no-mac-learning
```

- To verify the status of MAC learning on the QFX series:

```
user@switch> show ethernet-switching table
```

Related Documentation

- [Understanding MAC Learning on page 1876](#)
- [Example: Disabling MAC Learning in a VLAN on page 2003](#)
- [no-mac-learning \(Per VLAN\) on page 2109](#)

Creating a Private VLAN on a Single Switch

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature allows you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a secondary VLAN inside a primary VLAN. This topic describes how to configure a PVLAN on a single switch.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (You do not need to preconfigure the primary VLAN—it is configured as part of this procedure.) You do not need to create VLAN IDs (tags) for the secondary VLANs. It does not impair functioning if you tag the secondary VLANs, but tags are not used when secondary VLANs are configured on a single switch.

Keep these rules in mind when configuring a PVLAN:

- The primary VLAN must be a tagged VLAN.
- If you are going to configure a community VLAN, you must first configure the primary VLAN and the PVLAN trunk port. You must also configure the primary VLAN to be private using the **pvlan** statement.
- If you are going to configure an isolated VLAN, you must first configure the primary VLAN and the PVLAN trunk port.

If you complete your configuration steps in the order shown, you will not violate these PVLAN rules. To configure a private VLAN on a single switch:

1. Set the name and VLAN ID (802.1Q tag) for the primary VLAN:

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

2. Configure the VLAN to be private:

```
[edit vlans]
user@switch# set primary-vlan-name pvlan
```

3. Configure the trunk interfaces for the primary VLAN:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

4. Add the trunk interfaces to the primary VLAN:

```
[edit vlans]
user@switch# set primary-vlan-name interface interface-name
```

5. Configure the access interfaces for the community (secondary) VLANs:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode access
```

6. Add the access interfaces to the community VLANs:

```
[edit vlans]
user@switch# set community-vlan-name interface interface-name
```

7. For each community VLAN, set the primary VLAN:

```
[edit vlans]
```

```
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

8. Configure isolated ports:

```
[edit vlans]
user@switch# set primary-vlan-name interface interface-name isolated
```

**Related
Documentation**

- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Verifying That a Private VLAN Is Working on page 2176](#)

Creating a Private VLAN Spanning Multiple Switches

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature allows you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a secondary VLAN inside a primary VLAN. This topic describes how to configure a PVLAN to span multiple switches.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (You do not need to preconfigure the primary VLAN—it is configured as part of this procedure.) You do not need to create VLAN IDs (tags) for the secondary VLANs. It does not impair functioning if you tag the secondary VLANs, but tags are not used when secondary VLANs are configured on a single switch.

The following rules apply to creating PVLANS:

- The primary VLAN must be a tagged VLAN.
- If you are going to configure a community VLAN, you must first configure the primary VLAN and the PVLAN trunk port. You must also configure the primary VLAN to be private using the **pvlan** statement.
- If you are going to configure an isolated VLAN, you must first configure the primary VLAN and the PVLAN trunk port.

If you complete your configuration steps in the order shown, you will not violate these PVLAN rules. To configure a private VLAN to span multiple switches:

1. Set the name and VLAN ID (802.1Q tag) for the primary VLAN:

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

2. Configure the VLAN to be private:

```
[edit vlans]
user@switch# set primary-vlan-name pvlan
```

3. Configure the trunk interfaces for the primary VLAN:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

4. Add the trunk interfaces to the primary VLAN:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name interface interface-name
```

5. Configure the access interfaces for the community (secondary) VLANs:

```
[edit interfaces]
```

```
user@switch# set interface-name unit 0 family ethernet-switching port-mode access
```

6. Add the access interfaces to the community VLANs:

```
[edit vlans]
```

```
user@switch# set community-vlan-name interface interface-name
```

7. For each community VLAN, set the primary VLAN:

```
[edit vlans]
```

```
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

8. Configure an isolated VLAN ID to create an interswitch isolated domain that spans the switches:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name isolation-vlan-id number
```

9. Configure isolated ports:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name interface interface-name isolated
```

Related Documentation

- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Verifying That a Private VLAN Is Working on page 2176](#)

Configuring the Forwarding Mode

By default, QFX Series products forward packets using store-and-forward mode. You can configure all the interfaces on a QFX series switch or QFabric to use cut-through mode instead.

To enable cut-through switching mode, enter the following statement:

```
[edit forwarding-options]
```

```
user@switch# set cut-through
```

Related Documentation

- [cut-through on page 2080](#)

Private VLAN Configuration Tasks

- [Creating a Private VLAN on a Single Switch on page 2060](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2061](#)

Creating a Private VLAN on a Single Switch

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature allows you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a secondary VLAN inside a primary VLAN. This topic describes how to configure a PVLAN on a single switch.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (You do not need to preconfigure the primary VLAN—it is configured as part of this procedure.) You do not need to create VLAN IDs (tags) for the secondary VLANs. It does not impair functioning if you tag the secondary VLANs, but tags are not used when secondary VLANs are configured on a single switch.

Keep these rules in mind when configuring a PVLAN:

- The primary VLAN must be a tagged VLAN.
- If you are going to configure a community VLAN, you must first configure the primary VLAN and the PVLAN trunk port. You must also configure the primary VLAN to be private using the **pvlan** statement.
- If you are going to configure an isolated VLAN, you must first configure the primary VLAN and the PVLAN trunk port.

If you complete your configuration steps in the order shown, you will not violate these PVLAN rules. To configure a private VLAN on a single switch:

1. Set the name and VLAN ID (802.1Q tag) for the primary VLAN:

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

2. Configure the VLAN to be private:

```
[edit vlans]
user@switch# set primary-vlan-name pvlan
```

3. Configure the trunk interfaces for the primary VLAN:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

4. Add the trunk interfaces to the primary VLAN:

```
[edit vlans]
user@switch# set primary-vlan-name interface interface-name
```

5. Configure the access interfaces for the community (secondary) VLANs:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode access
```

6. Add the access interfaces to the community VLANs:

```
[edit vlans]
user@switch# set community-vlan-name interface interface-name
```

7. For each community VLAN, set the primary VLAN:

```
[edit vlans]
```



```
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

8. Configure isolated ports:

```
[edit vlans]
user@switch# set primary-vlan-name interface interface-name isolated
```

Related Documentation

- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
- [Verifying That a Private VLAN Is Working on page 2176](#)

Creating a Private VLAN Spanning Multiple Switches

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature allows you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a secondary VLAN inside a primary VLAN. This topic describes how to configure a PVLAN to span multiple switches.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (You do not need to preconfigure the primary VLAN—it is configured as part of this procedure.) You do not need to create VLAN IDs (tags) for the secondary VLANs. It does not impair functioning if you tag the secondary VLANs, but tags are not used when secondary VLANs are configured on a single switch.

The following rules apply to creating PVLANS:

- The primary VLAN must be a tagged VLAN.
- If you are going to configure a community VLAN, you must first configure the primary VLAN and the PVLAN trunk port. You must also configure the primary VLAN to be private using the `pvlan` statement.
- If you are going to configure an isolated VLAN, you must first configure the primary VLAN and the PVLAN trunk port.

If you complete your configuration steps in the order shown, you will not violate these PVLAN rules. To configure a private VLAN to span multiple switches:

1. Set the name and VLAN ID (802.1Q tag) for the primary VLAN:

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

2. Configure the VLAN to be private:

```
[edit vlans]
user@switch# set primary-vlan-name pvlan
```

3. Configure the trunk interfaces for the primary VLAN:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

4. Add the trunk interfaces to the primary VLAN:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name interface interface-name
```

5. Configure the access interfaces for the community (secondary) VLANs:

```
[edit interfaces]
```

```
user@switch# set interface-name unit 0 family ethernet-switching port-mode access
```

6. Add the access interfaces to the community VLANs:

```
[edit vlans]
```

```
user@switch# set community-vlan-name interface interface-name
```

7. For each community VLAN, set the primary VLAN:

```
[edit vlans]
```

```
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

8. Configure an isolated VLAN ID to create an interswitch isolated domain that spans the switches:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name isolation-vlan-id number
```

9. Configure isolated ports:

```
[edit vlans]
```

```
user@switch# set primary-vlan-name interface interface-name isolated
```

Related Documentation

- [Understanding Private VLANs on page 1878](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 1883](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Verifying That a Private VLAN Is Working on page 2176](#)

Q-in-Q Tunneling Configuration Tasks

- [Configuring Q-in-Q Tunneling on page 2062](#)
- [Configuring Layer 2 Protocol Tunneling on page 2063](#)

Configuring Q-in-Q Tunneling

Q-in-Q tunneling and VLAN translation allow service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Data centers can use Q-in-Q tunneling to isolate customer traffic within a single site or when customer traffic flows between cloud data centers in different geographic locations.

Before you begin setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs on the neighboring switches. See [“Configuring VLANs” on page 2048](#).

To configure Q-in-Q tunneling:

1. Create the service VLAN (S-VLAN) and configure an ID for it:

```
[edit vlans]
```

```
user@switch# set s-vlan-name vlan-id s-vlan-ID
```

2. Enable Q-in-Q tunneling on the S-VLAN:

```
[edit vlans]
```

```
user@switch# set s-vlan-name dot1q-tunneling
```

3. Set the allowed customer VLANs (C-VLANs) on the S-VLAN (optional). Here, the C-VLANs are identified by a range:

```
[edit vlans]
```

```
user@switch# set s-vlan-name dot1q-tunneling customer-vlans range
```

4. Configure a global value for the tag protocol identifier (EtherType) of the service VLAN tags (optional):

```
[edit]
```

```
user@switch# set ethernet-switching-options dot1q-tunneling ether-type ether-type-value
```

Depending on your interface configuration, you might need to adjust the MTU value on your trunk or access ports to accommodate the 4 bytes used for the tag added by Q-in-Q tunneling. For example, if you use the default MTU value of 1514 bytes on your access and trunk ports, you need to make one of the following adjustments:

- Reduce the MTU on the access links by at least 4 bytes so that the frames do not exceed the MTU of the trunk link when S-VLAN tags are added.
- Increase the MTU on the trunk link so that the link can handle the larger frame size.

Related Documentation

- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 2266](#)
- [mtu on page 2603](#)

Configuring Layer 2 Protocol Tunneling

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to switches at a remote location. This feature is useful when you have a network that includes remote sites that are connected across a service provider network and you want to run Layer 2 protocols on switches connected across the service provider network.

Tunneled Layer 2 PDUs do not normally arrive at high rate. If the tunneled Layer 2 PDUs do arrive at high rate, there might be a problem in the network. Typically, you would want to shut down the interface that is receiving a high rate of tunneled Layer 2 PDUs so that the problem can be isolated. You do so using the **shutdown-threshold** statement. However, if you do not want to completely shut down the interface, you can configure the switch to drop tunneled Layer 2 PDUs that exceed a certain threshold using the **drop-threshold** statement.

L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate-limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets are not reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because

of the rate limit will not be reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

There are no default settings for **drop-threshold** and **shutdown-threshold**. If you do not specify these thresholds, then no thresholds are enforced. As a result, the switch tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

You can specify a drop threshold value without specifying a shutdown threshold value, and you can specify a shutdown threshold value without specifying a drop threshold value. If you specify both threshold values, then the drop threshold value must be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit operation fails.



NOTE: L2PT and VLAN translation configured with the **mapping** statement cannot both be configured on the same switch.



NOTE: If the switch receives untagged Layer 2 control PDUs to be tunneled, then you must configure the switch to map untagged (native) packets to an L2PT-enabled VLAN. Otherwise, the untagged Layer 2 control PDU packets are discarded. For more information, see “[Understanding Q-in-Q Tunneling and VLAN Translation](#)” on page 1906.

To configure L2PT:

1. Because L2PT operates under the Q-in-Q tunneling configuration, you must enable Q-in-Q tunneling before you can configure L2PT. Enable Q-in-Q tunneling on VLAN **customer-1**:

[edit]

user@switch# **set vlans customer-1 dot1q-tunneling**

2. Enable L2PT for the Layer 2 protocol you want to tunnel, on the VLAN:

- To enable L2PT for a specific protocol (here, STP):

[edit]

user@switch# **set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp**

- To enable L2PT for all supported protocols:

[edit]

user@switch# **set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling all**

3. (Optional) Configure the drop threshold:



NOTE: If you also configure the shutdown threshold, ensure that you configure the drop threshold value to be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration changes, the commit operation fails.

[edit]

```
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
drop-threshold 50
```

4. (Optional) Configure the shutdown threshold:



NOTE: If you also configure the drop threshold, ensure that you configure the shutdown threshold value to be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value and you try to commit the configuration changes, the commit operation fails.

[edit]

```
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
shutdown-threshold 100
```



NOTE: Once an interface is disabled, you must explicitly reenabling it using the `clear ethernet-switching layer2-protocol-tunneling error` command. Otherwise, the interface remains disabled.

Related Documentation

- [Understanding Layer 2 Protocol Tunneling on page 1910](#)
- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)

Proxy ARP Configuration Task

- [Configuring Proxy ARP on page 2065](#)

Configuring Proxy ARP

You can configure proxy Address Resolution Protocol (ARP) to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

[edit interfaces]

```
user@switch# set xe-0/0/3 unit 0 proxy-arp restricted
```



BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch is not a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid the situation of the switch's response to a gratuitous ARP request appearing to the host to be an indication of an IP conflict:

To configure proxy ARP on a routed VLAN interface (RVI):

```
[edit interfaces]
user@switch# set vlan unit 100 proxy-arp restricted
```

**Related
Documentation**

- [Understanding Proxy ARP on page 1914](#)
- [Verifying That Proxy ARP Is Working Correctly on page 2181](#)
- [Understanding Routed VLAN Interfaces on page 1875](#)

Configuration Statements

- [address on page 2069](#)
- [alarm \(STP\) on page 2071](#)
- [arp \(Interfaces\) on page 2072](#)
- [block on page 2073](#)
- [bpdu-block on page 2074](#)
- [bpdu-block-on-edge on page 2075](#)
- [bpdu-timeout-action on page 2076](#)
- [bridge-priority on page 2077](#)
- [broadcast on page 2078](#)
- [cost \(STP\) on page 2079](#)
- [configuration-name \(MSTP\) on page 2080](#)
- [cut-through on page 2080](#)
- [description \(VLAN\) on page 2081](#)
- [disable \(STP\) on page 2082](#)
- [disable-timeout \(BPDU\) on page 2083](#)
- [drop-threshold on page 2084](#)
- [edge \(STP\) on page 2085](#)
- [ethernet-switching-options on page 2086](#)
- [eui-64 on page 2088](#)
- [filter \(VLANs\) on page 2088](#)
- [force-version on page 2089](#)
- [forward-delay on page 2090](#)

- [group-limit](#) on page 2091
- [hello-time](#) on page 2092
- [interface \(BPDU\)](#) on page 2093
- [interface \(Spanning Trees\)](#) on page 2094
- [interface \(Storm Control\)](#) on page 2095
- [interface \(STP\)](#) on page 2096
- [interface \(VLANs\)](#) on page 2097
- [interfaces](#) on page 2097
- [l3-interface \(VLAN\)](#) on page 2098
- [mac-limit](#) on page 2099
- [mac-notification](#) on page 2100
- [mac-table-aging-time](#) on page 2101
- [max-age](#) on page 2102
- [max-hops](#) on page 2103
- [members](#) on page 2104
- [mode \(STP\)](#) on page 2105
- [msti](#) on page 2106
- [mstp](#) on page 2107
- [native-vlan-id](#) on page 2108
- [no-mac-learning \(Global\)](#) on page 2109
- [no-mac-learning \(Per VLAN\)](#) on page 2109
- [no-root-port](#) on page 2110
- [notification-interval](#) on page 2111
- [port-mode](#) on page 2112
- [preferred](#) on page 2113
- [primary \(Address on Interface\)](#) on page 2114
- [priority \(STP\)](#) on page 2115
- [protocols](#) on page 2116
- [proxy-arp](#) on page 2129
- [reflective-relay](#) on page 2130
- [revision-level](#) on page 2130
- [rstp](#) on page 2131
- [storm-control](#) on page 2132
- [stp](#) on page 2133
- [traceoptions \(Ethernet Switching Options\)](#) on page 2134
- [traceoptions \(STP\)](#) on page 2137
- [unknown-unicast-forwarding](#) on page 2140

- [vlan \(Ethernet\) on page 2141](#)
- [vlan \(STP\) on page 2142](#)
- [vlan \(Unknown Unicast\) on page 2143](#)
- [vlan-id \(VLANs\) on page 2144](#)
- [vlan-range on page 2144](#)
- [vlans on page 2145](#)
- [vlan-tagging on page 2146](#)
- [vstp on page 2147](#)

address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcidlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 family *family*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the interface address.

Options *address*—Address of the interface.

The remaining statements are explained separately.



NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Protocol Family*
- *negotiate-address*
- *unnumbered-address (Ethernet)*
- *Junos OS System Basics Configuration Guide*
- *family*

alarm (STP)

Syntax	alarm;
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols rstp interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols stp interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols vstp vlan <i>vlan-id</i> interface (all <i>interface-name</i>) bpdu-timeout-action]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces configured for loop protection, configure the software to generate a message to be sent to the system log file to record the loop-protection event.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999 • Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface

arp (Interfaces)

Syntax	<code>arp ip-address (mac multicast-mac) mac-address publish;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.
Options	<p>ip-address—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.</p> <p>mac mac-address—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>multicast-mac mac-address—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.</p>



NOTE: The edit **logical-systems** hierarchy is not available on QFabric systems.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Static ARP Table Entries • Configuring Static ARP Entries on page 2045

block

Syntax	block;
Hierarchy Level	[edit protocols mstp (Spanning Trees) interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols rstp interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols stp interface (all <i>interface-name</i>) bpdu-timeout-action], [edit protocols vstp vlan <i>vlan-id</i> interface (all <i>interface-name</i>) bpdu-timeout-action]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure loop protection on a specific interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999 • Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • <i>show spanning-tree interface</i>

bpdu-block

Syntax bpdu-block {
 [interface](#) (all | [*interface-name*]);
 [disable-timeout](#) *timeout*;
 }

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure BPDU protection on an interface. If the interface receives BPDUs, it is disabled.



NOTE: BPDU block protection is disabled on Node devices.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
- [Unblocking an Interface That Receives BPDUs in Error on page 2054](#)
- [clear ethernet-switching bpdu-error on page 2185](#)
- [show spanning-tree bridge on page 2235](#)
- [show spanning-tree interface](#)

bpdu-block-on-edge

Syntax	bpdu-block-on-edge;
Hierarchy Level	[edit protocols mstp], [edit protocols rstp], [edit protocols vstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding VSTP on page 1902 • Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • clear ethernet-switching bpdu-error on page 2185 • show spanning-tree bridge on page 2235 • show spanning-tree interface

bpdu-timeout-action

Syntax	<pre>bpdu-timeout-action { alarm; block; }</pre>
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>)], [edit protocols rstp interface (all <i>interface-name</i>)], [edit protocols stp interface (all <i>interface-name</i>)], [edit protocols vstp vlan <i>vlan-id</i> interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the BPDU timeout action on a specific interface. You must configure at least one action (alarm, block, or both).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1999• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1904• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface

bridge-priority

Syntax	<code>bridge-priority <i>priority</i>;</code>
Hierarchy Level	[edit protocols mstp], [edit protocols mstp <i>msti-id</i>], [edit protocols rstp], [edit protocols stp], [edit protocols vstp <i>vlan</i> <i>vlan-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.
Options	<i>priority</i> —Bridge priority. It can be set only in increments of 4096. Range: 0 through 61,440 Default: 32,768
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Understanding MSTP on page 1901 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface

broadcast

Syntax	<code>broadcast address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.
Default	The default broadcast address has a host portion of all ones.
Options	address —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255 .



NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 2040

cost (STP)

Syntax	<code>cost cost;</code>
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>)], [edit protocols mstp msti msti-id interface interface-name], [edit protocols rstp (Spanning Trees) interface (all <i>interface-name</i>)], [edit protocols stp interface (all <i>interface-name</i>)], [edit protocols vstp vlan vlan-id interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.
Default	Link cost is determined by the link speed.
Options	cost —Link cost associated with the port. Range: 1 through 200,000,000
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding MSTP on page 1901 • Overview of Spanning-Tree Protocols on page 1900 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface

configuration-name (MSTP)

Syntax	configuration-name <i>configuration-name</i> ;
Hierarchy Level	[edit protocols mstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the configuration name. The configuration name is the MSTP region name carried in the MSTP BPDUs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Understanding MSTP on page 1901• show spanning-tree bridge on page 2235• <i>show spanning-tree interface</i>

cut-through

Syntax	cut-through;
Hierarchy Level	[edit forwarding-options]
Description	Configures all the interfaces in the QFX series switch or QFabric to use cut-through forwarding mode instead of store-and-forward mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Forwarding Mode on page 2059

description (VLAN)

Syntax	<code>description <i>text-description</i>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Provide a textual description for the VLAN. The text has no effect on the operation of the VLAN or switch.
Options	<i>text-description</i> —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962• Understanding Bridging on page 1869• show vlans on page 2251


disable (STP)

Syntax	disable;
Hierarchy Level	[edit protocols mstp], [edit protocols mstp interface <i>interface-name</i>], [edit protocols mstp msti <i>msti-id</i> vlan (<i>vlan-id</i> <i>vlan-name</i>) interface <i>interface-name</i>], [edit protocols rstp], [edit protocols rstp interface <i>interface-name</i>], [edit protocols stp], [edit protocols stp interface <i>interface-name</i>], [edit protocols vstp], [edit protocols vstp vlan <i>vlan-id</i> interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable STP, MSTP, RSTP, or VSTP on the switch or on a specific interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding MSTP on page 1901• Overview of Spanning-Tree Protocols on page 1900• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface

disable-timeout (BPDU)

Syntax	<code>disable-timeout <i>timeout</i>;</code>
Hierarchy Level	[edit ethernet-switching-options bpd-block]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.
Default	The disable timeout is not enabled.
Options	<p>timeout: Length of time, in seconds, that the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service.</p> <p>Range: 10 through 3600 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990 • Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903 • show spanning-tree bridge on page 2235 • show spanning-tree interface on page 2240

drop-threshold

Syntax	<code>drop-threshold <i>number</i>;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i> dot1q-tunneling layer2-protocol-tunneling (all <i>protocol-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold value must be less than or equal to the shutdown threshold value.</p> <p>L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate-limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets are not reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit will not be reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.</p>
	<p> NOTE: If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit operation fails.</p>
	You can specify a drop threshold value without specifying a shutdown threshold value.
Default	No drop threshold is specified.
Options	<p><i>number</i>—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs.</p> <p>Range: 1 through 1000</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</i> • <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i> • Configuring Layer 2 Protocol Tunneling on page 2063 • shutdown-threshold on page 2168

edge (STP)

Syntax	edge;
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>)], [edit protocols mstp msti msti-id interface interface-name], [edit protocols rstp interface (all <i>interface-name</i>)], [edit protocols stp interface (all <i>interface-name</i>)], [edit protocols vstp vlan vlan-id interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge interfaces. Edge interfaces immediately transition to a forwarding state.
Default	Edge interfaces are not enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Understanding MSTP on page 1901 • Overview of Spanning-Tree Protocols on page 1900 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface

ethernet-switching-options

```
Syntax ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
      output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds {
  }
  port-error-disable {
    disable-timeout timeout;
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      fcoe-trusted;
      mac-limit limit action action;
      no-allowed-mac-log;
    }
    vlan (all | vlan-name) {
      (arp-inspection | no-arp-inspection) [
        forwarding-class (for DHCP Snooping or DAI Packets) class-name;
      ]
    }
  }
}
```

```

dhcp-option82 {
  circuit-id {
    prefix (Circuit ID for Option 82) hostname;
    use-interface-description;
    use-vlan-id;
  }
  remote-id {
    prefix (Remote ID for Option 82) hostname | mac | none;
    use-interface-description;
    use-string string;
  }
  vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp) {
  forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
  examine-vn2vn {
    beacon-period milliseconds;
  }
  fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
 - [Overview of Access Port Protection on page 4674](#)
 - [Understanding Storm Control on page 4694](#)

eui-64

Syntax	eui-64;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring the Interface Address on page 2040

filter (VLANs)

Syntax	filter (input output) <i>filter-name</i> ;
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply a firewall filter to traffic ingressing or egressing a VLAN.
Default	All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.
Options	<i>filter-name</i> —Name of a firewall filter defined at the [edit firewall family <i>family-name</i> filter] hierarchy level. input —Apply a firewall filter to VLAN ingress traffic. output —Apply a firewall filter to VLAN egress traffic.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring Firewall Filters on page 4747 • Overview of Firewall Filters on page 4635

force-version

Syntax	force-version stp;
Hierarchy Level	[edit protocols vstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Force VLAN Spanning Tree Protocol (VSTP) to use the STP protocol instead of the default protocol, RSTP.
Options	stp —Spanning Tree Protocol
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show spanning-tree bridge on page 2235• <i>show spanning-tree interface</i>• Understanding VSTP on page 1902

forward-delay

Syntax	<code>forward-delay <i>seconds</i>;</code>
Hierarchy Level	[edit protocols mstp], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan <i>vlan-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify how long a bridge interface remains in the listening and learning states before transitioning to the forwarding state.
Options	<i>seconds</i> —Number of seconds the bridge interface remains in the listening and learning states. Range: 4 through 30 seconds Default: 15 seconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding MSTP on page 1901• Overview of Spanning-Tree Protocols on page 1900• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface

group-limit

Syntax	<code>group-limit <i>limit</i>;</code>
Hierarchy Level	[edit protocols igmp-snooping vlan (IGMP Snooping) <i>vlan-id</i> <i>vlan-number</i> interface (IGMP Snooping) <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a limit for the number of multicast groups allowed on the specified interface. After this limit is reached, new reports are ignored and related flows are not flooded on the interface.
Default	No group limits are configured.
Options	<i>limit</i> —Number that represents the maximum number of multicast groups allowed on the specified interface. Range: 0 through 65535
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IGMP Snooping on EX Series Switches</i> • <i>Configuring IGMP Snooping (CLI Procedure)</i> • <i>Configuring IGMP Snooping (J-Web Procedure)</i> • <i>group (IGMP Snooping)</i>

hello-time

Syntax	<code>hello-time <i>seconds</i>;</code>
Hierarchy Level	<code>[edit protocols <i>mstp</i>],</code> <code>[edit protocols <i>rstp</i>],</code> <code>[edit protocols <i>vstp</i>],</code> <code>[edit protocols <i>vstp</i> vlan <i>vlan-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the time interval at which the root bridge transmits configuration BPDUs.
Options	<i>seconds</i> —Number of seconds between transmissions of configuration BPDUs. Range: 1 through 10 seconds Default: 2 seconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding MSTP on page 1901• Overview of Spanning-Tree Protocols on page 1900• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface

interface (BPDU)

Syntax	interface (all <i>interface-name</i>);
Hierarchy Level	[edit ethernet-switching-options bpd-block]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply BPDU protection to all interfaces or one or more interfaces.
Options	<p>all—All interfaces.</p> <p><i>interface-name</i>—Name of the interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240

interface (Spanning Trees)

Syntax	<pre>interface <i>interface-name</i> { arp-on-stp; bpdu-timeout-action block; log; cost <i>cost</i>; disable; edge; mode <i>mode</i>; no-root-port; priority <i>priority</i>; }</pre>
Hierarchy Level	<pre>[edit protocols mstp], [edit protocols mstp msti <i>msti-id</i>], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan (all <i>vlan-id</i> <i>vlan-name</i>)]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.</p> <p>The edge, mode, and no-root-port options are not available at the <code>[edit protocols mstp msti <i>msti-id</i>]</code> hierarchy level.</p>
Options	<p><i>interface-name</i>—Name of an interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show spanning-tree bridge• show spanning-tree interface• Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches• Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Configuring VSTP (CLI Procedure)

- [show spanning-tree bridge on page 2235](#)

interface (Storm Control)

Syntax	<pre>interface (all <i>interface-name</i>) { <i>bandwidth bandwidth</i>; no-broadcast; no-multicast; no-unknown-unicast; }</pre>
Hierarchy Level	[edit ethernet-switching-options storm-control]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Apply storm control to all interfaces or to the specified interface.</p> <p>The remaining statement is explained separately.</p>
Default	Storm control is disabled.
Options	<p>all—Apply storm control to all interfaces.</p> <p><i>interface-name</i>—Apply storm control to the specified interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Storm Control on page 4694 • Example: Configuring Storm Control on page 4713

interface (STP)

Syntax	<pre>interface <i>interface-name</i> { disable; cost <i>cost</i>; edge; mode <i>mode</i>; no-root-port; priority <i>priority</i>; }</pre>
Hierarchy Level	[edit protocols mstp], [edit protocols mstp msti], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan <i>vlan-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.
Options	<p><i>interface-name</i>—Name of a Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding RSTP on page 1901• Understanding MSTP on page 1901• Overview of Spanning-Tree Protocols on page 1900• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface

interface (VLANs)

Syntax	<code>interface <i>interface-name</i> { mapping (native (push swap) tag (push swap)); }</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For a specific VLAN, configure an interface.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962 • Configuring VLANs on page 2048 • Understanding Bridging on page 1869

interfaces

Syntax	<code>interfaces <i>interface-name</i> { no-mac-learning; }</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure settings for interfaces that have been assigned to family ethernet-switching .
Options	<p><i>interface-name</i> —Name of an interface that is configured for family ethernet-switching.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

l3-interface (VLAN)

Syntax	<code>l3-interface vlan.logical-interface-number;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between VLANs. Traffic between VLANs must be routed, which requires a common Layer 3 interface.
Default	No Layer 3 (routing) interface is associated with the VLAN.
Options	<code>vlan.logical-interface-number</code> —Number of the logical interface defined with a set interfaces vlan unit command. For the logical interface number, use the same number you configure in the unit statement.



NOTE: This statement is not supported on QFabric systems.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show ethernet-switching interfaces on page 1838• show vlans on page 2251

mac-limit

Syntax	<code>mac-limit <i>number</i>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the number of MAC addresses allowed on a VLAN.
Default	MAC limit is disabled.
Options	<i>number</i> —Maximum number of MAC addresses. Range: 1 through 32768



NOTE: This statement is not supported on QFabric systems.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show vlans on page 2251 • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962 • Configuring MAC Table Aging on page 2044 • Understanding Bridging on page 1869

mac-notification

Syntax	<pre>mac-notification { notification-interval <i>seconds</i>; }</pre>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.</p> <p>The remaining statement is explained separately.</p>
Default	MAC notification is disabled by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring MAC Notification on page 2042

mac-table-aging-time

Syntax	<code>mac-table-aging-time seconds;</code>
Hierarchy Level	[edit ethernet-switching-options], [edit vlans <i>vlan-name</i>]
Release Information	Statement introduced for specific VLANs in Junos OS Release 11.1 for the QFX Series.
Description	<p>Define how long entries remain in the Ethernet switching table before expiring:</p> <ul style="list-style-type: none"> • If you specify this statement at the [ethernet-switching-options] hierarchy level, it applies to all VLANs on the switch. • If you specify this statement at the [vlans] hierarchy level, it applies to the specified VLAN.
Default	300 seconds
Options	<p>seconds—Time that entries remain in the Ethernet switching table before being removed.</p> <ul style="list-style-type: none"> • Range—60 to 1,000,000 seconds. • Default—300 seconds.



NOTE: The `vlans` hierarchy is not available on QFabric systems.

Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962 • Configuring MAC Table Aging on page 2044 • Understanding Bridging on page 1869 • show ethernet-switching statistics aging on page 2205

max-age

Syntax	<code>max-age seconds;</code>
Hierarchy Level	[edit protocols mstp], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan <i>vlan-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the maximum age of received protocol BPDUs.
Options	seconds —Maximum age of received protocol BPDUs. Range: 6 through 40 seconds Default: 20 seconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding MSTP on page 1901• Overview of Spanning-Tree Protocols on page 1900• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240

max-hops

Syntax	<code>max-hops hops;</code>
Hierarchy Level	[edit protocols mstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Multiple Spanning Tree Protocol (MSTP), configure the maximum number of hops that a BPDU can be forwarded in the MSTP region.
Options	<p>hops — Number of hops the BPDU can be forwarded.</p> <p>Range: 1 through 255 hops</p> <p>Default: 20 hops</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Understanding MSTP on page 1901• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240

members

Syntax	<code>members [(all <i>names</i> <i>vlan-ids</i>)];</code>
Hierarchy Level	[edit interfaces <i>ge-fpc/pic/port</i> unit 0 ethernet-switching <i>vlan</i>] [edit interfaces <i>xe-fpc/pic/port</i> unit 0 ethernet-switching <i>vlan</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For trunk interfaces, configure the VLANs for which the interface can carry traffic.



TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after *vlan* or *vlangs* in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options *all*—Specify that this trunk interface be a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.



NOTE: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, *all* cannot be the name of a VLAN on the switch.

names—Names of one or more VLANs.

vlan-ids—Numeric identifiers of one or more VLANs.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Understanding Bridging on page 1869](#)
- [show ethernet-switching interfaces on page 1838](#)
- [show vlans on page 2251](#)

mode (STP)

Syntax	<code>mode mode;</code>
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>)], [edit protocols mstp msti msti-id interface interface-name], [edit protocols rstp interface (all <i>interface-name</i>)], [edit protocols stp interface (all <i>interface-name</i>)], [edit protocols vstp vlan vlan-id interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link mode to identify point-to-point links.
Default	For a full-duplex link, the default link mode is point-to-point . For a half-duplex link, the default link mode is shared .
Options	<i>mode</i> —Link mode: <ul style="list-style-type: none"> • point-to-point—Link is point to point. • shared—Link is shared media.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Understanding MSTP on page 1901 • Overview of Spanning-Tree Protocols on page 1900 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface on page 2240

msti

Syntax	<pre>msti <i>msti-id</i> { vlan (<i>vlan-id</i> <i>vlan-name</i>); interface <i>interface-name</i> { disable; cost <i>cost</i>; edge; mode <i>mode</i>; priority <i>priority</i>; } }</pre>
Hierarchy Level	[edit protocols mstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.
Default	MSTI is disabled.
Options	<p><i>msti-id</i> —MSTI identifier.</p> <p>Range: 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Understanding MSTP on page 1901

mstp

```

Syntax  mstp {
        disable;
        bpdtimeout-action;
        bridge-priority priority;
        configuration-name (MSTP) name;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
            bpdtimeout-action {
                block;
                alarm;
            }
            disable;
            cost cost;
            edge;
            mode mode;
            no-root-port;
            priority priority;
        }
        max-age seconds;
        max-hops hops;
        msti msti-id {
            vlan (vlan-id | vlan-name);
            interface interface-name {
                disable;
                cost cost;
                edge;
                mode mode;
                priority priority;
            }
        }
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
              <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
        revision-level revision-level;
    }

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Multiple Spanning Tree Protocol (MSTP). MSTP is defined in the IEEE 802.1Q-2003 specification and is used to create a loop-free topology in networks with multiple spanning-tree regions.

The statements are explained separately.

Default MSTP is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
- [Understanding MSTP on page 1901](#)
- [show spanning-tree bridge on page 2235](#)
- [show spanning-tree interface on page 2240](#)

native-vlan-id

Syntax native-vlan-id *vlan-id*;

Hierarchy Level [edit [interfaces](#) ge-*fpc/pic/port* unit 0 [ethernet-switching](#)],
[edit [interfaces](#) xe-*fpc/pic/port* unit 0 [ethernet-switching](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the VLAN identifier to associate with untagged packets received on the interface.

Options *vlan-id*—Numeric identifier of the VLAN.
Range: 1 through 4094

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Junos OS Network Interfaces Configuration Guide](#)
- [Understanding Bridging on page 1869](#)
- [show ethernet-switching interfaces on page 1838](#)
- [show vlans on page 2251](#)

no-mac-learning (Global)

Syntax	no-mac-learning;
Hierarchy Level	[edit ethernet-switching-options interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.
Default	MAC learning is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	•

no-mac-learning (Per VLAN)

Syntax	no-mac-learning;
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disables MAC address learning for the specified VLAN.
Default	MAC learning is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-root-port

Syntax	no-root-port;
Hierarchy Level	[edit protocols mstp interface (all <i>interface-name</i>)], [edit protocols rstp interface (all <i>interface-name</i>)], [edit protocols stp interface (all <i>interface-name</i>)], [edit protocols vstp vlan <i>vlan-id</i> interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an interface to be a spanning tree designated port. If the bridge receives more STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. When the bridge stops receiving more STP BPDUs on the root-protected interface, interface traffic is no longer blocked.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1994• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918• Understanding VSTP on page 1902• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240


notification-interval

Syntax	notification-interval <i>seconds</i> ;
Hierarchy Level	[edit ethernet-switching-options mac-notification]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure the MAC notification interval for a switch.</p> <p>The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications are sent to the network management system every 10 seconds.</p>
Options	<p><i>seconds</i>—The MAC notification interval, in seconds.</p> <p>Range: 1 through 60</p> <p>Default: 30</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring MAC Notification on page 2042


port-mode

Syntax	port-mode (access tagged-access trunk);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure whether an interface on the switch operates in access, tagged access, or trunk mode.
Default	All switch interfaces are in access mode.
Options	<p>access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.</p> <p>tagged-access—Have the interface operate in access mode. In this mode, the interface can be in multiple VLANs. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.</p> <p>trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Reflective Relay on page 2044• Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958

preferred

Syntax	preferred;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.
<div>  NOTE: The edit logical-systems hierarchy is not available on QFabric systems. </div>	
Default	The lowest-numbered address on the subnet is the preferred address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interface Address on page 2040

primary (Address on Interface)

Syntax	primary;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</div>	
Default	For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 2040

priority (STP)

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	[edit protocols <code>mstp interface</code> (all <i>interface-name</i>)], [edit protocols <code>mstp msti msti-id interface interface-name</code>], [edit protocols <code>rstp interface</code> (all <i>interface-name</i>)], [edit protocols <code>stp interface</code> (all <i>interface-name</i>)], [edit protocols <code>vstp vlan vlan-id interface</code> (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the interface priority to control which interface is elected as the root port.
Options	priority —Interface priority. The interface priority must be set in increments of 16. Range: 0 through 240 Default: 128
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Network Regions for VLANs with MSTP on page 1931 • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Understanding MSTP on page 1901 • Overview of Spanning-Tree Protocols on page 1900 • Understanding VSTP on page 1902 • show spanning-tree bridge on page 2235 • show spanning-tree interface on page 2240

protocols

```
Syntax protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            hold-down-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family family-name {
            ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
        }
        graceful-restart {
            disable;
            restart-time seconds;
            stale-routes-time seconds;
        }
        group group-name {
            ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
        }
        hold-time seconds;
        import [ policy-names ];
        include-mp-next-hop;
        keep (all | none);
        local-address address;
```



```

local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tll-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}

```

```

    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (Liveness Detection) (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable {
        interface interface-name;
    }
    interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
            group ip-address;
        }
    }
}

```

```

    }
    robust-count number;
  }
}
isis {
  disable;
  export [ policy-names ];
  ignore-attached-bit;
  interface interface-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
    }
    version (1 | automatic);
  }
  checksum;
  csnp-interval (seconds | disable);
  disable;
  hello-padding (adaptive | loose | strict);
  level (1 | 2) {
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    metric metric;
    passive;
    priority number;
  }
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-unicast-topology;
  passive;
  point-to-point;
}
level (1 | 2) {
  disable;
  authentication-key key;
  authentication-type authentication;
  external-preference preference;

```

```
no-csnp-authentication;
no-hello-authentication;
no-psnp-authentication;
preference preference;
prefix-export-limit number;
wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
}
topologies {
    ipv4-multicast;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
```

```

forward-delay seconds;
hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size > <no-stamp | world-readable |
        no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                        meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}

```

```
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
}
(summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
```

```

authentication {
    md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
    simple-password key-string;
}
dead-interval seconds;
demand-circuit;
flood-reduction;
hello-interval seconds;
ipsec-sa sa-name;
no-neighbor-down-notification;
retransmit-interval seconds;
topology (name | default | ipv4-multicast) {
    disable;
    metric metric;
}
transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
}

```

```
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
}
```



```

bootstrap-import [ policy-names ];
bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
prefix-list prefix-list-addresses {
    source source-address {
        next-hop next-hop-address;
    }
    wildcard-source {
        next-hop next-hop-address;
    }
}

```

```

    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    tunnel-devices [ mt-fpc/pic/port ];
}
rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
    }
    export [ policy-names ];
    import [ policy-names ];
    metric-out metric;
    neighbor neighbor-name {
        any-sender;
        authentication-key password;
        authentication-type type;
        bfd-liveness-detection {
            ... same statements as at the [edit protocols rip group group-name
                bfd-liveness-detection] hierarchy level ...
        }
        (check-zero | no-check-zero);
        import [ policy-names ];
        message-size number;
        metric-in metric;
        receive (both | none | version-1 | version-2);
        route-timeout seconds;
        send (broadcast | multicast | none | version-1);
        update-interval seconds;
    }
    preference preference;
    route-timeout seconds;
    update-interval seconds;
}

```

```

holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number> <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
    }
}

```

```

    priority priority;
  }
  max-age seconds;
}
traceoptions {
  file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
  flag flag;
}
uplink-failure-detection {
  group group-name {
    link-to-monitor interface-name;
    link-to-disable interface-name;
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number > <size size > <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure protocols.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Junos OS Routing Protocols Configuration Guide](#)

proxy-arp

Syntax	proxy-arp (restricted unrestricted);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. restricted added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series..
Description	For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.
Default	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
Options	<ul style="list-style-type: none"> • none—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. • restricted—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address. • unrestricted—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
Default: unrestricted	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Restricted and Unrestricted Proxy ARP • Configuring Proxy ARP (CLI Procedure) • Example: Configuring Proxy ARP on an EX Series Switch • Configuring Gratuitous ARP

reflective-relay

Syntax	reflective-relay;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.
Default	Switch interfaces are not configured for reflective relay.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958• Configuring Reflective Relay on page 2044

revision-level

Syntax	revision-level <i>revision-level</i> ;
Hierarchy Level	[edit protocols mstp]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.
Default	The revision number is disabled.
Options	<i>revision-level</i> —Revision number of the MSTP region configuration. Range: 0 through 65535
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Network Regions for VLANs with MSTP on page 1931• Understanding MSTP on page 1901• show spanning-tree bridge on page 2235• show spanning-tree interface on page 2240

rstp

Syntax	<pre> rstp { disable; bpdu-block-on-edge; bridge-priority <i>priority</i>; forward-delay <i>seconds</i>; hello-time <i>seconds</i>; interface (all <i>interface-name</i>) { bpdu-timeout-action { block; alarm; } disable; cost <i>cost</i>; edge; mode <i>mode</i>; no-root-port; priority <i>priority</i>; } max-age <i>seconds</i>; traceoptions { file <i>name</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, providing shorter convergence times than those provided with basic STP.</p> <p>The statements are explained separately.</p>
Default	RSTP is enabled on all Ethernet switching interfaces.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918 • Understanding RSTP on page 1901 • show spanning-tree bridge on page 2235 • show spanning-tree interface on page 2240

storm-control

Syntax storm-control {
 action-shutdown;
 interface (all | *interface-name*) {
 bandwidth *bandwidth*;
 no-broadcast;
 no-multicast;
 no-unknown-unicast;
 }
 }

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Apply storm control to all interfaces or to the specified interfaces.

The statements are explained separately.



NOTE: The **no-multicast** option is not supported on QFabric systems.

Default By default, storm control is enabled on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. You can change the storm control level by configuring it as a specific bandwidth value.

When you configure storm control bandwidth on an aggregated Ethernet interface, each member of the aggregated interface is assigned that bandwidth. For example, if you configure 7000000 Kbps on aggregated interface **ae1**, and **ae1** has two members, **xe-2:0/0/0** and **xe-2:0/0/1**, each member is allowed a bandwidth level of 7000000 Kbps. Thus, the storm control bandwidth on **ae1** could be as much as 14000000 Kbps of combined broadcast and unknown unicast traffic.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding Storm Control on page 4694](#)
- [Example: Configuring Storm Control on page 4713](#)
- [port-error-disable on page 4823](#)
- [disable-timeout on page 4809](#)
- [clear ethernet-switching port-error on page 4859](#)

stp

Syntax	<pre> stp { disable; bridge-priority <i>priority</i>; forward-delay <i>seconds</i>; hello-time <i>seconds</i>; interface (all <i>interface-name</i>) { disable; bpdu-timeout-action { block; alarm; } cost <i>cost</i>; edge; mode <i>mode</i>; no-root-port; priority <i>priority</i>; } max-age <i>seconds</i>; traceoptions { file <i>name</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>When you explicitly configure STP, a switch uses the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP (defined in the IEEE 802.1D 1998 specification).</p> <p>The remaining statements are explained separately.</p>
Default	STP is disabled; by default, RSTP is enabled on all Ethernet switching ports.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1990 • Configuring STP on page 2053 • Overview of Spanning-Tree Protocols on page 1900 • show spanning-tree bridge on page 2235 • show spanning-tree interface on page 2240

traceoptions (Ethernet Switching Options)

Syntax `traceoptions {
 file filename <files number> <no-stamp> <replace> <size size> <world-readable |
 no-world-readable>;
 flag flag <disable>;
 }`

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.



NOTE: The `traceoptions` statement is not supported on the QFX3000 QFabric system.

Description Define global tracing operations for access security features on Ethernet switches.

Default The `traceoptions` feature is disabled by default.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename* —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

files *number* —(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (**xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag* —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **access-security**—Trace access security events.
- **all**—All tracing operations.
- **analyzer**—Trace analyzer events.
- **config-internal**—Trace internal configuration operations.
- **filter**—Trace filter transaction events.
- **forwarding-database**—Trace forwarding database events.

- **general**—Trace general events.
- **interface**—Trace interface events.
- **krt**—Trace communications over routing sockets.
- **lib**—Trace library calls.
- **nexthop**—Trace next-hop events.
- **normal**—Trace normal events.
- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **rtg**—Trace redundant trunk group events.
- **state**—Trace state transitions.
- **stp**—Trace spanning-tree events.
- **task**—Trace Ethernet-switching task processing.
- **timer**—Trace Ethernet-switching timer processing.
- **unknown-unicast-forwarding**—Trace unknown unicast forwarding events.
- **vlan**—Trace VLAN events.

no-stamp—(Optional) Do not timestamp the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending to it.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes

Range: 10 KB through 1 gigabyte

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

- Related Documentation**
- [Overview of Spanning-Tree Protocols on page 1900](#)
 - [Understanding Bridging on page 1869](#)

traceoptions (STP)

Syntax	<pre> traceoptions { file <i>name</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } </pre>
Hierarchy Level	[edit protocols mstp], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan <i>vlan-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.



NOTE: traceoptions is not support on QFabric systems.

Description	Set STP protocol-level tracing options.
Default	Traceoptions is disabled.
Options	<p>disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file <code>/var/log/stp-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 1 trace file only</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the STP-specific tracing options:</p> <ul style="list-style-type: none"> • all—Trace all operations. • all-failures—Trace all failure conditions. • bpdu—Trace BPDU reception and transmission.

- **bridge-detection-state-machine**—Trace the bridge detection state machine.
- **events**—Trace events of the protocol state machine.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **ppmd**—Trace the state and events for the ppm process.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.O**. When the **trace-file** again reaches its maximum size, **trace-file.O** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.O**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
 - [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1918](#)
 - [Understanding RSTP on page 1901](#)
 - [Understanding MSTP on page 1901](#)
 - [Overview of Spanning-Tree Protocols on page 1900](#)
 - [Understanding VSTP on page 1902](#)
 - [show spanning-tree bridge on page 2235](#)
 - [show spanning-tree interface on page 2240](#)

unknown-unicast-forwarding

Syntax unknown-unicast-forwarding {
 vlan (all | *vlan-name*){
 interface *interface-name*;
 }
 }

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.



NOTE: Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

The remaining statements are explained separately.

Default Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Unknown Unicast Forwarding](#)
 - [Understanding Unknown Unicast Forwarding](#)
 - [show ethernet-switching table on page 2211](#)
 - [show vlans on page 2251](#)

vlan (Ethernet)

Syntax	<pre> vlan { members [(all names vlan-ids)]; } </pre>
Hierarchy Level	[edit interfaces ge-chassis/slot/port unit logical-unit-number ethernet-switching], [edit interfaces xe-chassis/slot/port unit logical-unit-number ethernet-switching]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>For Gigabit Ethernet and aggregated Ethernet interfaces, assign an 802.1Q VLAN tag ID to a logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Bridging with Multiple VLANs on page 1971 • Junos OS Network Interfaces Configuration Guide

vlan (STP)

```
Syntax  vlan (vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
            bpdu-timeout-action {
                block;
                alarm;
            }
            cost cost;
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
        }
        max-age seconds;
        traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
            flag flag;
        }
    }
```

Hierarchy Level [edit protocols **mstp** **msti** *msti-id*],
[edit protocols **vstp**]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the VLANs for a Multiple Spanning Tree Instance (MSTI).

The remaining statements are explained separately.



TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Default Not enabled.

Options *vlan-id*—Numeric VLAN identifier.

vlan-name—Name of the VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Network Regions for VLANs with MSTP on page 1931](#)
 - [Understanding MSTP on page 1901](#)
 - [Understanding VSTP on page 1902](#)

vlan (Unknown Unicast)

Syntax `vlan (all | vlan-name) {
 interface interface-name;
}`

Hierarchy Level [edit [ethernet-switching-options unknown-unicast-forwarding](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Specify a VLAN from which unknown unicast packets will be forwarded, or specify that the packets should be forwarded from *all* VLANs. Unknown unicast packets are forwarded from a VLAN to a specific trunk interface.

The remaining statement is explained separately.



TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlangs` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options `all`—All VLANs.

`vlan-name`—Name of a VLAN.

Required Privilege Level `routing`—To view this statement in the configuration.
`routing-control`—To add this statement to the configuration.

- Related Documentation**
- [Configuring Unknown Unicast Forwarding](#)
 - [Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface](#)
 - [Understanding Unknown Unicast Forwarding](#)
 - [show ethernet-switching table on page 2211](#)
 - [show vlans on page 2251](#)

vlan-id (VLANs)

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.
Default	If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.
Options	<i>number</i> —VLAN tag identifier. Range: 0 through 4093.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Bridging with Multiple VLANs on page 1971• Understanding Bridging on page 1869

vlan-range

Syntax	<code>vlan-range <i>vlan-id-low-vlan-id-high</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.
Default	None.
Options	<i>vlan-id-low-vlan-id-high</i> —Specify the first and last VLAN ID number for the group of VLANs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VLANs on page 2048• Configuring Routed VLAN Interfaces on page 2046• Understanding Bridging on page 1869

vlan

Syntax	<pre> vlan { vlan-name { description text-description; dot1q-tunneling { customer-vlans (id range); } filter input filter-name; filter output filter-name; interface interface-name { isolated; mapping (policy tag push native push); promiscuous; } isolation-vlan-id; l3-interface vlan.logical-interface-number; mac-limit number; mac-table-aging-time seconds; no-local-switching; no-mac-learning; primary-vlan vlan-name; pvlan extend-secondary-vlan-id vlan-id; vlan-id number; vlan-range vlan-id-low-vlan-id-high; } } </pre>
Hierarchy Level	[edit]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure VLAN properties on the QFX Series.
Default	If you use the default factory configuration, all switch interfaces become part of the VLAN default.
Options	<p>vlan-name—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VLANs on page 2048 • Configuring Q-in-Q Tunneling on page 2062 • Creating a Series of Tagged VLANs on page 2050

- [Configuring Routed VLAN Interfaces on page 2046](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Understanding Bridging on page 1869](#)

[vlan-tagging](#)

Syntax	<code>vlan-tagging;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>] [edit interfaces interface-range <i>interface-range-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.
Default	VLAN tagging is disabled by default.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• vlan-id on page 2624• Configuring a Layer 3 Logical Interface on page 2537

vstp

```

Syntax  vstp {
        disable;
        bpdu-block-on-edge;
        force-version stp;
        vlan (vlan-id | vlan-name) {
            bridge-priority priority;
            forward-delay seconds;
            hello-time seconds;
            interface (all | interface-name) {
                disable;
                bpdu-timeout-action {
                    alarm;
                    block;
                }
                cost cost;
                edge;
                mode mode;
                no-root-port;
                priority priority;
            }
            max-age seconds;
            traceoptions {
                file name <replace> <size size> <files number> <no-stamp>
                  <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
    }

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding VSTP on page 1902](#)
- [show spanning-tree bridge on page 2235](#)
- [show spanning-tree interface on page 2240](#)

MVRP Configuration Statements

- [disable \(MVRP\) on page 2148](#)
- [interface \(MVRP\) on page 2149](#)
- [join-timer \(MVRP\) on page 2150](#)
- [leave-timer \(MVRP\) on page 2151](#)
- [leaveall-timer \(MVRP\) on page 2152](#)
- [passive \(MVRP\) on page 2153](#)

disable (MVRP)

Syntax	disable;
Hierarchy Level	[edit protocols mvrp], [edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Disable the MVRP configuration on the interface.
Default	MVRP is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)</i>

interface (MVRP)

Syntax	<pre>interface (all <i>interface-name</i>) { disable; join-timer <i>milliseconds</i>; leave-timer <i>milliseconds</i>; leaveall-timer <i>milliseconds</i>; registration (forbidden normal); }</pre>
Hierarchy Level	[edit protocols mvrp]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).



NOTE: On QFX Series switches, you must configure specific interfaces—you cannot specify interface all. You can enable MVRP on an interface range.

Default	By default, MVRP is disabled.
Options	<p>all—All interfaces on the switch.</p> <p><i>interface-name</i>—Names of interface to be configured for MVRP.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i> • <i>Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)</i>

join-timer (MVRP)

Syntax	join-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	<p>Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs).</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	200 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds that the interface must wait before sending MVRP PDUs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• leave-timer on page 2151• leaveall-timer on page 2152• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i>• <i>Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)</i>

leave-timer (MVRP)

Syntax	<code>leave-timer <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit protocols mvrp interface (all <i>interface-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	<p>For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	1000 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • join-timer on page 2150 • leaveall-timer on page 2152 • <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i> • <i>Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)</i>

leaveall-timer (MVRP)

Syntax	<code>leaveall-timer <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	<p>For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	10000 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds between the sending of Leave All messages.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• join-timer on page 2150• leave-timer on page 2151• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i>• <i>Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)</i>

passive (MVRP)

Syntax	<code>passive;</code>
Hierarchy Level	<code>[edit protocols mvrp],</code> <code>[edit protocols mvrp interface(all <i>interface-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 13.1 for the QFX Series.
Description	Configure an MVRP interface to not announce its membership in a VLAN or send any VLAN declarations (updates) unless it receives registration for that VLAN from a peer (server).
Default	Passive mode is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP on page 1953

Private VLAN Configuration Statements

- [extend-secondary-vlan-id on page 2154](#)
- [isolated on page 2154](#)
- [isolation-vlan-id on page 2155](#)
- [primary-vlan on page 2155](#)
- [pvlan on page 2156](#)
- [promiscuous on page 2156](#)
- [pvlan-trunk on page 2157](#)
- [vlans on page 2158](#)

extend-secondary-vlan-id

Syntax	<code>extend-secondary-vlan-id <i>number</i>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> pvlan]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	Configure traffic that egresses from a secondary VLAN trunk port to retain its secondary VLAN tag instead of getting the tag of the primary VLAN that the secondary port is a member of.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887• Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024

isolated

Syntax	<code>isolated;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an access or trunk port to be isolated. You configure a trunk port to be isolated so that it can be a secondary VLAN trunk port—that is, it can carry secondary VLAN traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating a Private VLAN on a Single Switch on page 2057• Creating a Private VLAN Spanning Multiple Switches on page 2058• Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887• Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024

isolation-vlan-id

Syntax	<code>isolation-vlan-id <i>number</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i> pvlan]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an interswitch isolated VLAN within a private VLAN that spans multiple switches.
Options	<i>number</i> —VLAN tag identifier. Range: 0 through 4093
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Creating a Private VLAN on a Single Switch on page 2057 • Creating a Private VLAN Spanning Multiple Switches on page 2058

primary-vlan

Syntax	<code>primary-vlan <i>vlan-name</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the primary VLAN for this community VLAN. The primary VLAN must be tagged, and the community VLAN must be untagged.</p> <p>If you want to create a community VLAN, you must configure the primary VLAN to be private using the pvlan statement.</p>
	<div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; margin-right: 10px;">  </div> <div> <p>TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.</p> </div> </div>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Creating a Private VLAN on a Single Switch on page 2057 • Creating a Private VLAN Spanning Multiple Switches on page 2058

pvlan

Syntax	pvlan;
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify that the VLAN is private and access ports in the VLAN do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.
Options	none
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating a Private VLAN on a Single Switch on page 2057• Creating a Private VLAN Spanning Multiple Switches on page 2058

promiscuous

Syntax	promiscuous;
Hierarchy Level	[edit vlan <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an access or trunk port to be promiscuous.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating a Private VLAN on a Single Switch on page 2057• Creating a Private VLAN Spanning Multiple Switches on page 2058

pvlan-trunk

Syntax	pvlan-trunk;
Hierarchy Level	[edit vlans <i>vlan-name</i> interface]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an interface to be a private VLAN trunk port.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating a Private VLAN on a Single Switch on page 2057• Creating a Private VLAN Spanning Multiple Switches on page 2058

vlan

Syntax	<pre> vlan { vlan-name { description text-description; dot1q-tunneling { customer-vlans (id range); } filter input filter-name; filter output filter-name; interface interface-name { isolated; mapping (policy tag push native push); promiscuous; } isolation-vlan-id; l3-interface vlan.logical-interface-number; mac-limit number; mac-table-aging-time seconds; no-local-switching; no-mac-learning; primary-vlan vlan-name; pvlan extend-secondary-vlan-id vlan-id; vlan-id number; vlan-range vlan-id-low-vlan-id-high; } } </pre>
Hierarchy Level	[edit]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure VLAN properties on the QFX Series.
Default	If you use the default factory configuration, all switch interfaces become part of the VLAN default.
Options	<p>vlan-name—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VLANs on page 2048 • Configuring Q-in-Q Tunneling on page 2062 • Creating a Series of Tagged VLANs on page 2050

- [Configuring Routed VLAN Interfaces on page 2046](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Understanding Bridging on page 1869](#)

Q-in-Q Tunneling Configuration Statements

- [customer-vlans on page 2160](#)
- [dot1q-tunneling \(Ethernet Switching\) on page 2161](#)
- [dot1q-tunneling \(VLANs\) on page 2162](#)
- [ether-type on page 2163](#)
- [layer2-protocol-tunneling on page 2164](#)
- [mapping on page 2166](#)
- [mapping-range on page 2167](#)
- [no-local-switching on page 2167](#)
- [shutdown-threshold on page 2168](#)
- [vlan-id-start on page 2169](#)
- [vlans on page 2170](#)

customer-vlans

Syntax	customer-vlans (<i>id</i> native <i>range</i>);
Hierarchy Level	[edit vlans <i>vlan-name</i> dot1q-tunneling]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Option native introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Limit the set of accepted customer VLAN tags to a range or to discrete values when mapping customer VLANs to service VLANs.
Options	<p>id—Numeric identifier for a VLAN.</p> <p>native—Accepts untagged and priority-tagged packets from access interfaces and assigns the configured S-VLAN to the packet.</p> <p>range—Range of numeric identifiers for VLANs. On the QFX series, you can include as many as eight separate customer VLAN ranges for a given service VLAN. Do not configure more than this number of ranges.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• dot1q-tunneling (Ethernet Switching) on page 2161• <i>ether-type</i>• <i>Example: Setting Up Q-in-Q Tunneling on EX Series Switches</i>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i>• <i>Understanding Q-in-Q Tunneling on EX Series Switches</i>• Configuring Q-in-Q Tunneling on page 2062• <i>Example: Setting Up Q-in-Q Tunneling on page 2036</i>• dot1q-tunneling (Ethernet Switching) on page 2161• ether-type on page 2163

dot1q-tunneling (Ethernet Switching)

Syntax	dot1q-tunneling { ether-type (0x8100 0x88a8 0x9100); }
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set a global value for the EtherType for Q-in-Q tunneling. The remaining statement is explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • dot1q-tunneling on page 2162 • <i>Example: Setting Up Q-in-Q Tunneling on EX Series Switches</i> • <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i> • Configuring Q-in-Q Tunneling on page 2062 • <i>Example: Setting Up Q-in-Q Tunneling on page 2036</i> • dot1q-tunneling on page 2162

dot1q-tunneling (VLANs)

Syntax	<pre>dot1q-tunneling { customer-vlans (id native range); layer2-protocol-tunneling all protocol-name { drop-threshold number; shutdown-threshold number; } }</pre>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Option native introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Options layer2-protocol-tunneling, drop-threshold, and shutdown-threshold introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Enable Q-in-Q tunneling on the specified VLAN.

**NOTE:**

- The VLAN on which you enable Q-in-Q tunneling must be a tagged VLAN.
 - You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.
-

The remaining statements are explained separately.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Q-in-Q Tunneling on EX Series Switches• Configuring Q-in-Q Tunneling (CLI Procedure)• Configuring Q-in-Q Tunneling on page 2062• Example: Setting Up Q-in-Q Tunneling on page 2036• Configuring Layer 2 Protocol Tunneling on page 2063• dot1q-tunneling (Ethernet Switching) on page 2161

ether-type

Syntax	ether-type (0x8100 0x88a8 0x9100)
Hierarchy Level	[edit ethernet-switching-options dot1q-tunneling]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure a global value for the tag protocol identifier (EtherType) of the service VLAN tags (outer tags) in Q-in-Q tunneling. Only one EtherType value is supported at a time.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Q-in-Q Tunneling on page 2062• Example: Setting Up Q-in-Q Tunneling on page 2036

layer2-protocol-tunneling

Syntax layer2-protocol-tunneling all | *protocol-name* {
 drop-threshold number;
 shutdown-threshold number;
 }

Hierarchy Level [edit vlans *vlan-name* dot1q-tunneling]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.
 Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable Layer 2 protocol tunneling (L2PT) on the VLAN.

 The remaining statements are explained separately.

Default L2PT is not enabled.

Options all—Enable all supported Layer 2 protocols.

protocol-name—Name of the Layer 2 protocol. Values are:

- 802.1x—IEEE 802.1X authentication
- 802.3ah—IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)



NOTE: If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.

- cdp—Cisco Discovery Protocol
- e-lmi—Ethernet local management interface
- gvrp—GARP VLAN Registration Protocol
- lacp—Link Aggregation Control Protocol



NOTE: If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.

- lldp—Link Layer Discovery Protocol
- mmrp—Multiple MAC Registration Protocol
- mvrp—Multiple VLAN Registration Protocol
- stp—Spanning Tree Protocol, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol

- **udld**—Unidirectional Link Detection (UDLD)
- **vstp**—VLAN Spanning Tree Protocol
- **vtp**—VLAN Trunking Protocol

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [show ethernet-switching layer2-protocol-tunneling interface on page 2195](#)
- [show ethernet-switching layer2-protocol-tunneling statistics on page 2197](#)
- [show ethernet-switching layer2-protocol-tunneling vlan on page 2200](#)
- *Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches*
- *Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)*
- [Configuring Layer 2 Protocol Tunneling on page 2063](#)

mapping

Syntax	<code>mapping (native (push swap) tag (push swap));</code> <code>mapping native inner-tag tag push;</code> <code>mapping native push inner-tag tag;</code>
Hierarchy Level	[edit vlangs <i>vlan-name</i> interface <i>interface-name</i> egress], [edit vlangs <i>vlan-name</i> interface <i>interface-name</i> ingress], [edit vlangs <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Map a specific C-VLAN to an S-VLAN. By default, the received incoming or outgoing tag is replaced with the new tag.
Options	<p>inner-tag (QFabric systems only)—apply the specified tag as an inner tag to packets that are received as untagged on an access interface.</p> <p>native—Map untagged and priority-tagged packets to an S-VLAN.</p> <p>push—Retain the incoming tag (as an inner tag) and adds an additional VLAN tag. When you use this option, the TPID of the outer tag is set as follows:</p> <ul style="list-style-type: none">• If Q-in-Q tunneling is not enabled in the VLAN, then the Ethertype for outer tag is set to 0x8100.• If Q-in-Q tunneling is enabled in the VLAN and a packet is egressing from a trunk port, then the Ethertype is set to 0x88a8 (or as configured by an ether-type statement). <p>swap—Replaces the incoming VLAN tag with the VLAN ID tag of the S-VLAN. Using this option is also referred to as VLAN ID translation. When you use this option on a trunk port for which Q-in-Q tunneling is enabled, use the ether-type statement to set the Ethertype.</p> <p>tag—Original VLAN tag that will be replaced (with swap) or that will become an inner tag (with push).</p>
Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Q-in-Q Tunneling on page 2062• Example: Setting Up Q-in-Q Tunneling on page 2036

mapping-range

Syntax	<code>mapping-range C-VLAN-range (push swap) <vlan-id-start S-VLAN-ID>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> interface (VLANs) <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an access interface to map a range of C-VLANs to a range of S-VLANs. Use this statement instead of using multiple set vlans <i>VLAN-name</i> interface <i>interface-name</i> mapping (push swap) statements to configure Q-in-Q tunneling or VLAN translation on a per-VLAN basis. This statement is particularly useful if you have used the vlan-range statement to create multiple VLANs.
Options	<p>push—Retain the incoming tag and adds an additional VLAN tag (Q-in-Q tunneling).</p> <p>swap—Swap the incoming VLAN tag with the VLAN ID tag of the S-VLAN (VLAN translation).</p> <p>vlan-ID-start <i>S-VLAN-ID</i>—(Optional) Set the start of the S-VLAN range that the C-VLANs will be mapped to. If you omit this option, mapping begins with the first ID in the range of S-VLAN IDs (which you configure using the set vlans <i>vlan-range</i> statement).</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Q-in-Q Tunneling on page 2062 • Example: Setting Up Q-in-Q Tunneling on page 2036 • vlan-range on page 2144

no-local-switching

Syntax	<code>no-local-switching;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify that access ports in this VLAN domain do not forward packets to one another. You use this statement with primary VLANs and isolated secondary VLANs.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating a Private VLAN on a Single Switch on page 2057 • Creating a Private VLAN Spanning Multiple Switches on page 2058

shutdown-threshold

Syntax	<code>shutdown-threshold <i>number</i>;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i> dot1q-tunneling layer2-protocol-tunneling (all <i>protocol-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled. Once an interface is disabled, you must explicitly reenable it using the clear ethernet-switching layer2-protocol-tunneling error command. Otherwise, the interface remains disabled.</p> <p>The shutdown threshold value must be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value, the drop threshold value has no effect.</p> <p>You can specify a shutdown threshold value without specifying a drop threshold value.</p>
Default	No shutdown threshold is specified.
Options	<p><i>number</i>—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled.</p> <p>Range: 1 through 1000</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• drop-threshold on page 2084• <i>Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</i>• <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i>• Configuring Layer 2 Protocol Tunneling on page 2063

vlan-id-start

Syntax	<code>vlan-id-start S-VLAN-ID;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> interface (VLANs) <i>interface-name</i> mapping-range <i>C-VLAN-range</i> (push swap)]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an access interface to map a range of C-VLANs to a range of S-VLANs. Use this statement instead of using multiple set vlans <i>VLAN-name</i> interface <i>interface-name</i> mapping (push swap) statements to configure Q-in-Q tunneling or VLAN translation on a per-VLAN basis. This statement sets the start of the S-VLAN range that the C-VLANs are mapped to. If you omit this option, mapping begins with the first ID in the range of S-VLAN IDs (which you configure using the set vlans <i>vlan-range</i> statement).
Options	None
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Q-in-Q Tunneling on page 2062 • Example: Setting Up Q-in-Q Tunneling on page 2036

vlan

Syntax	<pre>vlan { vlan-name { description text-description; dot1q-tunneling { customer-vlans (id range); } filter input filter-name; filter output filter-name; interface interface-name { isolated; mapping (policy tag push native push); promiscuous; } isolation-vlan-id; l3-interface vlan.logical-interface-number; mac-limit number; mac-table-aging-time seconds; no-local-switching; no-mac-learning; primary-vlan vlan-name; pvlan extend-secondary-vlan-id vlan-id; vlan-id number; vlan-range vlan-id-low-vlan-id-high; } }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure VLAN properties on the QFX Series.
Default	If you use the default factory configuration, all switch interfaces become part of the VLAN default.
Options	vlan-name —Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long. The remaining statements are described separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VLANs on page 2048• Configuring Q-in-Q Tunneling on page 2062• Creating a Series of Tagged VLANs on page 2050

- [Configuring Routed VLAN Interfaces on page 2046](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Understanding Bridging on page 1869](#)

CHAPTER 22

Administration

- [Routine Monitoring on page 2173](#)
- [Monitoring Commands on page 2183](#)

Routine Monitoring

- [Verifying That MAC Notification Is Working Properly on page 2173](#)
- [Verifying That a Series of Tagged VLANs Has Been Created on page 2173](#)
- [Verifying That Q-in-Q Tunneling Is Working on page 2175](#)
- [Verifying That a Private VLAN Is Working on page 2176](#)
- [Verifying That Proxy ARP Is Working Correctly on page 2181](#)
- [Verifying That MVRP Is Working Correctly on page 2182](#)

Verifying That MAC Notification Is Working Properly

Purpose Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

Action To verify that MAC notification is enabled or disabled and also to verify the MAC notification interval setting.

```
user@switch> show ethernet-switching mac-notification
Notification Status: Enabled
Notification Interval: 30
```

Meaning The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 30 seconds.

Related Documentation • [Configuring MAC Notification on page 2042](#)

Verifying That a Series of Tagged VLANs Has Been Created

Purpose Verify that a series of tagged VLANs has been created on the switch.

- Action** 1. Display the VLANs in the ascending order of their VLAN ID:

```
user@switch> show vlans sort-by tag
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

2. Display the VLANs by the alphabetical order of the VLAN name:

```
user@switch> show vlans sort-by name
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

3. Display the VLANs by specifying the VLAN range name (here, the VLAN range name is **employee**):

```
user@switch> show vlans employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

Meaning The sample output shows the VLANs configured on the switch. The series of tagged VLANs is displayed: **__employee_120__** through **__employee_130__**. Each of the tagged VLANs is configured on the trunk interface **xe-0/0/22.0**. The asterisk (*) next to the interface name indicates that the interface is **UP**.

When a series of VLANs is created using the **vlan-range** statement, the VLAN names are preceded and followed by a double underscore.

Related Documentation

- [Creating a Series of Tagged VLANs on page 2050](#)

Verifying That Q-in-Q Tunneling Is Working

Purpose After creating a Q-in-Q VLAN, verify that it is set up properly.

Action 1. Use the **show configuration vlans** command to determine if you successfully created the primary and secondary VLAN configurations:

```
user@switch> show configuration vlans
svlan {
  vlan-id 300;
  dot1q-tunneling {
    customer-vlans [ 101-200 ];
  }
}
```

2. Use the **show vlans** command to view VLAN information and link status:

```
user@switch> show vlans s-vlan-name extensive
VLAN: svlan, Created at: Thu Oct 23 16:53:20 2008
802.1Q Tag: 300, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
```

```
Customer VLAN ranges:
                        101-200
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 0), Untagged 1 (Active = 0)
                    xe-0/0/1, tagged, trunk
                    xe-0/0/2, untagged, access
```

Meaning The output confirms that Q-in-Q tunneling is enabled and that the VLAN is tagged, and lists the customer VLANs that are associated with the tagged VLAN.

Related Documentation

- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*

Verifying That a Private VLAN Is Working

Purpose After creating and configuring private VLANs (PVLANS), verify that they are set up properly.

Action 1. To determine whether you successfully created the primary and secondary VLAN configurations:

- For a PVLAN on a single switch, use the **show configuration vlans** command:

```
user@switch> show configuration vlans
community1 {
    interface {
        interface a;
        interface b;
    }
    primary-vlan pvlan;
}
community2 {
    interface {
        interface d;
        interface e;
    }
    primary-vlan pvlan;
}
pvlan {
    vlan-id 1000;
    interface {
        isolated1;
        isolated2;
        trunk1;
        trunk2;
    }
    no-local-switching;
}
```

- For a PVLAN spanning multiple switches, use the **show vlans extensive** command:

```
user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
                    ge-0/0/20.0*, tagged, trunk
```

```

ge-0/0/22.0*, tagged, trunk, pvlan-trunk
ge-0/0/23.0*, tagged, trunk, pvlan-trunk
ge-0/0/7.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/2.0, untagged, access

```

```

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk

```

```

VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, untagged, access
    ge-1/0/6.0*, untagged, access

```

```

VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access
    ge-0/0/1.0*, untagged, access
    ge-0/0/2.0, untagged, access
    ge-0/0/7.0*, untagged, access
    ge-1/0/6.0*, untagged, access

```

```
Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
  __pvlan_primary_ge-0/0/0.0__
  __pvlan_primary_ge-0/0/2.0__
Community VLANs :
  COM1
  community2
Inter-switch-isolated VLAN :
  __pvlan_primary_isiv__
```

2. Use the **show vlans extensive** command to view VLAN information and link status for a PVLAN on a single switch or for a PVLAN spanning multiple switches.

- For a PVLAN on a single switch:

```
user@switch> show vlans pvlan extensive
VLAN: pvlan, Created at: time
802.1Q Tag: vlan-id, Internal index: index-number, Admin State: Enabled,
Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 6 (Active = 0)
  trunk1, tagged, trunk
  interface a, untagged, access
  interface b, untagged, access
  interface c, untagged, access
  interface d, untagged, access
  interface e, untagged, access
  interface f, untagged, access
  trunk2, tagged, trunk
Secondary VLANs: Isolated 2, Community 2
Isolated VLANs :
  __pvlan_pvlan_isolated1__
  __pvlan_pvlan_isolated2__
Community VLANs :
  community1
  community2
```

- For a PVLAN spanning multiple switches:

```
user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
  ge-0/0/20.0*, tagged, trunk
  ge-0/0/22.0*, tagged, trunk, pvlan-trunk
  ge-0/0/23.0*, tagged, trunk, pvlan-trunk
  ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
  ge-0/0/20.0*, tagged, trunk
  ge-0/0/22.0*, tagged, trunk, pvlan-trunk
  ge-0/0/23.0*, tagged, trunk, pvlan-trunk
  ge-0/0/0.0*, untagged, access
```

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
 Internal index: 6, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Isolated, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/2.0, untagged, access

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk

VLAN: community2, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Community, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, untagged, access
 ge-1/0/6.0*, untagged, access

VLAN: primary, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/0.0*, untagged, access
 ge-0/0/1.0*, untagged, access
 ge-0/0/2.0, untagged, access
 ge-0/0/7.0*, untagged, access
 ge-1/0/6.0*, untagged, access

Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
 Isolated VLANs :
 __pvlan_primary_ge-0/0/0.0__
 __pvlan_primary_ge-0/0/2.0__
 Community VLANs :
 COM1
 community2
 Inter-switch-isolated VLAN :
 __pvlan_primary_isiv__

3. Use the **show ethernet-switching table** command to view logs for MAC learning on the VLANs:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 8 entries, 1 learned
```

VLAN	MAC address	Type	Age	Interfaces
default	*	Flood	-	All-members
pvlan	*	Flood	-	All-members
pvlan	MAC1	Replicated	-	interface a
pvlan	MAC2	Replicated	-	interface c
pvlan	MAC3	Replicated	-	isolated2
pvlan	MAC4	Learn	0	trunk1
__pvlan_pvlan_isolated1__	*	Flood	-	All-members
__pvlan_pvlan_isolated1__	MAC4	Replicated	-	trunk1
__pvlan_pvlan_isolated2__	*	Flood	-	All-members
__pvlan_pvlan_isolated2__	MAC3	Learn	0	isolated2
__pvlan_pvlan_isolated2__	MAC4	Replicated	-	trunk1
community1	*	Flood	-	All-members
community1	MAC1	Learn	0	interface a
community1	MAC4	Replicated	-	trunk1
community2	*	Flood	-	All-members
community2	MAC2	Learn	0	interface c
community2	MAC4	Replicated	-	trunk1



NOTE: If you have configured a PVLAN spanning multiple switches, you can use the same command on all the switches to check the logs for MAC learning on those switches.

Meaning In the output displays for a PVLAN on a single switch, you can see that the primary VLAN contains two community domains (**community1** and **community2**), two isolated ports, and two trunk ports. The PVLAN on a single switch has only one tag (1000), which is for the primary VLAN.

The PVLAN that spans multiple switches contains multiple tags:

- The community domain **COM1** is identified with tag **100**.
- The community domain **community2** is identified with tag **20**.
- The interswitch isolated domain is identified with tag **50**.
- The primary VLAN **primary** is identified with tag **10**.

Also, for the PVLAN that spans multiple switches, the trunk interfaces are identified as **pvlan-trunk**.

Related Documentation

- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\)](#)
- [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\)](#)
- [Creating a Private VLAN on a Single Switch on page 2057](#)
- [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)

Verifying That Proxy ARP Is Working Correctly

Purpose Verify that the switch is sending proxy ARP messages.

Action List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
198319 datagrams received
45 ARP requests received
12 ARP replies received
2 resolution requests received
2 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted-proxy requests not proxied
0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
168705 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29555 which were not for me
0 packets discarded waiting for resolution
4 packets sent after waiting for resolution
27 ARP requests sent
47 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
```

```

0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

Meaning The statistics show that two proxy ARP requests were received, and the **proxy requests not proxied** field indicates that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- *Configuring Proxy ARP (CLI Procedure)*

Verifying That MVRP Is Working Correctly

Purpose After configuring your switch to participate in MVRP, verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

Action 1. Confirm that MVRP is enabled on your switch.

```
user@switch> show mvrp
```

Global MVRP configuration

```

MVRP status           : Enabled
MVRP dynamic vlan creation: Enabled
MVRP Timers (ms):
Interface      Join   Leave  LeaveAll
-----
all            200   600    10000
xe-0/1/1.0     200   600    10000

```

Interface based configuration:

Interface	Status	Registration	Dynamic VLAN Creation
all	Disabled	Fixed	Enabled
xe-0/1/1.0	Enabled	Normal	Enabled

2. Confirm that MVRP messages are being sent and received on your switch.

```
user@switch> show mvrp statistics interface xe-0/1/1.0
```

```

MVRP statistics
MRPDU received           : 3342
Invalid PDU received     : 0
New received             : 2
Join Empty received      : 1116
Join In received         : 2219
Empty received           : 2
In received              : 2
Leave received            : 1
LeaveAll received         : 1117
MRPDU transmitted       : 3280
MRPDU transmit failures : 0
New transmitted          : 0
Join Empty transmitted   : 1114
Join In transmitted      : 2163
Empty transmitted        : 1
In transmitted           : 1
Leave transmitted         : 1
LeaveAll transmitted      : 1111

```

Meaning The output of `show mvrp` shows that interface `xe-0/1/1.0` is enabled for MVRP participation as shown in the status in the **Interface based configuration** field.

The output for `show mvrp statistics interface xe-0/1/1.0` confirms that MVRP messages are being transmitted and received on the interface.



NOTE: You can identify an MVRP compatibility issue on EX Series switches by looking at the output from this command. If *Join Empty received* and *Join In received* incorrectly display zero, even though the value for *MRPDU received* has been increased, you are probably running different versions of Junos OS, including Release 11.3, on the switches in this network. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the earlier release version. To remedy these problems, see *Configuring Multiple VLAN Registration Protocol on EX Series Switches (CLI Procedure)*.

- Related Documentation**
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches](#)
 - [Example: Configuring Automatic VLAN Administration Using MVRP on page 1953](#)
 - [Configuring Multiple VLAN Registration Protocol on EX Series Switches \(CLI Procedure\)](#)
 - [Configuring Multiple VLAN Registration Protocol on page 2051](#)

Monitoring Commands

- `clear ethernet-switching bpdu-error`
- `clear ethernet-switching layer2-protocol-tunneling error`
- `clear ethernet-switching layer2-protocol-tunneling statistics`
- `clear ethernet-switching table`
- `clear spanning-tree statistics`
- `show ethernet-switching interfaces`
- `show ethernet-switching layer2-protocol-tunneling interface`
- `show ethernet-switching layer2-protocol-tunneling statistics`
- `show ethernet-switching layer2-protocol-tunneling vlan`
- `show ethernet-switching mac-learning-log`
- `show ethernet-switching mac-notification`
- `show ethernet-switching statistics aging`
- `show ethernet-switching statistics mac-learning`
- `show ethernet-switching table`
- `show interfaces xe`
- `show spanning-tree bridge`

- [show spanning-tree interface](#)
- [show spanning-tree mstp configuration](#)
- [show spanning-tree statistics](#)
- [show system statistics arp](#)
- [show vlans](#)

clear ethernet-switching bpdu-error

Syntax	<code>clear ethernet-switching bpdu-error interface <i>interface-name</i></code>
Release Information	Command introduced in Junos OS Release 9.1 for EX Series switches. Command updated in Junos OS Release 11.1 for EX Series switches—a BPDU error shuts down the interface and this command brings the interface back up. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear bridge protocol data unit (BPDU) errors from an interface and bring up the interface.
Options	<i>interface-name</i> —Clear BPDU errors on the specified interface.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show spanning-tree interface on page 2240 • Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches • Understanding BPDU Protection for STP, RSTP, and MSTP on page 1903
List of Sample Output	clear ethernet-switching bpdu-error interface on page 2185

Sample Output

clear ethernet-switching bpdu-error interface

```
user@switch> clear ethernet-switching bpdu-error interface xe-0/0/1.0
```

clear ethernet-switching layer2-protocol-tunneling error

Syntax	clear ethernet-switching layer2-protocol-tunneling error <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear Layer 2 protocol tunneling (L2PT) errors on one or more interfaces. If an interface has been disabled because the amount of Layer 2 protocol traffic exceeded the shutdown threshold or because the switch has detected an error in the network topology or configuration, use this command to reenable the interface.
Options	none —Clears L2PT errors on all interfaces. interface <i>interface-name</i> —(Optional) Clear L2PT errors on the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches• Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)• Configuring Layer 2 Protocol Tunneling on page 2063
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling error on page 2186 clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0 on page 2186

Sample Output

clear ethernet-switching layer2-protocol-tunneling error

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error
```

clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0
```

clear ethernet-switching layer2-protocol-tunneling statistics

Syntax	clear ethernet-switching layer2-protocol-tunneling statistics <interface <i>interface-name</i> > <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear Layer 2 protocol tunneling (L2PT) statistics on one or more interfaces or VLANs.
Options	<p>none—Clear L2PT statistics on all interfaces and VLANs.</p> <p>interface <i>interface-name</i>—(Optional) Clear L2PT statistics on the specified interface.</p> <p>vlan <i>vlan-name</i>—(Optional) Clear L2PT statistics on the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching layer2-protocol-tunneling statistics on page 2197 • <i>Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</i> • <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i> • Configuring Layer 2 Protocol Tunneling on page 2063
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling statistics on page 2187 clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 2187 clear ethernet-switching layer2-protocol-tunneling error vlan v2 on page 2187

Sample Output

clear ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics
```


clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics interface xe-0/1/1.0
```

clear ethernet-switching layer2-protocol-tunneling error vlan v2

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

clear ethernet-switching table

Syntax	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <management-vlan> <persistent-mac < <i>interface</i> <i>mac-address</i> >> <vlan <i>vlan-name</i> >
Syntax (QFX Series)	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <persistent-mac < <i>interface</i> <i>mac-address</i> >> <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<div>  <p>NOTE: On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.</p> </div> <p>Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).</p>
Options	<p>none—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.</p> <p>interface <i>interface-name</i>—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.</p> <p>mac <i>mac-address</i>—(Optional) Clear the specified learned MAC address from the Ethernet switching table.</p> <p>management-vlan—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.</p> <p>persistent-mac <<i>interface</i> <i>mac-address</i>>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the interface option to clear all MAC addresses on an interface, or use the mac-address option to clear all entries for a specific MAC address.</p> <p>Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port</p>

will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

vlan *vlan-name*—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

Required Privilege Level

view

Related Documentation

- *show ethernet-switching table*
- [show ethernet-switching table on page 2211](#)
- *Verifying That Persistent MAC Learning Is Working Correctly*

List of Sample Output [clear ethernet-switching table on page 2189](#)


Output Fields This command produces no output.

Sample Output

[clear ethernet-switching table](#)

```
user@switch> clear ethernet-switching table
```

clear spanning-tree statistics

List of Syntax	Syntax on page 2190 Syntax (EX Series Switches and the QFX Series) on page 2190
Syntax	clear spanning-tree statistics <interface <i>interface-name</i> > <logical-system <i>logical-system-name</i> >
Syntax (EX Series Switches and the QFX Series)	clear spanning-tree statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear Spanning Tree Protocol statistics.
Options	none —Reset STP counters for all interfaces for all routing instances. interface <i>interface-name</i> —(Optional) Clear STP statistics for the specified interface only. logical-system <i>logical-system-name</i> —(Optional) Clear STP statistics on a particular logical system.
<div> NOTE: The logical-system option is not available on QFabric systems.</div>	
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">show spanning-tree statistics on page 2248
List of Sample Output	clear stp statistics on page 2190

Sample Output

clear stp statistics

```
user@host> clear stp statistics
```

show ethernet-switching interfaces

Syntax	show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about switched Ethernet interfaces.
Options	<p>none—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display Ethernet-switching information for a specific interface.</p>
Required Privilege Level	view
List of Sample Output	show ethernet-switching interfaces on page 2192 show ethernet-switching interfaces summary on page 2193 show ethernet-switching interfaces brief on page 2193 show ethernet-switching interfaces detail on page 2193 show ethernet-switching interfaces interface-name on page 2194
Output Fields	Table 172 on page 1838 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 203: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are up or down .	none, brief , detail , summary
VLAN members	Name of a VLAN.	none, brief , detail , summary

Table 203: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Blocking	<p>Forwarding state of the interface:</p> <ul style="list-style-type: none"> • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface. • MAC limit exceeded—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • MAC move limit exceeded—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control in effect —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control shutdown in effect —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires. 	none, brief , detail , summary
Index	VLAN index internal to Junos OS software.	detail
untagged tagged	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	detail

Sample Output

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
xe-0/0/0.0	up	T1122	unblocked
xe-0/0/1.0	down	default	– MAC limit exceeded
xe-0/0/2.0	down	default	– MAC move limit exceeded
xe-0/0/3.0	down	default	– Storm control in effect
xe-0/0/4.0	down	default	unblocked
xe-0/0/5.0	down	default	unblocked
xe-0/0/6.0	down	default	unblocked
xe-0/0/7.0	down	default	unblocked
xe-0/0/8.0	down	default	unblocked
xe-0/0/9.0	up	T111	unblocked
xe-0/0/10.0	down	default	unblocked
xe-0/0/11.0	down	default	unblocked
xe-0/0/12.0	down	default	unblocked
xe-0/0/13.0	down	default	unblocked
xe-0/0/14.0	down	default	unblocked
xe-0/0/15.0	down	default	unblocked
xe-0/0/16.0	down	default	unblocked
xe-0/0/17.0	down	default	unblocked
xe-0/0/18.0	down	default	unblocked
xe-0/0/19.0	up	T111	unblocked
xe-0/1/0.0	down	default	unblocked
xe-0/1/1.0	down	default	unblocked
xe-0/1/2.0	down	default	unblocked
xe-0/1/3.0	down	default	unblocked

show ethernet-switching interfaces summary

```

user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0

```

show ethernet-switching interfaces brief

```

user@switch> show ethernet-switching interfaces brief
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down  default       unblocked
xe-0/0/1.0  down  employee-vlan unblocked
xe-0/0/2.0  down  employee-vlan unblocked
xe-0/0/3.0  down  employee-vlan unblocked
xe-0/0/8.0  down  employee-vlan unblocked
xe-0/0/10.0 down  default       unblocked
xe-0/0/11.0 down  employee-vlan unblocked

```

show ethernet-switching interfaces detail

```

user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/1.0 Index: 66
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/2.0 Index: 67
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/3.0 Index: 68
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/8.0 Index: 69
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/10.0 Index: 70
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/11.0 Index: 71
State: down
VLANs:
    employee-vlan          tagged      unblocked

```

show ethernet-switching interfaces interface-name

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down   default       unblocked
```

show ethernet-switching layer2-protocol-tunneling interface

Syntax	<code>show ethernet-switching-layer2-protocol-tunneling interface</code> <code><interface-name></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about Layer 2 protocol tunneling (L2PT) on interfaces that have been configured for L2PT.
Options	none —Display L2PT information about all interfaces on which L2PT is enabled. interface-name —(Optional) Display L2PT information for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching layer2-protocol-tunneling statistics on page 2197 • show ethernet-switching layer2-protocol-tunneling vlan on page 2200 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) • show ethernet-switching layer2-protocol-tunneling statistics on page 2197 • show ethernet-switching layer2-protocol-tunneling vlan on page 2200 • Configuring Layer 2 Protocol Tunneling on page 2063
List of Sample Output	show ethernet-switching layer2-protocol-tunneling interface on page 2196 show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0 on page 2196
Output Fields	Table 204 on page 2195 lists the output fields for the show ethernet-switching layer2-protocol-tunneling interface command. Output fields are listed in the approximate order in which they appear.

Table 204: show ethernet-switching layer2-protocol-tunneling interface Output Fields

Field Name	Field Description
Interface	Name of an interface on the switch.
Operation	Type of operation being performed on the interface. Values are Encapsulation and Decapsulation .
State	State of the interface. Values are active and shutdown .
Description	If the interface state is shutdown , displays why the interface is shut down. If the description says Loop detected , it means that the interface is an access interface that has received L2PT-enabled PDUs. Access interfaces should not receive L2PT-enabled PDUs. This scenario might mean that there is a loop in the network.

Sample Output

show ethernet-switching layer2-protocol-tunneling interface

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface
```

```
Layer2 Protocol Tunneling information:
```

Interface	Operation	State	Description
xe-0/0/0.0	Encapsulation	Shutdown	Shutdown threshold exceeded
xe-0/0/1.0	Decapsulation	Shutdown	Loop detected
xe-0/0/2.0	Decapsulation	Active	

show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0
```

```
Layer2 Protocol Tunneling information:
```

Interface	Operation	State	Description
xe-0/0/0.0	Encapsulation	Shutdown	Shutdown threshold exceeded

show ethernet-switching layer2-protocol-tunneling statistics


Syntax	show ethernet-switching-layer2-protocol-tunneling statistics <interface <i>interface-name</i> > <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.
<div>  NOTE: The show ethernet-switching-layer2-protocol-tunneling statistics command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch. </div>	
Options	<p>none—Display L2PT statistics for all interfaces on which you enabled L2PT.</p> <p>interface <i>interface-name</i>—(Optional) Display L2PT statistics for the specified interface.</p> <p>vlan <i>vlan-name</i>—(Optional) Display L2PT statistics for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ethernet-switching layer2-protocol-tunneling statistics on page 2187 • show ethernet-switching layer2-protocol-tunneling interface on page 2195 • show ethernet-switching layer2-protocol-tunneling vlan on page 2200 • show vlans • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) • show vlans on page 2251 • Configuring Layer 2 Protocol Tunneling on page 2063
List of Sample Output	show ethernet-switching layer2-protocol-tunneling statistics on page 2198 show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0 on page 2198 show ethernet-switching layer2-protocol-tunneling statistics vlan v2 on page 2198
Output Fields	Table 205 on page 2198 lists the output fields for the show ethernet-switching layer2-protocol-tunneling statistics command. Output fields are listed in the approximate order in which they appear.

Table 205: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
VLAN	Name of a VLAN on which L2PT has been configured.
Interface	Name of an interface on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are all , 802.1x , 802.3ah , cdp , e-lmi , gvrp , lACP , lldp , mmrp , mvrp , stp , udld , vstp , and vtp .
Operation	Type of operation being performed on the interface. Values are Encapsulation and Decapsulation .
Packets	Number of packets that have been encapsulated or de-encapsulated.
Drops	Number of packets that have exceeded the drop threshold and have been dropped.
Shutdowns	Number of times that packets have exceeded the shutdown threshold and the interface has been shut down.

Sample Output

show ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics
```

```

Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v1    xe-0/0/0.0  mvrp     Encapsulation  0        0        0
v1    xe-0/0/1.0  mvrp     Decapsulation  0        0        0
v1    xe-0/0/2.0  mvrp     Decapsulation  60634    0        0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0        0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0        0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0        0

```

show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0
```

```

Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v1    xe-0/0/0.0  mvrp     Encapsulation  0        0        0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0        0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0        0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0        0
v2    xe-0/0/0.0  mvrp     Encapsulation  0        0        0
v2    xe-0/0/0.0  stp      Encapsulation  0        0        0
v2    xe-0/0/0.0  vtp      Encapsulation  0        0        0
v2    xe-0/0/0.0  vstp     Encapsulation  0        0        0

```

show ethernet-switching layer2-protocol-tunneling statistics vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

```

Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v2    xe-0/0/0.0  cdp      Encapsulation  0        0        0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0        0

```

v2	xe-0/0/0.0	lldp	Encapsulation	0	0	0
v2	xe-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	xe-0/0/0.0	stp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vtp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vstp	Encapsulation	0	0	0
v2	xe-0/0/1.0	cdp	Decapsulation	0	0	0
v2	xe-0/0/1.0	gvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	lldp	Decapsulation	0	0	0
v2	xe-0/0/1.0	mvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	stp	Decapsulation	0	0	0
v2	xe-0/0/1.0	vtp	Decapsulation	0	0	0

show ethernet-switching layer2-protocol-tunneling vlan

Syntax	<code>show ethernet-switching-layer2-protocol-tunneling vlan <vlan-name></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.
Options	none —Display information about L2PT for the VLANs on which you have configured L2PT. vlan-name —(Optional) Display information about L2PT for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching layer2-protocol-tunneling interface on page 2195 • show ethernet-switching layer2-protocol-tunneling statistics on page 2197 • show vlans • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) • show vlans on page 2251 • Configuring Layer 2 Protocol Tunneling on page 2063
List of Sample Output	show ethernet-switching layer2-protocol-tunneling vlan on page 2201 show ethernet-switching layer2-protocol-tunneling vlan v2 on page 2201
Output Fields	Table 206 on page 2200 lists the output fields for the show ethernet-switching layer2-protocol-tunneling vlan command. Output fields are listed in the approximate order in which they appear.

Table 206: show ethernet-switching layer2-protocol-tunneling vlan Output Fields

Field Name	Field Description
VLAN	Name of the VLAN on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are all , 802.1x , 802.3ah , cdp , e-lmi , gvrp , lACP , lldp , mmrp , mvrp , stp , vstp , and vtp .
Drop Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the switch begins dropping the Layer 2 PDUs.
Shutdown Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the interface is disabled.

Sample Output

show ethernet-switching layer2-protocol-tunneling vlan

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan
```

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
                Threshold Threshold
v1             mvrp          100           200
v2             cdp            0             0
v2             cdp            0             0
v2             gvrp           0             0
```

show ethernet-switching layer2-protocol-tunneling vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan v2
```

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
                Threshold Threshold
v2             cdp            0             0
v2             cdp            0             0
v2             gvrp           0             0
```

show ethernet-switching mac-learning-log

Syntax	show ethernet-switching mac-learning-log
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Displays the event log of learned MAC addresses.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching table on page 2211 • show ethernet-switching interfaces on page 1838
List of Sample Output	show ethernet-switching mac-learning-log on page 2202
Output Fields	Table 207 on page 2202 lists the output fields for the show ethernet-switching mac-learning-log command. Output fields are listed in the approximate order in which they appear.

Table 207: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted Added	MAC address deleted from or added to the MAC learning log.
Blocking	Forwarding state of the interface: <ul style="list-style-type: none"> • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log

```

user@switch> show ethernet-switching mac-learning-log
Mon Feb 25 08:07:05 2008
  vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v12 mac 00:00:00:00:00:00 was deleted

```

```
Mon Feb 25 08:07:05 2008
  vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
  vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching mac-notification

Syntax	show ethernet-switching mac-notification
Release Information	Command introduced in Junos OS Release 9.6 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about MAC notification.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Verifying That MAC Notification Is Working Properly</i>
List of Sample Output	show ethernet-switching mac-notification (MAC Notification Enabled) on page 2204 show ethernet-switching mac-notification (MAC Notification Disabled) on page 2204
Output Fields	Table 208 on page 2204 lists the output fields for the show ethernet-switching mac-notification command. Output fields are listed in the order in which they appear.

Table 208: show ethernet-switching mac-notification Output Fields

Field Name	Field Description
Notification Status	MAC notification status: <ul style="list-style-type: none"> • Enabled—MAC notification is enabled. • Disabled—MAC notification is disabled.
Notification Interval	MAC notification interval in seconds.

Sample Output

show ethernet-switching mac-notification (MAC Notification Enabled)

```
user@switch> show ethernet-switching mac-notification
Notification Status      : Enabled
Notification Interval    : 30
```

Sample Output

show ethernet-switching mac-notification (MAC Notification Disabled)

```
user@switch> show ethernet-switching mac-notification
Notification Status      : Disabled
Notification Interval    : 0
```


show ethernet-switching statistics aging

Syntax	show ethernet-switching statistics aging <brief detail>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display media access control (MAC) aging statistics.
Options	none —(Optional) Display MAC aging statistics. brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching statistics mac-learning on page 2207 • mac-table-aging-time on page 2101 • Configuring MAC Table Aging on page 2044
List of Sample Output	show ethernet-switching statistics aging on page 2206
Output Fields	Table 209 on page 2205 lists the output fields for the show ethernet-switching statistics aging command. Output fields are listed in the approximate order in which they appear.

Table 209: show ethernet-switching statistics aging Output Fields

Field Name	Field Description	Level of Output
Total age messages received	Total number of aging messages received from the hardware.	All levels
Immediate aging	Aging message indicating that the entry should be removed immediately.	All levels
MAC address seen	Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.	All levels
MAC address not seen	Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.	All levels
Error age messages	The received aging message contains the following errors: <ul style="list-style-type: none"> • Invalid VLAN—The VLAN of the packet does not exist. • No such entry—The MAC address and VLAN pair provided by the aging message does not exist. • Static entry—An unsuccessful attempt was made to age out a static MAC entry. 	All levels

Sample Output

show ethernet-switching statistics aging

```
user@switch> show ethernet-switching statistics aging
```

```
Total age messages received: 0
```

```
Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
```

```
Error age messages: 0
```

```
Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

show ethernet-switching statistics mac-learning

Syntax	show ethernet-switching statistics mac-learning <brief detail> <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display media access control (MAC) learning statistics.
Options	<p>none—(Optional) Display MAC learning statistics for all interfaces.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>interface <i>interface-name</i>—(Optional) Display MAC learning statistics for the specified interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>show ethernet-switching statistics aging</i> • <i>show ethernet-switching mac-learning-log</i> • <i>show ethernet-switching table</i> • <i>show ethernet-switching interfaces</i> • <i>Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</i> • <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i> • show ethernet-switching statistics aging on page 2205 • show ethernet-switching mac-learning-log on page 2202 • show ethernet-switching table on page 2211 • show ethernet-switching interfaces on page 1838 • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962 • Example: Setting Up Bridging with Multiple VLANs on page 1971
List of Sample Output	show ethernet-switching statistics mac-learning on page 2208 show ethernet-switching statistics mac-learning detail on page 2209 show ethernet-switching statistics mac-learning interface ge-0/0/28 detail on page 2209 show ethernet-switching statistics mac-learning interface on page 2209 show ethernet-switching statistics mac-learning detail (QFX Series) on page 2209
Output Fields	Table 210 on page 2208 lists the output fields for the show ethernet-switching statistics mac-learning command. Output fields are listed in the approximate order in which they appear.

Table 210: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface for which statistics are being reported. (Displayed in the output under the heading Interface .)	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading Local pkts .)	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading Transit pkts .)	All levels
Learning message with error	<p>MAC learning messages received with errors (Displayed under the heading Error):</p> <ul style="list-style-type: none"> • Invalid VLAN—The VLAN of the packet does not exist. • Invalid MAC—The MAC address is either NULL or a multicast MAC address. • Security violation—The MAC address is not an allowed MAC address. • Interface down—The MAC address is learned on an interface that is down. • Incorrect membership—The MAC address is learned on an interface that is not a member of the VLAN. • Interface limit—The number of MAC addresses learned on the interface has exceeded the limit. • MAC move limit—This MAC address has moved among multiple interfaces too many times in a given interval. • VLAN limit—The number of MAC addresses learned on the VLAN has exceeded the limit. • VLAN membership limit—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit. • Invalid VLAN index—The VLAN of the packet, although configured, does not yet exist in the kernel. • Interface not learning—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked. • No nexthop—The MAC address is learned on an interface that does not have a unicast next hop. • MAC learning disabled—The MAC address is learned on an interface on which MAC learning has been disabled. • Others—The message contains some other error. 	All levels

Sample Output

show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning
```

```
Learning stats: 0 learn msg rcvd, 0 error
Interface      Local pkts      Transit pkts      Error
ge-0/0/0.0     0               0                 0
ge-0/0/1.0     0               0                 0
ge-0/0/2.0     0               0                 0
ge-0/0/3.0     0               0                 0
```

show ethernet-switching statistics mac-learning detail

```

user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

Interface: ge-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0

```

```

Interface: ge-0/0/1.0
Learning message from local packets: 0
Learning message from transit packets: 2
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0

```

show ethernet-switching statistics mac-learning interface ge-0/0/28 detail

```

user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/28 detail

Interface: ge-0/0/28.0
Learning message from local packets: 0
Learning message from transit packets: 5
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
                          VLAN membership limit: 20
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0

```

show ethernet-switching statistics mac-learning interface

```

user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1

```

Interface	Local pkts	Transit pkts	Error
ge-0/0/1.0	0	1	1

show ethernet-switching statistics mac-learning detail (QFX Series)

```

user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

Interface: xe-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0

```

Invalid VLAN:	0	Invalid MAC:	0
Security violation:	0	Interface down:	0
Incorrect membership:	0	Interface limit:	0
MAC move limit:	0	VLAN limit:	0
Invalid VLAN index:	0	Interface not learning:	0
No nexthop:	0	MAC learning disabled:	0
Others:	0		

Interface: xe-0/0/1.0

Learning message from local packets: 0

Learning message from transit packets: 2

Learning message with error: 0

Invalid VLAN:	0	Invalid MAC:	0
Security violation:	0	Interface down:	0
Incorrect membership:	0	Interface limit:	0
MAC move limit:	0	VLAN limit:	0
Invalid VLAN index:	0	Interface not learning:	0
No nexthop:	0	MAC learning disabled:	0
Others:	0		

show ethernet-switching table

Syntax	show ethernet-switching table <brief detail extensive summary> <interface <i>interface-name</i> > <management-vlan> <sort-by (<i>name</i> <i>tag</i>)> <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.
Description	Displays the Ethernet switching table.
Options	<p>none—(Optional) Display brief information about the Ethernet switching table.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p>management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p>sort-by (<i>name</i> <i>tag</i>)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.</p> <p>vlan <i>vlan-name</i>—(Optional) Display the Ethernet switching table for a specific VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962 • Example: Setting Up Bridging with Multiple VLANs on page 1971
List of Sample Output	show ethernet-switching table on page 2212 show ethernet-switching table (Private VLANs) on page 2213 show ethernet-switching table brief on page 2213 show ethernet-switching table detail on page 2213 show ethernet-switching table extensive on page 2215 show ethernet-switching table interface on page 2216
Output Fields	Table 211 on page 2211 lists the output fields for the show ethernet-switching table command. Output fields are listed in the approximate order in which they appear.

Table 211: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of a VLAN.	All levels

Table 211: show ethernet-switching table Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC address	MAC address associated with the VLAN.	All levels
Type	Type of MAC address: <ul style="list-style-type: none"> • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members. 	All levels
Age	Time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or with the All-members option (flood entry).	All levels
Learned	For learned entries, the time at which the entry was added to the Ethernet switching table.	detail, extensive

Sample Output

show ethernet-switching table

```

user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age  Interfaces
F2         *                Flood     -    All-members
F2         00:00:05:00:00:03 Learn     0    xe-0/0/44.0
F2         00:19:e2:50:7d:e0 Static    -    Router
Linux      *                Flood     -    All-members
Linux      00:19:e2:50:7d:e0 Static    -    Router
Linux      00:30:48:90:54:89 Learn     0    xe-0/0/47.0
T1         *                Flood     -    All-members
T1         00:00:05:00:00:01 Learn     0    xe-0/0/46.0
T1         00:00:5e:00:01:00 Static    -    Router
T1         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T1         00:19:e2:50:7d:e0 Static    -    Router
T10        *                Flood     -    All-members
T10        00:00:5e:00:01:09 Static    -    Router
T10        00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T10        00:19:e2:50:7d:e0 Static    -    Router
T111       *                Flood     -    All-members
T111       00:19:e2:50:63:e0 Learn     0    xe-0/0/15.0
T111       00:19:e2:50:7d:e0 Static    -    Router
T111       00:19:e2:50:ac:00 Learn     0    xe-0/0/15.0
T2         *                Flood     -    All-members
T2         00:00:5e:00:01:01 Static    -    Router
T2         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T2         00:19:e2:50:7d:e0 Static    -    Router
T3         *                Flood     -    All-members
T3         00:00:5e:00:01:02 Static    -    Router
T3         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T3         00:19:e2:50:7d:e0 Static    -    Router
T4         *                Flood     -    All-members

```



```

T4          00:00:5e:00:01:03 Static      - Router
T4          00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
[output truncated]

```

show ethernet-switching table (Private VLANs)

```

user@switch> show ethernet-switching table
Ethernet-switching table: 10 entries, 3 learned
VLAN      MAC address      Type      Age Interfaces
pvlan     *                Flood     - All-members
pvlan     00:10:94:00:00:02 Replicated - xe-0/0/28.0
pvlan     00:10:94:00:00:35 Replicated - xe-0/0/46.0
pvlan     00:10:94:00:00:46 Replicated - xe-0/0/4.0
c2        *                Flood     - All-members
c2        00:10:94:00:00:02 Learn        0 xe-0/0/28.0
c1        *                Flood     - All-members
c1        00:10:94:00:00:46 Learn        0 xe-0/0/4.0
__pvlan_pvlan_xe-0/0/46.0__ *          Flood     - All-members
__pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn    0 xe-0/0/46.0

```

show ethernet-switching table brief

```

user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age Interfaces
F2        *                Flood     - All-members
F2        00:00:05:00:00:03 Learn        0 xe-0/0/44.0
F2        00:19:e2:50:7d:e0 Static      - Router
Linux     *                Flood     - All-members
Linux     00:19:e2:50:7d:e0 Static      - Router
Linux     00:30:48:90:54:89 Learn        0 xe-0/0/47.0
T1        *                Flood     - All-members
T1        00:00:05:00:00:01 Learn        0 xe-0/0/46.0
T1        00:00:5e:00:01:00 Static      - Router
T1        00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
T1        00:19:e2:50:7d:e0 Static      - Router
T10       *                Flood     - All-members
T10       00:00:5e:00:01:09 Static      - Router
T10       00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
T10       00:19:e2:50:7d:e0 Static      - Router
T111     *                Flood     - All-members
T111     00:19:e2:50:63:e0 Learn        0 xe-0/0/15.0
T111     00:19:e2:50:7d:e0 Static      - Router
T111     00:19:e2:50:ac:00 Learn        0 xe-0/0/15.0
T2        *                Flood     - All-members
T2        00:00:5e:00:01:01 Static      - Router
T2        00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
T2        00:19:e2:50:7d:e0 Static      - Router
T3        *                Flood     - All-members
T3        00:00:5e:00:01:02 Static      - Router
T3        00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
T3        00:19:e2:50:7d:e0 Static      - Router
T4        *                Flood     - All-members
T4        00:00:5e:00:01:03 Static      - Router
T4        00:19:e2:50:63:e0 Learn        0 xe-0/0/46.0
[output truncated]

```

show ethernet-switching table detail

```

user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *

```

```
Interface(s): xe-0/0/44.0
Type: Flood
Nexthop index: 0

F2, 00:00:05:00:00:03
Interface(s): xe-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
Nexthop index: 0

F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, *
Interface(s): xe-0/0/47.0
Type: Flood
Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, 00:30:48:90:54:89
Interface(s): xe-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T1, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T1, 00:00:05:00:00:01
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
Nexthop index: 0

T1, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T10, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T10, 00:00:5e:00:01:09
Interface(s): Router
```

```

Type: Static
Nexthop index: 0

T10, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T111, *
Interface(s): xe-0/0/15.0
Type: Flood
Nexthop index: 0
[output truncated]

```

show ethernet-switching table extensive

```

user@switch> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): xe-0/0/44.0
Type: Flood
Nexthop index: 0

F2, 00:00:05:00:00:03
Interface(s): xe-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
Nexthop index: 0

F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, *
Interface(s): xe-0/0/47.0
Type: Flood
Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, 00:30:48:90:54:89
Interface(s): xe-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T1, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T1, 00:00:05:00:00:01
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07

```

```
Nexthop index: 0

T1, 00:00:5e:00:01:00
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T1, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T10, 00:00:5e:00:01:09
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:08
  Nexthop index: 0

T10, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T111, *
  Interface(s): xe-0/0/15.0
  Type: Flood
  Nexthop index: 0
[output truncated]
```

show ethernet-switching table interface

```
user@switch> show ethernet-switching table interface xe-0/0/1
Ethernet-switching table: 1 unicast entries

```

VLAN	MAC address	Type	Age	Interfaces
V1	*	Flood	-	All-members
V1	00:00:05:00:00:05	Learn	0	xe-0/0/1.0

show interfaces xe

Syntax	<pre>show interfaces <i>device-name:type-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <routing-instance (all <i>instance-name</i>)> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.
Options	<p><i>device-name:type-fpc/pic/port</i>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>routing-instance (all <i>instance-name</i>)—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 317 • Troubleshooting Network Interfaces on page 1160 • Troubleshooting an Aggregated Ethernet Interface on page 1161 • Junos® OS Network Interfaces
List of Sample Output	<p>show interfaces on page 2225</p> <p>show interfaces (Asymmetric Flow Control) on page 2226</p> <p>show interfaces brief on page 2226</p> <p>show interfaces detail on page 2226</p> <p>show interfaces detail (Asymmetric Flow Control) on page 2228</p> <p>show interfaces extensive on page 2229</p> <p>show interfaces extensive (Asymmetric Flow Control) on page 2231</p>

[show interfaces terse on page 2233](#)

[show interfaces \(QFabric System\) on page 2233](#)

Output Fields Table 212 on page 2218 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 212: show interfaces xe Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex .	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
NOTE: This field is only displayed if asymmetric flow control is not configured.		

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured-flow-control	Configured flow control for the interface transmit buffers (tx-buffers) and receive buffers (rx-buffers): <ul style="list-style-type: none"> tx-buffers—On if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. Off if the interface is not configured to respond to received PAUSE messages. rx-buffers—On if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. Off if the interface is not configured to generate and send PAUSE messages. <p>NOTE: This field is only displayed if asymmetric flow control is configured.</p>	All levels
Auto-negotiation	Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	Remote fault status: <ul style="list-style-type: none"> Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The Mapped forwarding class column lists the forwarding classes mapped to each CoS queue.	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	detail extensive
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of packets that exceeds the configured MTU. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Filter statistics	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	extensive

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are enable or disable. • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> Destination slot—FPC slot number. CoS transmit queue—Queue number and its associated user-configured forwarding class name. Bandwidth %—Percentage of bandwidth allocated to the queue. Bandwidth bps—Bandwidth allocated to the queue (in bps). Buffer %—Percentage of buffer space allocated to the queue. Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. Priority—Queue priority: low or high. Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	detail extensive
IPv6 transit statistics	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	extensive
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none

Table 212: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	detail extensive
Addresses, Flags	Information about the address flags.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0

```

```
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode
```

show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode
```

show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch
```

show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)

Traffic statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps
Input packets:	0	0 pps
Output packets:	0	0 pps

IPv6 transit statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Egress queues: 12 supported, 9 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 fc7	0	0	0
2 no-loss	0	0	0
3 fcoe	0	0	0
4 fc4	0	0	0
5 fc5	0	0	0
6 fc6	0	0	0
7 network-cont	0	0	0
8 mcast	0	0	0

Queue number:	Mapped forwarding classes
0	best-effort
1	fc7
2	no-loss
3	fcoe
4	fc4
5	fc5
6	fc6
7	network-control
8	mcast

Active alarms : None

Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Local statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Transit statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps

```

Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
    Input bytes :          0          0 bps
    Output bytes :          0          0 bps
    Input packets:          0          0 pps
    Output packets:          0          0 pps
  IPv6 transit statistics:
    Input bytes :          0
    Output bytes :          0
    Input packets:          0
    Output packets:          0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort          0              0              0
    1 fc7                 0              0              0
    2 no-loss              0              0              0
    3 fcoe                 0              0              0
    4 fc4                  0              0              0
    5 fc5                  0              0              0
    6 fc6                  0              0              0
    7 network-cont         0              0              0
    8 mcast                0              0              0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                fc7
    2                no-loss
    3                fcoe
    4                fc4
    5                fc5
    6                fc6

```



```

7          network-control
8          mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

0 best-effort	0	0	0
1 fc7	0	0	0
2 no-loss	0	0	0
3 fcoe	0	0	0
4 fc4	0	0	0
5 fc5	0	0	0
6 fc6	0	0	0
7 network-cont	0	0	0
8 mcast	0	0	0

Queue number: Mapped forwarding classes

0	best-effort
1	fc7
2	no-loss
3	fcoe
4	fc4
5	fc5
6	fc6
7	network-control
8	mcast

Active alarms : None

Active defects : None

MAC statistics:

	Receive	Transmit
Total octets	0	0
Total packets	0	0
Unicast packets	0	0
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	

MAC Priority Flow Control Statistics:

Priority : 0	0	0
Priority : 1	0	0
Priority : 2	0	0
Priority : 3	0	0
Priority : 4	0	0
Priority : 5	0	0
Priority : 6	0	0
Priority : 7	0	0

Filter statistics:

Input packet count	0	
Input packet rejects	0	
Input DA rejects	0	
Input SA rejects	0	
Output packet count		0

```

Output packet pad count                                0
Output packet error count                              0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      %      bps      %      usec
0 best-effort      75      7500000000      75      0      low
none
7 network-control      5      500000000      5      0      low
none
8 mcast      20      2000000000      20      0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
IPv6 transit statistics:
  Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0                0                0
1 fc7                  0                0                0
2 no-loss              0                0                0
3 fcoe                 0                0                0
4 fc4                  0                0                0
5 fc5                  0                0                0
6 fc6                  0                0                0
7 network-cont         0                0                0
8 mcast                0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    0            0
Unicast packets  0            0
Broadcast packets 0            0
Multicast packets 0            0
CRC/Align errors 0            0
FIFO errors       0            0
MAC control frames 0            0
MAC pause frames  0            0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
MAC Priority Flow Control Statistics:
Priority : 0       0            0
Priority : 1       0            0

```

```

Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         75    7500000000    75      0    low    none
7 network-control     5     500000000     5      0    low    none
8 mcast               20    2000000000    20      0    low    none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0          0 bps
Output bytes : 0          0 bps
Input packets: 0          0 pps
Output packets: 0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface      Admin Link Proto  Local      Remote

xe-0/0/1       up    up
xe-0/0/1.0     up    up    eth-switch

```

show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000
CoS queues : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

show spanning-tree bridge

List of Syntax [Syntax on page 2235](#)
[Syntax \(QFX Series\) on page 2235](#)

Syntax `show spanning-tree bridge`
 `<brief | detail>`
 `<msti msti-id>`
 `<routing-instance routing-instance-name>`
 `<vlan-id vlan-id>`

Syntax (QFX Series) `show spanning-tree bridge`
 `<brief | detail>`
 `<msti msti-id>`
 `<vlan-id vlan-id>`

Release Information Command introduced in Junos OS Release 8.4.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the configured or calculated Spanning Tree Protocol (STP) parameters.

Options **none**—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).

brief | detail—(Optional) Display the specified level of output.

msti *msti-id*—(Optional) Display STP bridge information for the specified MSTI.

routing-instance *routing-instance-name*—(Optional) Display STP bridge information for the specified routing instance.

vlan-id *vlan-id*—(Optional) Display STP bridge information for the specified VLAN.

Required Privilege Level view

List of Sample Output [show spanning-tree bridge routing-instance on page 2236](#)
 [show spanning-tree bridge msti on page 2237](#)
 [show spanning-tree bridge vlan-id \(MSTP\) on page 2238](#)
 [show spanning-tree bridge \(RSTP\) on page 2238](#)
 [show spanning-tree bridge vlan-id \(RSTP\) on page 2239](#)

Output Fields [Table 213 on page 2235](#) lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

Table 213: show spanning-tree bridge Output Fields

Field Name	Field Description
Routing instance name	Name of the routing instance under which the bridge is configured.
Enabled protocol	Spanning Tree Protocol type enabled.

Table 213: show spanning-tree bridge Output Fields (*continued*)

Field Name	Field Description
Root ID	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs).
Maximum age	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
Forward delay	How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
Message age	Number of elapsed seconds since the most recent BPDU was received.
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

Sample Output

show spanning-tree bridge routing-instance

```

user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name           : vs1
Enabled protocol                 : MSTP

```



```

STP bridge parameters for CIST
  Root ID          : 32768.00:13:c3:9e:c8:80
  Root cost        : 0
  Root port        : ge-10/2/0
  CIST regional root : 32768.00:13:c3:9e:c8:80
  CIST internal root cost : 22000
  Hello time       : 2 seconds
  Maximum age      : 20 seconds
  Forward delay    : 15 seconds
  Hop count        : 18
  Message age      : 0
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID      : 32768.00:90:69:0b:7f:d1
    Extended system ID : 1

STP bridge parameters for MSTI 1
  MSTI regional root : 32769.00:13:c3:9e:c8:80
  Root cost          : 22000
  Root port          : ge-10/2/0
  Hello time         : 2 seconds
  Maximum age        : 20 seconds
  Forward delay      : 15 seconds
  Hop count          : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID      : 32769.00:90:69:0b:7f:d1
    Extended system ID : 1

STP bridge parameters for MSTI 2
  MSTI regional root : 32770.00:13:c3:9e:c8:80
  Root cost          : 22000
  Root port          : ge-10/2/0
  Hello time         : 2 seconds
  Maximum age        : 20 seconds
  Forward delay      : 15 seconds
  Hop count          : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID      : 32770.00:90:69:0b:7f:d1
    Extended system ID : 1

```

show spanning-tree bridge msti

```

user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name      : vs1
Enabled protocol          : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root      : 32769.00:13:c3:9e:c8:80
  Root cost               : 22000
  Root port               : xe-10/2/0
  Hello time              : 2 seconds
  Maximum age             : 20 seconds
  Forward delay           : 15 seconds
  Hop count               : 18

```

```
Number of topology changes      : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                     : 32769.00:90:69:0b:7f:d1
  Extended system ID            : 1
```

show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1101 routing-instance vs1 detail
```

```
STP bridge parameters
Routing instance name          : vs1
Enabled protocol               : MSTP

STP bridge parameters for CIST
Root ID                       : 32768.00:13:c3:9e:c8:80
Root cost                     : 0
Root port                     : xe-10/2/0
CIST regional root            : 32768.00:13:c3:9e:c8:80
CIST internal root cost       : 22000
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Hop count                    : 18
Message age                   : 0
Number of topology changes    : 0
Local parameters
  Bridge ID                   : 32768.00:90:69:0b:7f:d1
  Extended system ID          : 1
  Hello time                  : 2 seconds
  Maximum age                 : 20 seconds
  Forward delay               : 15 seconds
  Path cost method            : 32 bit
  Maximum hop count           : 20
```

show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
```

```
STP bridge parameters
Routing instance name          : GLOBAL
Enabled protocol               : RSTP
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0
  Extended system ID          : 0

STP bridge parameters for bridge VLAN 10
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0
```

```
Extended system ID          : 0

STP bridge parameters for bridge VLAN 20
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID         : 0
```

show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name        : GLOBAL
Enabled protocol              : RSTP

STP bridge parameters for VLAN 10
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID         : 0
```

show spanning-tree interface

List of Syntax	Syntax on page 2240 Syntax (EX Series Switches and the QFX Series) on page 2240
Syntax	<pre>show spanning-tree interface <brief detail> <msti <i>msti-id</i>> <routing-instance <i>routing-instance-name</i>> <vlan-id <i>vlan-id</i>></pre>
Syntax (EX Series Switches and the QFX Series)	<pre>show spanning-tree interface <brief detail> <msti <i>msti-id</i>> <vlan-id <i>vlan-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display the configured or calculated interface-level STP parameters.
Options	<p>none—Display brief STP interface information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>msti <i>msti-id</i>—(Optional) Display STP interface information for the specified MST instance.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP interface information for the specified routing instance.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display STP interface information for the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	show spanning-tree interface on page 2241 show spanning-tree interface (QFX Series) on page 2242 show spanning-tree interface detail on page 2242 show spanning-tree interface msti on page 2244 show spanning-tree interface vlan-id on page 2244 show spanning-tree interface (VSTP) on page 2245 show spanning-tree interface vlan-id (VSTP) on page 2245
Output Fields	<p>Table 214 on page 2240 lists the output fields for the show spanning-tree interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 214: show spanning-tree Interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.

Table 214: show spanning-tree Interface Output Fields (*continued*)

Field Name	Field Description
Port ID	Logical interface identifier configured to participate in the MSTP or VSTP instance.
Designated port ID	Port ID of the designated port for the LAN segment to which this interface is attached.
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: forwarding (FWD), blocking (BLK), listening, learning, or disabled.
Port Role	MSTP, VSTP, or RSTP port role: designated (DESG), backup (BKUP), alternate (ALT), (ROOT), or Root Prevented (Root-Prev).
Link type	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
Alternate	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (Yes) or nonalternate root port (No).
Boundary Port	Identifies the interface as an MSTP regional boundary port (Yes) or nonboundary port (No).

Sample Output

show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

```
Spanning tree interface parameters for instance 1
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 1

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
```

Spanning tree interface parameters for instance 0

```
Interface name           : ae1
Port identifier          : 128.1
Designated port ID      : 128.1
Port cost                : 1000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
```

```

Boundary port                : No

Interface name                : ge-2/1/2
Port identifier               : 128.2
Designated port ID           : 128.2
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

```

```

Interface name                : ge-2/1/5
Port identifier               : 128.3
Designated port ID           : 128.3
Port cost                     : 29999
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

```

```

Interface name                : ge-2/2/1
Port identifier               : 128.4
Designated port ID           : 128.26
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:13:c3:9e:c8:80
Port role                     : Root
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

```

```

Interface name                : xe-9/2/0
Port identifier               : 128.5
Designated port ID           : 128.5
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

```

```

Interface name                : xe-9/3/0
Port identifier               : 128.6
Designated port ID           : 128.6
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

```

Spanning tree interface parameters for instance 1

```

Interface name                : ae1
Port identifier               : 128.1
Designated port ID           : 128.1
Port cost                     : 1000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1

```

```

Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/2
Port identifier      : 128.2
Designated port ID   : 128.2
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/5
Port identifier      : 128.3
Designated port ID   : 128.3
Port cost            : 29999
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/2/1
Port identifier      : 128.4
Designated port ID   : 128.26
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role           : Root
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

...

```

show spanning-tree interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-7/0/0	128:1	128:1	32769.0090690b4fd1	2000	FWD	DESG
ge-5/1/0	128:2	128:2	32769.0090690b4fd1	20000	FWD	DESG
ge-5/1/1	128:3	128:3	32769.0090690b4fd1	20000	FWD	DESG
ae1	128:4	128:1	32769.0090690b47d1	10000	BLK	ALT
ge-5/1/4	128:5	128:3	32769.0090690b47d1	20000	BLK	ALT
xe-7/2/0	128:6	128:6	32769.0090690b47d1	2000	FWD	ROOT

show spanning-tree interface vlan-id

```

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-11/0/5	128:1	128:1	32768.0090690b7fd1	20000	FWD	DESG
ge-11/0/6	128:2	128:1	32768.0090690b7fd1	20000	BLK	BKUP
ge-11/1/0	128:3	128:2	32768.0090690b4fd1	20000	BLK	ALT
ge-11/1/1	128:4	128:3	32768.0090690b4fd1	20000	BLK	ALT

ge-11/1/4	128:5	128:1	32768.0090690b47d1	20000	BLK	ALT
xe-10/0/0	128:6	128:5	32768.0090690b4fd1	2000	BLK	ALT
xe-10/2/0	128:7	128:4	32768.0090690b47d1	2000	FWD	ROOT

show spanning-tree interface (VSTP)

```
user@host> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

Spanning tree interface parameters for VLAN 10

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

Spanning tree interface parameters for VLAN 20

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

show spanning-tree interface vlan-id (VSTP)

```
user@host> show spanning-tree interface vlan-id 10
```

Spanning tree interface parameters for VLAN 10

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

show spanning-tree mstp configuration

List of Syntax	Syntax on page 2246 Syntax (EX Series Switch and the QFX Series) on page 2246
Syntax	show spanning-tree mstp configuration <brief detail> <routing-instance <i>routing-instance-name</i> >
Syntax (EX Series Switch and the QFX Series)	show spanning-tree mstp configuration <brief detail>
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the MSTP configuration.
Options	none—Display MSTP configuration information. brief detail—(Optional) Display the specified level of output. routing-instance <i>routing-instance-name</i> —(Optional) Display MSTP configuration information for the specified routing instance.
Required Privilege Level	view
List of Sample Output	show spanning-tree mstp configuration detail on page 2247 show spanning-tree mstp configuration detail (QFX Series) on page 2247
Output Fields	Table 215 on page 2246 lists the output fields for the show spanning-tree mstp configuration command. Output fields are listed in the approximate order in which they appear.

Table 215: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context id	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI	MST instance identifier.
Member VLANs	VLAN identifiers associated with the MSTI.

Sample Output

show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

show spanning-tree mstp configuration detail (QFX Series)

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

show spanning-tree statistics

List of Syntax	Syntax on page 2248 Syntax (EX Series Switch and the QFX Series) on page 2248
Syntax	<pre>show spanning-tree statistics <brief detail> <interface <i>interface-name</i>> <routing-instance <i>routing-instance-name</i>></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show spanning-tree statistics <brief detail> <interface <i>interface-name</i> vlan <i>vlan-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series switches.</p>
Description	Display STP statistics.
Options	<p>none—Display brief STP statistics.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display STP statistics for the specified interface.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP statistics for the specified routing instance.</p>
Required Privilege Level	view
List of Sample Output	show spanning-tree statistics routing-instance on page 2249 show spanning-tree statistics interface routing-instance detail on page 2249
Output Fields	<p>Table 216 on page 2248 lists the output fields for the show spanning-tree statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 216: show spanning-tree statistics Output Fields

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last 5 secs	Number of BPDUs sent within a specified interval.
BPDUs received in last 5 secs	Number of BPDUs received within a specified interval.

Table 216: show spanning-tree statistics Output Fields (*continued*)

Field Name	Field Description
Interface	Interface for which the statistics are being displayed.
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

Sample Output

show spanning-tree statistics routing-instance

```

user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type           : bpdus
BPDUs sent             : 1396
BPDUs received         : 1027
BPDUs sent in last interval : 5      (duration: 4 sec)
BPDUs received in last interval: 4    (duration: 4 sec)

```

show spanning-tree statistics interface routing-instance detail

```

user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface  BPDUs sent  BPDUs received  Next BPDU
                                     transmission
ge-11/1/4      7           190           0

```

show system statistics arp

Syntax	show system statistics arp
Release Information	Command introduced in Junos OS Release 9.6 for EX Series switches.
Description	Display system-wide Address Resolution Protocol (ARP) statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Proxy ARP on an EX Series Switch</i>• Verifying That Proxy ARP Is Working Correctly on page 2181

Sample Output

```
user@switch> show system statistics arp
arp:
    90060 datagrams received
    34 ARP requests received
    610 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast source address
    0 datagrams with my own hardware address
    0 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    294 datagrams with source address duplicate to mine
    89113 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    309 ARP requests sent
    35 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor
```

show vlans

Syntax `show vlans`
`<brief | detail | extensive>`
`<dot1q-tunneling>`
`<sort-by (tag | name)>`
`<vlan-range-name>`

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.
Option **dot1q-tunneling** added in Junos OS Release 12.1 for the QFX Series.

Description Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.



NOTE: When a series of VLANs is created using the `vlan-range` statement, such VLAN names are preceded and followed by a double underscore. For example, a series of VLANs using the VLAN range 1 through 3 and the base VLAN name `marketing` would be displayed as `__marketing_1__`, `__marketing_2__`, and `__marketing_3__`.



NOTE: To display an 802.1X supplicant successfully authenticated in multiple-supplicant mode with dynamic VLAN movement, use the `show vlans vlan-name extensive` operational mode command, where *vlan-name* is the dynamic VLAN.

Options **none**—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

brief | detail | extensive—(Optional) Display the specified level of output.

sort-by (tag | name)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan-range-name—(Optional) Display VLANs in ascending order of VLAN range names.

Required Privilege Level `view`

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1962](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971](#)
- [Understanding Bridging on page 1869](#)
- [show ethernet-switching interfaces on page 1838](#)

List of Sample Output

- [show vlans on page 2254](#)
- [show vlans \(Private VLANs\) on page 2254](#)
- [show vlans brief on page 2255](#)
- [show vlans detail on page 2255](#)
- [show vlans extensive \(Port-Based\) on page 2256](#)
- [show vlans \(Q-in-Q Tunneling\) on page 2257](#)
- [show vlans extensive \(Q-in-Q Tunneling\) on page 2257](#)
- [show vlans extensive \(Q-in-Q Tunneling and L2TP\) on page 2257](#)
- [show vlans sort-by tag on page 2257](#)
- [show vlans sort-by name on page 2258](#)
- [show vlans tag on page 2259](#)

Output Fields Table 176 on page 1853 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

Table 217: show vlans Output Fields

Field Name	Field Description	Level of Output
Name	Name of a VLAN.	none, brief
Tag	802.1Q tag applied to this VLAN. If none is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or All-members option (flood entry). An asterisk (*) beside the interface indicates that the interface is UP .	All levels
Address	IP address.	none, brief
Ports Active /Total	Number of interfaces associated with a VLAN: Active indicates interfaces that are UP , and Total indicates interfaces that are active and inactive.	brief
VLAN	Name of a VLAN.	detail, extensive
Admin state	State of the interface. Values are: enabled —The interface is turned on, and the physical link is operational and can pass packets.	detail,extensive
MAC learning Status	Indicates if MAC learning is disabled.	detail, extensive
Description	Description for the VLAN.	detail,extensive
Primary IP	Primary IP address associated with a VLAN.	detail
Number of interfaces	Number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	detail, extensive
STP	Spanning tree associated with a VLAN.	detail,extensive
Tagged interfaces	Tagged interfaces with which a VLAN is associated.	detail,extensive

Table 217: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Untagged interfaces	Untagged interfaces with which a VLAN is associated.	detail, extensive
Dot1q Tunneling Status	Indicates if Q-in-Q tunneling is enabled.	extensive
Customer VLAN ranges	List of customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	extensive
Private VLAN Mode	The private VLAN mode for this VLAN. Values include Primary , Isolated , and Community .	extensive
Primary VLAN	Primary VLAN tag for this secondary VLAN.	extensive
Internal Index	VLAN index internal to Junos OS software.	extensive
Origin	Manner in which the VLAN was created: static or learn .	extensive
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X,	extensive
IP addresses	IP address associated with a VLAN.	extensive
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	extensive
Number of mapping rules	Number of mapping rules for Q-in-Q tunneling (Push) and VLAN translation (Swap).	
Secondary VLANs	Secondary VLANs associated with a primary VLAN.	extensive
Isolated VLANs	Isolated VLANs associated with a primary VLAN.	extensive
Community VLANs	Community VLANs associated with a primary VLAN.	extensive
VLANs summary	VLAN counts: <ul style="list-style-type: none"> • Total—Total number of VLANs on the switch. • Configured VLANs—Number of VLANs that are based on user-configured settings. • Internal VLANs—Number of VLANs created by the system with no explicit configuration or protocol—for example, the default VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership. • Temporary VLANs—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them. 	All levels

Table 217: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dot1q VLANs summary	802.1Q VLAN counts: <ul style="list-style-type: none"> • Total—Total number of 802.1Q-tagged and untagged VLANs on the switch. • Tagged VLANs—Number of 802.1Q-tagged VLANs. • Untagged VLANs—Number of untagged 802.1Q VLANs. • Private VLAN—Counts of the following kinds of 802.1Q private VLANs (PVLANS): <ul style="list-style-type: none"> • Primary VLANs—Number of primary forwarding private VLANs. • Community VLANs—Number of community transporting and forwarding private VLANs. • Isolated VLANs—Number of isolated receiving and forwarding private VLANs. • Inter-switch-isolated VLANs—Number of inter-switch isolated receiving and forwarding private VLANs. 	All levels
Dot1q Tunneled VLANs summary	Q-in-Q-tunneled VLAN counts: <ul style="list-style-type: none"> • Total—Total number of Q-in-Q-tunneled VLANs on the switch. • Private VLAN—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANS). 	All levels

Sample Output

show vlans

```
user@switch> show vlans
```

Name	Tag	Interfaces
default	None	xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0, xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0, xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0, xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0, xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0
v0001	1	xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0
v0002	2	None
v0003	3	None
v0004	4	None
v0005	5	None

show vlans (Private VLANs)

```
user@switch> show vlans
```

Name	Tag	Interfaces
__pvlan_pvlan_xe-0/0/46.0__		

```

c1                xe-0/0/44.0*, xe-0/0/46.0*
c2                xe-0/0/4.0*, xe-0/0/44.0*
default           xe-0/0/28.0*, xe-0/0/44.0*
pvlan            500      None
                  xe-0/0/4.0*, xe-0/0/28.0*, xe-0/0/44.0*, xe-0/0/46.0*

```

show vlans brief

```
user@switch> show vlans brief
```

Name	Tag	Address	Ports Active/Total
default	None		0/23
v0001	1		0/4
v0002	2		0/0
v0003	3		0/0
v0004	4		0/0
v0005	5		0/0
v0006	6		0/0
v0007	7		0/0
v0008	8		0/0
v0009	9		0/0
v0010	10		0/2
v0011	11		0/0
v0012	12		0/0
v0013	13		0/0
v0014	14		0/0
v0015	15		0/0
v0016	16		0/0

show vlans detail

```
user@switch> show vlans detail
```

```
VLAN: default, Tag: Untagged, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 23 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0,
xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0,
xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0,
```

```
Tagged interfaces: None
```

```
VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 4 (Active = 0)
```

```
Dot1q Tunneling Status: Enabled
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0,
```

```
VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 0 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: None
```

```
VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
Description: None
Primary IP: None, Number of interfaces: 0 (Active = 0)
STP: None, RTG: None
Untagged interfaces: None
Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

show vlans extensive (Port-Based)

```
user@switch> show vlans extensive
VLAN: default, created at Mon Feb  4 12:13:47 2008
Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
Description: None
Customer VLAN ranges:
    1-4100
Protocol: Port based
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    xe-0/0/34.0 (untagged, access)
    xe-0/0/33.0 (untagged, access)
    xe-0/0/32.0 (untagged, access)
    xe-0/0/31.0 (untagged, access)
    xe-0/0/30.0 (untagged, access)
    xe-0/0/29.0 (untagged, access)
    xe-0/0/28.0 (untagged, access)
    xe-0/0/27.0 (untagged, access)
    xe-0/0/26.0 (untagged, access)
    xe-0/0/25.0 (untagged, access)
    xe-0/0/19.0 (untagged, access)
    xe-0/0/18.0 (untagged, access)
    xe-0/0/17.0 (untagged, access)
    xe-0/0/16.0 (untagged, access)
    xe-0/0/15.0 (untagged, access)
    xe-0/0/14.0 (untagged, access)
    xe-0/0/13.0 (untagged, access)
    xe-0/0/11.0 (untagged, access)
    xe-0/0/9.0 (untagged, access)
    xe-0/0/8.0 (untagged, access)
    xe-0/0/3.0 (untagged, access)
    xe-0/0/2.0 (untagged, access)
    xe-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community 1
Isolated VLANs :
    __pvlan_pvlan_xe-0/0/3.0__
Community VLANs :
    comm1

VLAN: v0001, created at Mon Feb  4 12:13:47 2008
Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
```

```

xe-0/0/24.0 (tagged, trunk)
xe-0/0/23.0 (tagged, trunk)
xe-0/0/22.0 (tagged, trunk)
xe-0/0/21.0 (tagged, trunk)

```

```

VLAN: v0002, created at Mon Feb  4 12:13:47 2008
Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

```

VLAN: v0003, created at Mon Feb  4 12:13:47 2008
Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

show vlans (Q-in-Q Tunneling)

```

user@switch> show vlans dot1q-tunneling
Name      Tag      Interfaces
sv100     100      xe-0/0/4.0*, xe-0/0/15.0*

```

show vlans extensive (Q-in-Q Tunneling)

```

user@switch> show vlans sv100 extensive
VLAN: sv100, Created at: Sat Sep 10 12:53:52 2011
802.1Q Tag: 100, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    10-20
    40-50
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 1), Untagged 0 (Active = 0)
    ge-0/0/0.0, tagged, trunk

Number of mapping rules:
    Push 1 (Active = 0), Policy 0 (Active = 0), Swap 0 (Active = 0)

    xe-0/0/3.0*, 300, push

```

show vlans extensive (Q-in-Q Tunneling and L2TP)

```

user@switch> show vlans v1 extensive
VLAN: v1, Created at: Fri Mar 2 05:07:38 2012
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled

```

show vlans sort-by tag

```

user@switch> show vlans sort-by tag
Name      Tag      Interfaces
default           None
__vlan-x_1__  1

```

__vlan-x_2__	2	None
__vlan-x_3__	3	None
__vlan-x_4__	4	None
__vlan-x_5__	5	None
__vlan-x_6__	6	None
__vlan-x_7__	7	None
__vlan-x_8__	8	None
__vlan-x_9__	9	None
__vlan-x_10__	10	None
__vlan-x_11__	11	None
__vlan-x_12__	12	None
__vlan-x_13__	13	None
__vlan-x_14__	14	None
__vlan-x_15__	15	None
__vlan-x_16__	16	None
__vlan-x_17__	17	None
__vlan-x_18__	18	None
__vlan-x_19__	19	None
__vlan-x_20__	20	None

show vlans sort-by name

```
user@switch> show vlans sort-by employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*

```

__employee_128__ 128    xe-0/0/22.0*
__employee_129__ 129    xe-0/0/22.0*
__employee_130__ 130    xe-0/0/22.0*

```

show vlans tag

```
user@switch> show vlans employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

CHAPTER 23

Troubleshooting

- [Troubleshooting Procedures on page 2261](#)

Troubleshooting Procedures

- [Troubleshooting Ethernet Switching on page 2261](#)
- [Troubleshooting Layer 2 Protocol Tunneling on page 2262](#)
- [Troubleshooting Private VLANs on page 2263](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 2266](#)

Troubleshooting Ethernet Switching

Problem **Description:** Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

Solution Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

- Related Documentation**
- [arp on page 2072](#)
 - [mac-table-aging-time on page 2101](#)

Troubleshooting Layer 2 Protocol Tunneling

- [Drop Threshold Statistics Might Be Incorrect on page 2262](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 2262](#)

Drop Threshold Statistics Might Be Incorrect

Problem **Description:** L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets will not be reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit are not reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

Solution This is expected behavior.

Egress Filtering of L2PT Traffic Not Supported

Problem **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

Solution This is expected behavior.

- Related Documentation**
- [Understanding Layer 2 Protocol Tunneling on page 1910](#)
 - [Configuring Layer 2 Protocol Tunneling on page 2063](#)

Troubleshooting Private VLANs

Use the following information to troubleshoot a private VLAN configuration.

- [Limitations of Private VLANs on page 2263](#)
- [Forwarding with Private VLANs on page 2263](#)
- [Egress Firewall Filters with Private VLANs on page 2264](#)
- [Egress Port Mirroring with Private VLANs on page 2265](#)

Limitations of Private VLANs

The following constraints apply to private VLAN configurations:

- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.

Forwarding with Private VLANs

Problem	Description:
	<ul style="list-style-type: none"> • When isolated VLAN or community VLAN tagged traffic is received on a PVLAN trunk port, MAC addresses are learned from the primary VLAN. This means that output from the <code>show ethernet-switching table</code> command shows that MAC addresses are learned from the primary VLAN and replicated to secondary VLANs. This behavior has no effect on forwarding decisions. • If a packet with a secondary VLAN tag is received on a promiscuous port, it is accepted and forwarded. • If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped. <ul style="list-style-type: none"> • The packet has a community VLAN tag. • The packet is destined to a unicast MAC address or multicast group MAC address that was learned on an isolated VLAN. • If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped. <ul style="list-style-type: none"> • The packet has an isolated VLAN tag. • The packet is destined to a unicast MAC address or multicast group MAC address that was learned on a community VLAN. • If a packet with a primary VLAN tag is received by a secondary (isolated or community) VLAN port, the secondary port forwards the packet. • If you configure a community VLAN on one device and configure another community VLAN on a second device and both community VLANs use the same VLAN ID, traffic for one of the VLANs can be forwarded to the other VLAN. For example, assume the following configuration:

- Community VLAN comm1 on switch 1 has VLAN ID 50 and is a member of primary VLAN pvlan100.
- Community VLAN comm2 on switch 2 also has VLAN ID 50 and is a member of primary VLAN pvlan200.
- Primary VLAN pvlan100 exists on both switches.

If traffic for comm1 is sent from switch 1 to switch 2, it will be sent to the ports participating in comm2. (The traffic will also be forwarded to the ports in comm1, as you would expect.)

Solution These are expected behaviors.

Egress Firewall Filters with Private VLANs

Problem **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).

- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

Solution These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

Egress Port Mirroring with Private VLANs

Problem **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

Solution This is expected behavior.

- Related Documentation**
- [Understanding Private VLANs on page 1878](#)
 - [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887](#)
 - [Creating a Private VLAN on a Single Switch on page 2057](#)
 - [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
 - [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024](#)

Troubleshooting Q-in-Q and VLAN Translation Configuration

- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 2266](#)
- [Egress Port Mirroring with VLAN Translation on page 2266](#)

Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

Problem **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

Solution This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

Egress Port Mirroring with VLAN Translation

Problem **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

Solution This is expected behavior.

Related Documentation

- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
- [Example: Setting Up Q-in-Q Tunneling on page 2036](#)

PART 9

High Availability

- [Overview on page 2269](#)
- [Configuration on page 2281](#)
- [Administration on page 2343](#)
- [Troubleshooting on page 2411](#)

CHAPTER 24

Overview

- [Software Features Overview on page 2269](#)

Software Features Overview

- [Graceful Restart Concepts on page 2269](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 2270](#)
- [Understanding VRRP on page 2274](#)
- [Understanding VRRP Between QFabric Systems on page 2277](#)

Graceful Restart Concepts

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC).
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits

of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

**Related
Documentation**

- *Understanding High Availability Features on Juniper Networks Routers*
- *Graceful Restart System Requirements*
- *Aggregate and Static Routes*
- *Graceful Restart and Routing Protocols*
- *Graceful Restart and MPLS-Related Protocols*
- *Graceful Restart and Layer 2 and Layer 3 VPNs*
- *Graceful Restart on Logical Systems*
- *Example: Configuring Graceful Restart*
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)

Understanding Nonstop Software Upgrade for QFabric Systems

The framework that underlies a nonstop software upgrade in a QFabric system enables you to upgrade the system in a step-by-step manner and minimize the impact to the continuous operation of the system. This topic explains how a nonstop software upgrade works in a QFabric system, the steps that are involved, and the procedures that you need to implement to experience the benefits of this style of software upgrade.

Nonstop software upgrade enables some QFabric system components to continue operating while similar components in the system are being upgraded. In general, the QFabric system upgrades redundant components in stages so that some components remain operational and continue forwarding traffic while their equivalent counterparts upgrade to a new version of software.



TIP: Use the following guidelines to decide when to implement a nonstop software upgrade:

- If you need to upgrade all components of the system in the shortest amount of time (approximately one hour) and you do not need to retain the forwarding resiliency of the data plane, issue the **request system software add component all** command to perform a standard software upgrade. All components of the QFabric system upgrade simultaneously and expediently, but this type of upgrade does not provide resiliency or switchover capabilities.

- If you need to minimize service impact, preserve the forwarding operations of the data plane during the upgrade, and are willing to take the extra time required for component switchovers (in many cases, several hours), issue the three nonstop software upgrade commands (`request system software nonstop-upgrade` `director-group` | `fabric` | `node-group`) described in this topic in the correct order.



NOTE:

- Before you begin a nonstop software upgrade, issue the `request system software download` command to copy the software to the QFabric system.
- Each of the 3 nonstop software upgrade steps must be considered parts of the whole process. You must complete all 3 steps of a nonstop software upgrade in the correct order to ensure the proper operation of the QFabric system.
- Open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected. For more information on verification of the upgrade, see [“Verifying Nonstop Software Upgrade for QFabric Systems” on page 1410](#).
- Issue the `show fabric administration inventory` command to verify that all upgraded components are operational at the end of a step before beginning the next step.
- Once you start the nonstop software upgrade process, we strongly recommend that you complete all 3 steps within 12 hours.

The three steps to a successful nonstop software upgrade must be performed in the following order:

- Director group—The first step upgrades the Director devices, the fabric manager Routing Engine, and the diagnostic Routing Engine. To perform the first step, issue the **`request system software nonstop-upgrade director-group`** command. The key actions that occur during a Director group upgrade are:
 1. Connecting to the QFabric system by way of an SSH connection. This action establishes a load-balanced CLI session on one of the Director devices in the Director group.
 2. The QFabric system downloads and installs the new software in both Director devices.
 3. The Director device hosting the CLI session becomes the master for all QFabric system processes running on the Director group, such as the fabric manager and network Node group Routing Engines.
 4. The QFabric system installs the new software for the backup fabric manager Routing Engine on the backup Director device.

5. The backup Director device reboots to activate the new software.
6. The master Director device begins a 15 minute sequence that includes a temporary suspension of QFabric services and a QFabric database transfer. You cannot issue operational mode commands in the QFabric CLI during this period.
7. The QFabric system installs the new software for the fabric manager and diagnostic Routing Engines on the Director group master.
8. The QFabric system switches mastership of all QFabric processes from the master Director device to the backup Director device.
9. The master Director device reboots to activate the new software.
10. The CLI session terminates, and logging back in to the QFabric system with a new SSH connection establishes the session on the new master Director device (the original backup).
11. The previous master Director device resumes operation as a backup and the associated processes (such as the fabric manager and network Node group Routing Engines) become backup as well. The fabric control Routing Engine associated with this Director device returns to active status.



NOTE: After the Director group nonstop software upgrade completes, any Interconnect device or Node device that reboots will automatically download the new software, install it, and reboot again. As a result, try not to restart any QFabric system devices before you complete the rest of the nonstop software upgrade steps.



TIP:

- To enable BGP and OSPF to continue operating on the network Node group during a Director group nonstop service upgrade, we recommend that you configure graceful restart for these routing protocols. For more information on graceful restart, see [“Configuring Graceful Restart for QFabric Systems” on page 1375](#).
 - Wait 15 minutes after the second Director device returns to service and hosts Routing Engine processes before proceeding to step 2—the fabric upgrade. You can verify the operational status of both Director devices by issuing the `show fabric administration inventory director-group status` command. Also, issue the `show fabric administration inventory infrastructure` command to verify when the Routing Engine processes become load balanced (typically, there will be three to four Routing Engines running on each Director device).
-
- Fabric—The second step upgrades the Interconnect devices and the fabric control Routing Engines. To perform the second step, issue the **request system software nonstop-upgrade fabric** command. The key actions that occur during a fabric upgrade are:

1. The QFabric system downloads, validates, and installs the new software in all Interconnect devices and fabric control Routing Engines (FC-0 and FC-1).
2. One fabric control Routing Engine reboots and comes back online.
3. The other fabric control Routing Engine reboots and comes back online.
4. The first Interconnect device reboots, comes back online, and resumes the forwarding of traffic.
5. Subsequent Interconnect devices reboot one at a time, come back online, and return to service.

**NOTE:**

- If the software does not load properly on any one of the fabric components, all components revert back to the original software version.
- If one of the components in a fabric upgrade does not reboot successfully, issue the **request system reboot fabric** command to reattempt the rebooting process for this fabric component and activate the new software.

- **Node group**—The third and final step upgrades Node groups. You can choose to upgrade a network Node group, a redundant server Node group, or individual server Node groups. You can upgrade the Node groups one at a time or in groups (known as upgrade groups). However, you must upgrade all Node groups in your QFabric system before you can complete the nonstop software upgrade process. To perform the third step, issue the **request system software nonstop-upgrade node-group** command.

The key actions that occur during a network Node group upgrade are:

1. The QFabric system copies the new software to each Node device one at a time.
2. The QFabric system validates and then installs the new software in all Node devices simultaneously.
3. The system copies the software to the network Node group Routing Engines.
4. The QFabric system validates and then installs the software in the network Node group Routing Engines one at a time -- first the backup, then the master.
5. The backup network Node group Routing Engine reboots and comes back online.
6. The supporting Node devices reboot and come back online one at a time.



NOTE: To reduce the total upgrade duration, configure an upgrade group. All Node devices within the upgrade group reboot at the same time.

7. The master network Node group Routing Engine relinquishes mastership to the backup, reboots, and comes back online.

The key actions that occur during a redundant server Node group upgrade are:

1. The QFabric system copies the new software to the backup Node device, then the master Node device.
2. The QFabric system validates and then installs the new software on the backup Node device, then the master Node device.
3. The backup Node device reboots, comes back online, and becomes the master Node device.
4. The previous master Node device reboots and comes back online as a backup Node device.



NOTE: For redundant server Node groups, both Node devices must be online before the upgrade will proceed. If one of the devices is no longer available, remove the Node device from the Node group configuration before you issue the nonstop software upgrade command.

The key actions that occur during a server Node group upgrade for a Node group that contains one member are:

1. The Node device downloads the software package and validates the software.
2. The Node device installs the software and reboots.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [request system software add on page 394](#)
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)

Understanding VRRP

The Juniper Networks QFX Series supports the Virtual Router Redundancy Protocol (VRRP). This topic covers:

- [Overview of VRRP on page 2274](#)
- [Sample VRRP Topology on page 2275](#)

Overview of VRRP

Configuring end hosts on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the end hosts. When hosts are configured with static routes, the failure of the default gateway normally results in a

catastrophic event, isolating all hosts that are unable to detect available alternate paths to their gateway. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for end hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Switches configured with VRRP share a virtual IP address, which is the address you configure as the default route on the hosts. At any time, one of the switches is the VRRP master, meaning that it owns the virtual IP address and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure (**1 through 255**). If the master fails, the backup switch with the highest priority becomes the master within a few seconds. This is done without any interaction with the hosts.

In VRRP operation, the master sends advertisements to the backup switches at regular intervals. The default interval is 1 second. If the backup switches do not receive an advertisement for a set period, the backup with the highest priority takes over as master within a few seconds and begins forwarding packets. This is done without any interaction with the hosts.

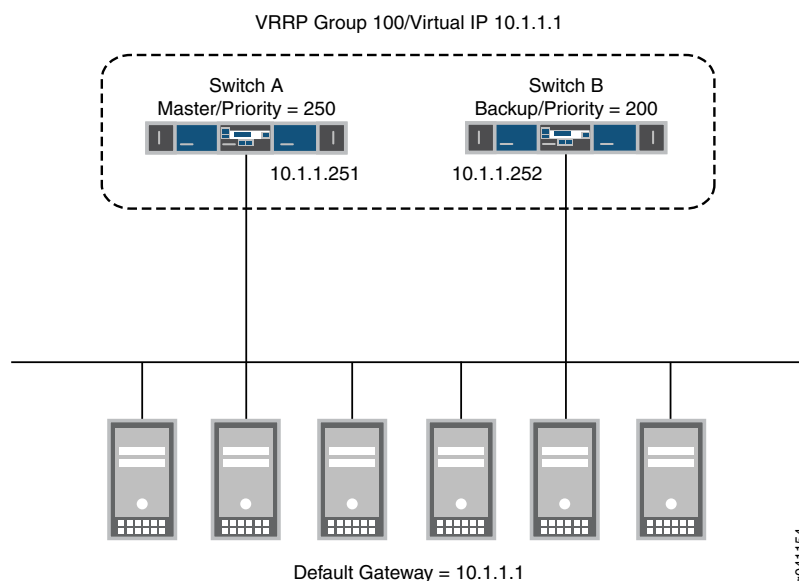


NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. One benefit of this configuration is if you use VMware's vMotion, virtual machines can transition between hosts connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a host connected to a QFabric system in data center A can transition to a host connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address.

Sample VRRP Topology

[Figure 61 on page 2276](#) illustrates a basic VRRP topology. In this example, switches A and B are running VRRP and share the virtual IP address 10.1.1.1. The default gateway for each of the clients is 10.1.1.1.

Figure 61: Basic VRRP Topology

The following illustrates basic VRRP behavior using [Figure 61 on page 2276](#) for reference:

1. When any of the servers wants to send traffic out of the LAN, it sends the traffic to the default gateway address of 10.1.1.1. This is a virtual IP address (VIP) owned by VRRP group 100. Because switch A is the master of the group, the VIP is associated with the “real” address 10.1.1.251 on switch A, and traffic from the servers is actually sent to this address. (Switch A is the master because it has been configured with a higher priority value.)
2. If there is a failure on switch A that prevents it from forwarding traffic to or from the servers—for example, if the interface connected to the LAN fails—switch B becomes the master and assumes ownership of the VIP. The servers continue to send traffic to the VIP, but because the VIP is now associated with the “real” address 10.1.1.252 on switch B (because of change of master), the traffic is sent to switch B instead of switch A.
3. If the problem that caused the failure on switch A is corrected, switch A becomes the master again and reasserts ownership of the VIP. In this case, the servers resume sending traffic to switch A.

Notice that no configuration changes are required on the servers for them to switch between sending traffic to switch A and switch B. When the VIP moves between 10.1.1.251 and 10.1.1.252, the change is detected by normal TCP-IP behavior and no configuration or intervention is required on the servers.

Related Documentation

- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)
- [Understanding VRRP Between QFabric Systems on page 2277](#)

Understanding VRRP Between QFabric Systems

Juniper Networks QFabric systems support the Virtual Router Redundancy Protocol (VRRP). This topic covers:

- [VRRP Differences on QFabric Systems on page 2277](#)
- [Configuration Details on page 2277](#)

VRRP Differences on QFabric Systems

Configuring servers on your network with static routes to a default gateway minimizes configuration effort and complexity and reduces processing overhead. However, a failure of the default gateway normally results in a catastrophic event, isolating the servers. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for servers if the primary gateway fails.

Switches configured with VRRP share a virtual IP (VIP) address, which is the address you configure as the default route on the servers. In normal VRRP operation, one of the switches is the VRRP master, meaning that it owns the VIP and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure. If the master fails, the backup switch with the highest priority becomes the master and takes ownership of the VIP within a few seconds. This is done without any interaction with the servers.

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. However, in normal VRRP operation, only one system can be the master for a given VRRP group at any one time, which means that only one system can act as a default gateway using the VIP configured for the group. When running VRRP over two QFabric systems, you might want both systems to simultaneously use the VIP to act as a gateway and forward traffic. To achieve this, you can configure a firewall filter to block the VRRP advertisement packets between the QFabric systems on the link between the two network Node groups. When you do this, both QFabric systems act as master and forward traffic received by the VIP (which is the default gateway address that you configure on servers connected to both QFabric systems). If you use VMware's vMotion, this configuration allows virtual machines to transition between servers connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a server connected to a QFabric system in data center A can transition to a server connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address because both QFabric systems use the same VIP.



NOTE: To use a firewall filter to block VRRP traffic, create a firewall term that matches traffic for protocol `vrrp` and discards that traffic.

Configuration Details

Configuring a VRRP group across two QFabric systems is similar to configuring VRRP on two switches. The main differences are listed here:

- All the interfaces in both QFabric systems that participate in VRRP must be members of the same VLAN.
- You must create routed VLAN interfaces (RVIs) in that VLAN on both QFabric systems.
- The IP addresses that you assign to both RVIs must be in the same subnet.
- You must configure VRRP on the RVIs.
- Both RVIs must be members of the same VRRP group. This is what allows the two QFabric systems to share a virtual IP address.

The following tables list the elements of an example VRRP configuration running on two QFabric systems—QFabric system A and QFabric system B. This example is configured so that both QFabric systems act as the VRRP master for VIP 10.1.1.50/24 and assumes that a firewall filter blocks the VRRP advertisements between the systems.

[Table 218 on page 2278](#) lists the required characteristics of the RVIs in the example configuration.



NOTE: Most of the configuration settings in the following tables would also apply in a traditional VRRP configuration. However, the advertisement interval and priority settings would need to be different (as noted).

Table 218: RVIs on QFabric systems in example VRRP configuration

RVI on QFabric System A	RVI on QFabric System B
vlan.100	vlan.200
Member of VLAN 100. (Note that the VLAN is the same on both QFabric systems.)	Member of VLAN 100
IP address 10.1.1.100/24	IP address 10.1.1.200/24
Member of VRRP group 500	Member of VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24

You must configure VRRP on the RVIs on both QFabric systems. [Table 219 on page 2278](#) lists the elements of a sample VRRP configuration on each RVI. Note that with the exception of the priority, the parameters *must* be the same on both systems.

Table 219: Sample VRRP configuration each RVI

VRRP on RVI on QFabric System A	VRRP on RVI on QFabric System B
VRRP group 500	VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24

Table 219: Sample VRRP configuration each RVI (*continued*)

Advertisement interval 60 seconds. (In a normal VRRP configuration, you would set this interval to be much smaller, such as 1 second. However, in this configuration these packets are blocked by the firewall filter on the interface that connects to QFabric system B, so there is no need to send them frequently.)	Advertisement interval 60 seconds
Authentication type md5	Authentication type md5
Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8	Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8
Priority 254. (In a normal VRRP configuration, this value would be different on the two systems and the system with the higher value would be the master. However, in this configuration both systems are acting as master, so you do not have to configure different values.)	Priority 254



NOTE: Priority 255 is not supported for RVIs.

Table 220 on page 2279 lists all the interfaces on QFabric system A in the example configuration and identifies what they connect to.

Table 220: Interfaces on QFabric system A. All interfaces are members of VLAN 100.

VLAN 100 Interfaces on QFabric System A	Connects To
vlan.100	vlan.200
Network Node group interface QFA-NNG:xe-0/0/0	QFB-NNG:xe-0/0/0 on QFabric system B
Network Node group interface QFA-NNG:xe-0/0/1	Redundant server Node group interface QFA-RSNG:xe-0/0/0
Redundant server Node group interface QFA-RSNG:xe-0/0/0	Connects to a network Node group interface QFA-NNG:xe-0/0/1
Redundant server Node group interface QFA-RSNG:xe-0/0/1	LAN with servers running virtual machines

Table 221 on page 2279 lists all the interfaces on QFabric system B in the example configuration and identifies what they connect to.

Table 221: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).

VLAN 100 Interfaces on QFabric System B	Connects To
vlan.200	vlan.100
Network Node group interface QFB-NNG:xe-0/0/0	QFA-NNG:xe-0/0/0 on QFabric system A
Network Node group interface QFB-NNG:xe-0/0/1	Redundant server Node group interface QFB-RSNG:xe-0/0/0

Table 221: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A). (*continued*)

Redundant server Node group interface QFB-RSNG:xe-0/0/0	Connects to a network Node group interface QFB-NNG:xe-0/0/1
Redundant server Node group interface QFB-RSNG:xe-0/0/1	LAN with servers running virtual machines

**Related
Documentation**

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuration

- [Configuration Tasks for Graceful Restart on page 2281](#)
- [Configuration Example for VRRP on page 2291](#)
- [Configuration Tasks for VRRP on page 2296](#)
- [Configuration Statements for Graceful Restart on page 2306](#)
- [Configuration Statements for VRRP on page 2318](#)

Configuration Tasks for Graceful Restart

- [Configuring Graceful Restart for QFabric Systems on page 2281](#)
- [Configuring Routing Protocols Graceful Restart on page 2285](#)

Configuring Graceful Restart for QFabric Systems

When you configure graceful restart in the QFabric CLI, the QFabric system applies the configuration to the network Node group to participate in graceful restart operations with devices external to the QFabric system. Such configuration preserves routing table state and helps neighboring routing devices to resume routing operations more quickly after a system restart. This also enables the network Node group to resume routing operations rapidly if there is a restart in the QFabric system (such as a software upgrade). As a result, we recommend enabling graceful restart for routing protocols in the QFabric CLI.



NOTE: The QFabric system also uses graceful restart internally within the fabric to facilitate interfabric resiliency and recovery. This internal feature is enabled by default with no configuration required.

- [Enabling Graceful Restart on page 2282](#)
- [Configuring Graceful Restart Options for BGP on page 2282](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2283](#)
- [Tracking Graceful Restart Events on page 2284](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
  graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]  
routing-options {  
  graceful-restart {  
    disable;  
    restart-duration seconds;  
  }  
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]  
protocols {  
  bgp {  
    graceful-restart {  
      disable;  
    }  
  }  
}
```

```

        restart-time seconds;
        stale-routes-time seconds;
    }
}
routing-options {
    graceful-restart;
}

```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group group-name graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group group-name neighbor ip-address graceful-restart]** hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

```

[edit]
protocols {
    ospf | ospfv3 {
        graceful-restart {
            disable;
            helper-disable
            no-strict-lsa-checking;
            notify-duration seconds;
            restart-duration seconds;
        }
    }
}
routing-options {
    graceful-restart;
}

```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }
```

To reenable the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 1378](#).



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols protocol traceoptions flag]** hierarchy level:

```
[edit protocols]
  isis {
    traceoptions {
      flag graceful-restart;
    }
  }
```



```

}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}

```

- Related Documentation**
- [Graceful Restart Concepts on page 2269](#)
 - [Verifying Graceful Restart Operation on page 2343](#)

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 2285](#)
- [Configuring Graceful Restart Options for BGP on page 2286](#)
- [Configuring Graceful Restart Options for ES-IS on page 2287](#)
- [Configuring Graceful Restart Options for IS-IS on page 2287](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2288](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 2289](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 2289](#)
- [Tracking Graceful Restart Events on page 2291](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```

routing-options {
  graceful-restart;
}

```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```

[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}

```

```
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group group-name graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group group-name neighbor ip-address graceful-restart]** hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

```
[edit]
protocols {
  esis {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable ES-IS graceful restart capability, include the **disable** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols isis graceful-restart]** hierarchy level.

```
[edit]
protocols {
  isis {
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



NOTE: If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



NOTE: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols isis]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 1378](#).

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
graceful-restart {
  helper-disable <both | restart-signaling | standard>
}
```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.

**NOTE:**

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospf3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 1378](#).



NOTE: You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Configuring Graceful Restart Options for RIP and RIPng

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  (rip | ripng) {
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for PIM Sparse Mode

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing

platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
  pim {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols *protocol* traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

Related Documentation

- [Graceful Restart Concepts on page 2269](#)
- [Graceful Restart System Requirements](#)
- [Graceful Restart and Routing Protocols](#)
- [Verifying Graceful Restart Operation on page 2343](#)
- [Example: Configuring Graceful Restart](#)

Configuration Example for VRRP

- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Example: Configuring VRRP for Load Sharing

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 2292](#)
- [Overview and Topology on page 2292](#)
- [Configuring VRRP on Both Switches on page 2293](#)
- [Verification on page 2295](#)

Requirements

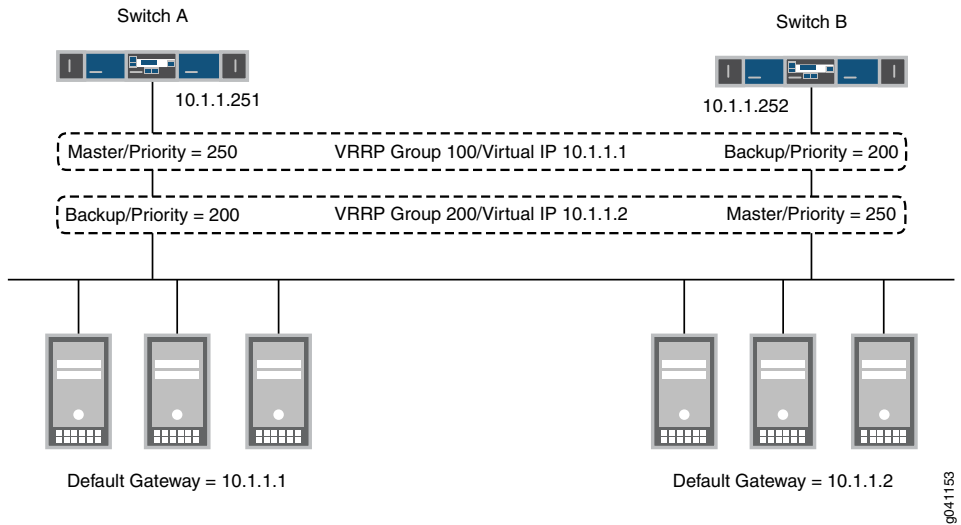
This example uses the following hardware and software components:

- Two QFX3500 switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the switches fails, the other switch takes over for it. In the topology shown in [Figure 62 on page 2292](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

Figure 62: VRRP Load-Sharing Configuration



This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 222 on page 2292](#) lists VRRP settings for each switch.

Table 222: Settings for VRRP Load-Sharing Example

Switch A	Switch B
VRRP Group 100:	VRRP Group 100:
<ul style="list-style-type: none">• Interface address: 10.1.1.251• VIP: 10.1.1.1• Priority: 250	<ul style="list-style-type: none">• Interface address: 10.1.1.252• VIP: 10.1.1.1• Priority: 200

Table 222: Settings for VRRP Load-Sharing Example (*continued*)

Switch A	Switch B
VRRP Group 200: <ul style="list-style-type: none"> Interface address: 10.1.1.251 VIP: 10.1.1.2 Priority: 200 	VRRP Group 200: <ul style="list-style-type: none"> Interface address: 10.1.1.252 VIP: 10.1.1.2 Priority: 250

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

Configuring VRRP on Both Switches

CLI Quick Configuration

Enter the following on Switch A:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
```

Enter the following on Switch B:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch A:

- Create VRRP group 100 on Switch A and configure the virtual IP address for the group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```
- Assign the VRRP priority for this interface in this group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```
- Create VRRP group 200 on Switch A and configure the virtual IP address for the group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```
- Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 200
```

Step-by-Step Procedure Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

Results Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.251 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 250
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 200
          }
        }
      }
    }
  }
}
```

Display the results of the configuration on Switch B:

```

user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.252 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 200
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 250
          }
        }
      }
    }
  }
}

```

Verification

- [Verifying that VRRP is Working on Switch A on page 2295](#)
- [Verifying that VRRP is Working on Switch B on page 2295](#)

Verifying that VRRP is Working on Switch A

Purpose Verify that VRRP is active on Switch A and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```

user@switch> show vrrp

```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	master	A .0327 1c1 10.1.1.251 vip 10.1.1.1	
xe-0/0/0.0	up	200	backup	A .0327 1c1 10.1.1.251 vip 10.1.1.2	

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `1c1` address is the physical address of the interface and the `vip` address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

Verifying that VRRP is Working on Switch B

Purpose Verify that VRRP is active on Switch B and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

Interface	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	backup	A .0327 lcl 10.1.1.252 vip 10.1.1.1	
xe-0/0/0.0	up	200	master	A .0327 lcl 10.1.1.252 vip 10.1.1.2	

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `lcl` address is the physical address of the interface and the `vip` address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)

Configuration Tasks for VRRP

- [Configuring Basic VRRP Support on page 2296](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2297](#)
- [Configuring the Startup Period for VRRP Operations on page 2298](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 2299](#)
- [Configuring VRRP Preemption and Hold Time on page 2300](#)
- [Configuring a Route to Be Tracked on page 2301](#)
- [Configuring a Logical Interface to Be Tracked on page 2302](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 2304](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 2304](#)
- [Configuring the Silent Period on page 2305](#)
- [Configuring Inheritance for a VRRP Group on page 2305](#)

Configuring Basic VRRP Support

To configure basic VRRP support, configure VRRP groups on interfaces by including the `vrrp-group` statement:

```
vrrp-group group-id {
    priority number;
    virtual-address [ addresses ];
}
```

An interface can be a member of multiple VRRP groups.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

For each interface, you must configure the following:

- Group identifier—Assign a value from 0 through 255. You must use the same identifier for each switch in the VRRP group.
- Priority—Assign a value from 1 through 255. The switch with the highest priority becomes the VRRP master. Assign different priorities to each switch in the VRRP group. If there are two or more switches with the same priority, the switch with the VRRP interface that has the highest IP address becomes the master.
- Virtual IP address—Normally, you configure only one address per group, but you can configure as many as eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
 - You must configure the same address on all the switches in the VRRP group.
 - If you configure a virtual IP address to be the same as a physical interface address, the switch with that interface becomes the master for the group. You must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
 - If the virtual IP address is not the same as the physical interface address, you must ensure that the address does not appear anywhere else in the switch configuration. For example, verify that you do not use this address for another interface (including an aggregated Ethernet interface) or for a static ARP entry.



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement at the **[edit interfaces *interface-name*]** hierarchy. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring the Startup Period for VRRP Operations on page 2298](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2297](#)

Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted switches participate in a VRRP group. By default, VRRP authentication is disabled. You can configure one of the following authentication methods for a group, and each switch in the same group must use the same method:

- Simple authentication—Uses a text password included in the transmitted packet. The receiving switch uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Adds an authentication header (AH) to the IP packet that encapsulates the VRRP packet. You create an authentication key that is used to create a hash of the packet, and the hash is stored in the AH. A receiving switch recalculates the hash on the incoming packet and compares the hashes. If they are identical, the packet is valid and is accepted. Otherwise the switch drops the incoming packet.

To enable authentication and specify an authentication method, include the **authentication-type** statement.

authentication-type *authentication*;

authentication can be **simple** or **md5**. The authentication type must be the same for all the switches in the VRRP group.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

authentication-key *key*;

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").



NOTE: The key must be the same for all switches in the VRRP group.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

**Related
Documentation**

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)

Configuring the Startup Period for VRRP Operations

Configure the startup-silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets while an interface is coming online. The period starts when the state of a VRRP interface is changed from down to up. During this period, Master Down Events are ignored.

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring the Advertisement Interval for the VRRP Master

By default, the master switch sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master switch is still operational. If the master switch fails or becomes unreachable, the backup switch with the highest priority value becomes the new master switch.

You can modify the advertisement interval in seconds or in milliseconds; the interval must be the same for all routing platforms in the VRRP group.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 2299](#)
- [Modifying the Advertisement Interval in Milliseconds on page 2299](#)

Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**

Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

```
fast-interval milliseconds;
```

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]



NOTE: Junos OS sets the advertisement interval to 0 in VRRP packets. When you configure VRRP with other vendors' equipment, the *fast-interval* statement works correctly only when the other equipment also has the advertisement interval set to 0 in the VRRP packet. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

**Related
Documentation**

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 2300](#)
- [Configuring the Preemption Hold Time on page 2300](#)
- [Overriding the Hold Time on page 2301](#)

Configuring VRRP Preemption

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

```
preempt;
```

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

```
no-preempt;
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address vrrp-group *group-id*] preempt**

Overriding the Hold Time

You can use the **asymmetric-hold-time** statement to configure a VRRP master to fail over to the backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.

When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.

You can include this statement at the following hierarchy level:

- **[edit protocols vrrp]**

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring a Route to Be Tracked

A VRRP master can track a route and dynamically trigger a new master router election if the route becomes unreachable. To enable this behavior, you must configure a cost that will be subtracted from the priority of the master if the tracked route becomes unreachable. The new priority must be less than the priority of one of the backups so that the backup becomes the new master.

To configure a route to be tracked, include the following statements:

```
track {
  priority-hold-time seconds;
  route prefix/prefix-length routing-instance default priority-cost priority;
}
```

You can include these statements at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**

The **prefix** and **prefix-length** values specify the route to be tracked. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority of the master changes because of a tracking event, the priority hold timer

begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **priority-cost** option is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through **254**. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).

**Related
Documentation**

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)
- [Configuring a Logical Interface to Be Tracked on page 2302](#)

Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can change the priority of the switch based on the state of the interface, which might trigger a new master election. VRRP can also track the operational speed of a logical interface and update the priority of the switch when the speed crosses a configured threshold. For each VRRP group, you can track as many as 10 logical interfaces.

When interface tracking is enabled on a switch, you cannot assign the switch a priority of 255 to make it the master for the group.

To configure a logical interface to be tracked, include the following statements:

```
track {  
  interface interface-name {  
    bandwidth-threshold bits-per-second priority-cost priority;  
    priority-cost priority;  
  }  
  priority-hold-time seconds;  
}
```

You can include these statements at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**

The interface specified is the interface to be tracked for the VRRP group. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **bandwidth-threshold** statement specifies a threshold for the tracked interface. If the bandwidth of the tracked interface drops below the threshold value, the system subtracts

the bandwidth threshold **priority-cost** value from the VRRP priority for the switch. You can create as many as five **bandwidth-threshold** statements for each tracked interface.

The interface **priority-cost** statement is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).



WARNING: On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

If you configure tracking for more than one interface, Junos OS subtracts the sum of the priority costs for the tracked interfaces from the VRRP priority if all the tracked interfaces fail. However, if you configure the interface **priority-cost** statement and the bandwidth threshold **priority-cost** statement, they are not added together. The switch uses only one priority cost for a tracked interface, as indicated in [Table 223 on page 2303](#):

Table 223: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	priority cost <i>priority</i>
Not down; media speed below one or more bandwidth thresholds	Priority-cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you do not configure any bandwidth thresholds. If you do not configure an interface **priority-cost** value and the interface fails, Junos OS subtracts the bandwidth threshold **priority-cost** value of the lowest bandwidth threshold from the priority of the switch.

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)
- [Configuring a Route to Be Tracked on page 2301](#)

Configuring a Backup to Accept Packets Destined for the Virtual IP Address

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

```
accept-data;
```

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group] *group-id***

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

```
no-accept-data;
```

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

Related Documentation

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring Passive ARP Learning for VRRP Backups

By default, VRRP backup switches drop ARP requests for the MAC address of the VRRP IP. This means that backups do not learn the ARP mappings (IP address to MAC address mappings) for the hosts sending the requests. If it becomes the master, the configured backup must learn all the entries that were present in the ARP cache of the original master. In environments with many directly attached hosts, the number of ARP entries to learn can be very large. This can cause a significant delay while the backup transitions to the master state, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup to learn approximately the same contents as the ARP cache in the master, thus preventing the problem of needing to learn many ARP entries quickly. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
```

`passive-learning;`

We recommend setting passive learning on both the backup and master VRRP switches. Doing so prevents the need to manually configure a master that fails and becomes a backup. While a switch operates as the master, the passive learning configuration has no impact. The configuration takes effect only when a switch operates as a backup.

- Related Documentation**
- [Understanding VRRP on page 2274](#)
 - [Configuring Basic VRRP Support on page 2296](#)
 - [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring the Silent Period

When the state of a VRRP interface changes from down to up, a silent period begins. During this period, any master down events are ignored. Configure the silent period interval to avoid problems that can be caused if incoming VRRP advertisement packets are delayed or interrupted while an interface starts up.

To configure the silent period, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```

- Related Documentation**
- [Understanding VRRP on page 2274](#)
 - [Configuring Basic VRRP Support on page 2296](#)
 - [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuring Inheritance for a VRRP Group

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]**:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
  vrrp-group group-id]
  vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, note the following conditions:

- Both inheriting groups and active groups must be on the same physical interface and logical system. However, the groups need not necessarily be on the same VLAN or logical interface.
- Both inheriting groups and active groups must be on the same routing instances; however, this limitation does not apply for groups on the integrated routing and bridging (IRB) interfaces.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

**Related
Documentation**

- [Understanding VRRP on page 2274](#)
- [Configuring Basic VRRP Support on page 2296](#)
- [Example: Configuring VRRP for Load Sharing on page 2291](#)

Configuration Statements for Graceful Restart

- [disable on page 2307](#)
- [disable \(BGP Graceful Restart\) on page 2308](#)
- [graceful-restart \(Enabling Globally\) on page 2309](#)
- [graceful-restart \(Protocols BGP\) on page 2310](#)
- [graceful-restart \(Protocols OSPF\) on page 2311](#)
- [helper-disable \(OSPF\) on page 2313](#)
- [no-strict-lsa-checking on page 2314](#)
- [notify-duration on page 2315](#)
- [restart-duration on page 2316](#)

- [restart-time \(BGP Graceful Restart\) on page 2317](#)
- [stale-routes-time on page 2318](#)

disable

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp isis ldp ospf ospf3 pim rip ripng rsvp) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (bgp isis isis ospf ospf3 ldp pim rip ripng rsvp) graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Disable graceful restart.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Graceful Restart • Configuring Routing Protocols Graceful Restart on page 2285 • Configuring Graceful Restart for MPLS-Related Protocols • Configuring VPN Graceful Restart • Configuring Logical System Graceful Restart • Graceful Restart Configuration Statements • Configuring Graceful Restart for QFabric Systems on page 1375

disable (BGP Graceful Restart)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart], [edit protocols bgp graceful-restart], [edit protocols bgp group <i>group-name</i> graceful-restart], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.




NOTE: When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• graceful-restart on page 1388

graceful-restart (Enabling Globally)

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.</p> <p>For VPNs, the graceful-restart statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</p> <p>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p>
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Graceful Restart</i> • Configuring Routing Protocols Graceful Restart on page 2285 • <i>Configuring Graceful Restart for MPLS-Related Protocols</i> • <i>Configuring VPN Graceful Restart</i> • <i>Configuring Logical System Graceful Restart</i> • <i>Graceful Restart Configuration Statements</i> • Configuring Graceful Restart for QFabric Systems on page 1375

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time seconds; stale-routes-time seconds; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2285 • Configuring Graceful Restart for QFabric Systems on page 1375 • <i>Junos OS High Availability Configuration Guide</i>

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....
notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds


Default: 180 seconds

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation

- [Example: Configuring Graceful Restart for OSPF on page 3894](#)
- [Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 3898](#)
- [Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 3901](#)
- [Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 3904](#)
- [Configuring Graceful Restart for QFabric Systems on page 1375](#)
- *Junos OS High Availability Configuration Guide*

helper-disable (OSPF)

Syntax	<code>helper-disable < both restart-signaling standard >;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart], [edit protocols ospf graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Options both , restart-signaling , and standard introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart. The last committed statement takes precedence over the previously configured statement.
Default	Helper mode is enabled by default for OSPF.
Options	both —(Optional) Disable helper mode for both standard and restart signaling-based graceful restart. restart-signaling —(Optional) Disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813).
<div>  <p>NOTE: Restart signaling-based helper mode is not supported for OSPFv3 configurations.</p> </div>	
	standard —(Optional) Disable helper mode for standard graceful restart (based on RFC 3623).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Protocols Graceful Restart on page 2285 • Configuring Graceful Restart for MPLS-Related Protocols • Configuring Graceful Restart for QFabric Systems on page 1375

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router.
Default	By default, LSA checking is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• <i>maximum-neighbor-recovery-time</i>• <i>recovery-time</i>

notify-duration

Syntax	<code>notify-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the length of time the router notifies helper OSPF routers that it has completed graceful restart.
Options	<p><i>seconds</i>—Length of time in the router notifies helper OSPF routers that it has completed graceful restart.</p> <p>Range: 1 through 3600</p> <p>Default: 30</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2285 • Configuring Graceful Restart for QFabric Systems on page 1375 • restart-duration on page 2316

restart-duration

Syntax	<code>restart-duration seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-options graceful-restart]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>
Options	<p>seconds—Time for the graceful restart period.</p> <p>Range:</p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—120 through 900• ES-IS—30 through 300• IS-IS—30 through 300• OSPF/OSPFv3—1 through 3600• PIM—30 through 300 <p>Default:</p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—300• ES-IS—180• IS-IS—210• OSPF/OSPFv3—180• PIM—60
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Enabling Graceful Restart*
 - [Configuring Routing Protocols Graceful Restart on page 2285](#)
 - *Configuring Graceful Restart for MPLS-Related Protocols*
 - *Configuring VPN Graceful Restart*
 - *Configuring Graceful Restart for VPNs*
 - *Configuring Logical System Graceful Restart*
 - *Graceful Restart Configuration Statements*
 - [Configuring Graceful Restart for QFabric Systems on page 1375](#)

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time seconds;</code>
Hierarchy Level	<p>[edit protocols (bgp rip ripng) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart (Enabling Globally)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	<p>seconds—Length of time for the graceful restart period.</p> <p>Range: 1 through 600 seconds</p> <p>Default: Varies by protocol:</p> <ul style="list-style-type: none"> • BGP—120 seconds • RIP and RIPng—60 seconds
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2285 • Configuring Graceful Restart Options for RIP and RIPng on page 2289 • Configuring Graceful Restart for QFabric Systems on page 1375 • stale-routes-time on page 2318

stale-routes-time


Syntax	<code>stale-routes-time</code> <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• restart-time (BGP Graceful Restart) on page 2317

Configuration Statements for VRRP


- [accept-data on page 2320](#)
- [advertise-interval on page 2321](#)
- [asymmetric-hold-time on page 2322](#)
- [authentication-key on page 2323](#)
- [authentication-type on page 2324](#)
- [bandwidth-threshold on page 2325](#)
- [failover-delay on page 2326](#)
- [fast-interval on page 2327](#)
- [hold-time \(VRRP\) on page 2328](#)
- [interface \(VRRP Group\) on page 2329](#)
- [preempt \(VRRP\) on page 2330](#)

- [priority \(Protocols VRRP\) on page 2331](#)
- [priority-cost \(VRRP\) on page 2332](#)
- [priority-hold-time on page 2333](#)
- [route \(Interfaces\) on page 2334](#)
- [startup-silent-period on page 2335](#)
- [traceoptions on page 2336](#)
- [track \(VRRP\) on page 2338](#)
- [virtual-address on page 2339](#)
- [vrrp-group on page 2340](#)

accept-data

Syntax	(accept-data no-accept-data);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the master router accepts all packets destined for the virtual IP address.</p> <ul style="list-style-type: none"> • accept-data—Enable the master router to accept all packets destined for the virtual IP address. • no-accept-data—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.
Default	<p>If the router acting as the master router is the IP address owner or has its priority set to 255, the master router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the master router does not own the IP address or has its priority set to a value less than 255, the master router responds only to ARP requests.</p>
<div>  NOTE: <ul style="list-style-type: none"> • If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets. • If you include the accept-data statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>). </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an Interface to Accept Packets Destined for the Virtual IP Address</i>


advertise-interval

Syntax	<code>advertise-interval seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
<div>  <p>NOTE: When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div>	
Options	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Advertisement Interval for the VRRP Master Router</i> • fast-interval on page 2327 • <i>inet6-advertise-interval</i> • <i>version-3</i>

asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS 11.3 for the QFX Series.
Description	<p>Configure a VRRP master to fail over to a backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.</p> <p>When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.</p>
Default	asymmetric-hold-time is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP Preemption and Hold Time on page 2300

authentication-key

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement. All routers in the VRRP group must use the same authentication scheme and password.
<div>  NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups. </div>	
Options	key —Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring VRRP Authentication (IPv4 Only)</i> • Configuring VRRP Authentication (IPv4 Only) on page 2297 • authentication-type on page 2324 • <i>version-3</i>

authentication-type

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement. All routers in the VRRP group must use the same authentication scheme and password.



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

Options	<i>authentication</i> —Authentication scheme: <ul style="list-style-type: none">• simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.• md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. Default: none (no authentication is performed).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP Authentication (IPv4 Only)• Configuring VRRP Authentication (IPv4 Only) on page 2297• authentication-key on page 2323• version-3


bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 10000000000000 bits per second</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked • Configuring a Logical Interface to Be Tracked on page 2302

failover-delay

Syntax	<code>failover-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p>
Options	<p><i>milliseconds</i>—Specify the failover delay time, in milliseconds.</p> <p>Range: 50 through 2000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

fast-interval

Syntax	<code>fast-interval milliseconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
Options	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p>Range: 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p>NOTE: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <i>fast-interval</i>. Commit check fails if a value less than 100 is configured.</p> </div>	
Default: 1 second	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Advertisement Interval for the VRRP Master Router</i> • Configuring the Advertisement Interval for the VRRP Master on page 2299 • advertise-interval on page 2321 • advertise-interval on page 2321 • <i>inet6-advertise-interval</i> • <i>version-3</i>

hold-time (VRRP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Backup Router to Preempt the Master Router• Configuring VRRP Preemption and Hold Time on page 2300

interface (VRRP Group)


Syntax	<pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>; priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>bandwidth-threshold statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.




WARNING: On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

Options	<p><i>interface-name</i>—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked • Configuring a Logical Interface to Be Tracked on page 2302 • Junos Services Interfaces Configuration Release 12.3

preempt (VRRP)

Syntax	(preempt no-preempt) { hold-time seconds; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router:</p> <ul style="list-style-type: none">• preempt—Allow the master router to be preempted. <div> NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.</div> <ul style="list-style-type: none">• no-preempt—Prohibit the preemption of the master router. When no-preempt is configured, the backup router cannot preempt the master router even if the backup router has a higher priority. <p>The remaining statement is explained separately.</p>
Default	By default the preempt statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the preempt statement is not explicitly configured.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Backup Router to Preempt the Master Router• Configuring VRRP Preemption and Hold Time on page 2300


priority (Protocols VRRP)

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
<div>  NOTE: Priority 255 is not supported for routed VLAN interfaces (RVIs). </div>	
Options	<p><i>priority</i>—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 1 through 255</p> <p>Default: 100 (for backup routers)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support • Configuring Basic VRRP Support on page 2296

priority-cost (VRRP)

Syntax	<code>priority-cost priority;</code>
Hierarchy Level	<code>[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	priority —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. Range: 1 through 254
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked• Configuring a Logical Interface to Be Tracked on page 2302

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
<div>  <p>NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>	
Options	<p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked • Configuring a Logical Interface to Be Tracked on page 2302

route (Interfaces)

Syntax	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</code>
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches.
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>prefix</i>—Route to be tracked for this VRRP group.</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><i>routing-instance instance-name</i>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Route to Be Tracked• Configuring a Route to Be Tracked on page 2301

startup-silent-period

Syntax	startup-silent-period <i>seconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	seconds —Number of seconds for the startup period. Default: 4 seconds Range: 1 through 2000 seconds
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Startup Period for VRRP Operations</i>• Configuring the Startup Period for VRRP Operations on page 2298

traceoptions

Syntax traceoptions {
 file <filename> <files number> <match regular-expression> <microsecond-stamp>
 <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple **flag** statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



NOTE: The traceoptions statement is not supported on a QFabric system.

Default If you do not include this statement, no VRRP-specific tracing operations are performed.

Options **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

Range: 0 through 4,294,967,296 files

Default: 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes
- **general**—General events
- **interfaces**—Interface changes

- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Tracing VRRP Operations</i>

track (VRRP)

Syntax	<pre>track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>priority-hold-time statement added in Junos OS Release 8.1.</p> <p>route statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Logical Interface to Be Tracked</i>• <i>Configuring a Route to Be Tracked</i>• Configuring a Logical Interface to Be Tracked on page 2302• Configuring a Route to Be Tracked on page 2301

virtual-address

Syntax	<code>virtual-address [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Basic VRRP Support</i> • Configuring Basic VRRP Support on page 2296

vrrp-group

Syntax	<pre> vrrp-group group-id { (accept-data no-accept-data); advertise-interval seconds; authentication-key key; authentication-type authentication; fast-interval milliseconds; (preempt no-preempt) { hold-time seconds; } priority number; track { interface interface-name { bandwidth-threshold bits-per-second priority-cost priority; priority-cost priority; } priority-hold-time seconds; route prefix/prefix-length routing-instance instance-name priority-cost priority; } virtual-address [addresses]; vrrp-inherit-from vrrp-group; } </pre>
Hierarchy Level	<pre> [edit interfaces interface-name unit logical-unit-number family inet address address], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group.
Options	<p>group-id—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support • Example: Configuring VRRP • Configuring Basic VRRP Support on page 2296 • Example: Configuring VRRP for Load Sharing on page 2291

- *vrrp-inet6-group*

CHAPTER 26

Administration

- [Operational Mode Commands for Graceful Restart on page 2343](#)
- [Operational Mode Commands for Nonstop Software Upgrade on page 2367](#)
- [Operational Mode Commands for VRRP on page 2399](#)

Operational Mode Commands for Graceful Restart

- [Verifying Graceful Restart Operation on page 2343](#)
- [show bgp neighbor](#)
- [show log](#)
- [show \(ospf | ospf3\) overview](#)

Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 2343](#)
- [Verifying BGP Graceful Restart on page 2344](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 2344](#)
- [Verifying CCC and TCC Graceful Restart on page 2345](#)

Graceful Restart Operational Mode Commands

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)
- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)

- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the [CLI Explorer](#).

Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
  Type: Internal   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

  Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.255.10.1      Local ID: 192.255.5.1      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 180
  Stale routes from peer are kept for: 180
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19   Sent 19   Checked 19
  Input messages: Total 2   Updates 1   Refreshes 0   Octets 42
  Output messages: Total 3   Updates 0   Refreshes 0   Octets 116
  Output Queue[0]: 0
```

Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 2285](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct  8 05:20:12 Restart mode - sending grace lsas
Oct  8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct  8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct  8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct  8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
```

CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
CE1-CE2-0	rmt-if	Restart	-----	0
fe-1/1/0.0	intf	Up		
PE1-PE2-0	tlsp	Up		
PE2-PE1-0	rlsp	Up		
CE1-CE2-1	rmt-if	Restart	-----	0
fe-1/1/0.1	intf	Up		
PE1-PE2-1	tlsp	Up		
PE2-PE1-1	rlsp	Up		
CE1-CE2-2	rmt-if	Restart	-----	0
fe-1/1/0.2	intf	Up		
PE1-PE2-2	tlsp	Up		
PE2-PE1-2	rlsp	Up		
CE1-CE2-3	rmt-if	Restart	-----	0
fe-1/1/0.3	intf	Up		
PE1-PE2-3	tlsp	Up		
PE2-PE1-3	rlsp	Up		

- Related Documentation**
- [Graceful Restart Concepts on page 2269](#)
 - [Configuring Graceful Restart for QFabric Systems on page 1375](#)

show bgp neighbor

List of Syntax	Syntax on page 2346 Syntax (EX Series Switch and QFX Series) on page 2346
Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <<i>neighbor-address</i>> <orf (detail <i>neighbor-address</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <<i>neighbor-address</i>> <orf (<i>neighbor-address</i> detail)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>neighbor-address</i>—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Configuration Guide</i> .
Required Privilege Level	view

Related Documentation • [clear bgp neighbor on page 3585](#)

List of Sample Output [show bgp neighbor on page 2353](#)
[show bgp neighbor \(CLNS\) on page 2354](#)
[show bgp neighbor \(Layer 2 VPN\) on page 2355](#)
[show bgp neighbor \(Layer 3 VPN\) on page 2357](#)
[show bgp neighbor neighbor-address on page 2357](#)
[show bgp neighbor neighbor-address on page 2358](#)
[show bgp neighbor orf neighbor-address detail on page 2359](#)

Output Fields [Table 224 on page 2347](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 224: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	Current state of the BGP session: <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Local Address	Address of the local routing device.
Remove-private options	Options associated with the <code>remove-private</code> statement.
Holdtime	Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the <code>preference</code> statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <code>out-delay</code> parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .
Table inet.number	Information about the routing table: <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	Information about dropped path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	Information about ignored path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.

Table 224: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)

```

```

Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages:  Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
  Keepalive Interval: 30      Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages:  Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000

```

```
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0
```


show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.13vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 15      Sent 20      Checked 20
  Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
  Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal      State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None

```

```

Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6    Local ID: 10.255.245.5    Active Holdtime: 60000
Keepalive Interval: 20000 Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)

```

```

Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:           1 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:           0 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show log

List of Syntax	Syntax on page 2360 Syntax (QFabric System) on page 2360 Syntax (TX Matrix Routers) on page 2360
Syntax	<code>show log</code> <code><filename user <username>></code>
Syntax (QFabric System)	<code>show log filename</code> <code><device-type (device-id device-alias)></code>
Syntax (TX Matrix Routers)	<code>show log</code> <code><all-lcc lcc <i>number</i> scc></code> <code><filename user <username>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.
Options	none —List all log files. <all-lcc lcc <i>number</i> scc> —(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis). device-type —(QFabric system only) (Optional) Display log messages for only one of the following device types: <ul style="list-style-type: none">• director-device—Display logs for Director devices.• infrastructure-device—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).• interconnect-device—Display logs for Interconnect devices.• node-device—Display logs for Node devices.



NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

Required Privilege Level trace

List of Sample Output [show log on page 2361](#)
[show log filename on page 2361](#)
[show log filename \(QFabric System\) on page 2362](#)
[show log user on page 2362](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

show log user

```

user@host> show log user
darius  mg2546          Thu Oct  1 19:37   still logged in
darius  mg2529          Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518          Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575          Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

show (ospf | ospf3) overview

List of Syntax	Syntax on page 2363 Syntax (EX Series Switch and QFX Series) on page 2363
Syntax	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview on page 2365 show ospf overview (With Database Protection) on page 2366 show ospf3 overview (With Database Protection) on page 2366 show ospf overview extensive on page 2366
Output Fields	<p>Table 225 on page 2364 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.</p>

Table 225: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 225: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the <code>[edit protocols ospf area area-id]</code> level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 70
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
  Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 7
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
  Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
```

```

Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1

```

Operational Mode Commands for Nonstop Software Upgrade

- [Performing a Nonstop Software Upgrade on the QFabric System on page 2367](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 2372](#)
- `request system software nonstop-upgrade`
- `show chassis nonstop-upgrade node-group`

Performing a Nonstop Software Upgrade on the QFabric System



NOTE: Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the `request system software add component all` command) does not require that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a nonstop software upgrade to Junos OS Release 12.2X50-D42.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device or one fabric control Routing Engine is always operational
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device

- Specifying an upgrade group if you want all Node devices in a Node group to be upgraded in parallel (which shortens the time of the upgrade)
- Rebooting devices automatically as part of the nonstop upgrade process

When performing a nonstop upgrade, start with the Director group upgrade, then issue the fabric upgrade, and end with the Node group upgrades.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade. For node-groups defined with two node-devices, both must be online in order for upgrade to succeed.



NOTE: Before you install the software, we recommend that you back up your current configuration files by issuing the `request system software configuration-backup` command.



NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the `request system software add component all` command.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 2368](#)
- [Downloading Software Files Using a Browser on page 2369](#)
- [Retrieving Software Files for Download on page 2370](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group on page 2370](#)
- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components on page 2370](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups on page 2371](#)
- [Performing a Nonstop Software Upgrade on a Node Group on page 2372](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@qfabric> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup/media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
 2. Click **Download Software**.
 3. In the **Switching** box, click **Junos OS Platforms**.
 4. In the **QFX Series** section, click the name of the platform for which you want to download software.
 5. Click the **Software** tab and select the release number from the **Release** drop-down list.
 6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see “Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device” on page 111.
- A login screen appears.
7. Enter your user ID and password and press **Enter**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@qfabric> request system software download /path/package-name
```

For example:

```
user@qfabric> request system software download  
ftp://server/files/jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group



NOTE: If you reboot any Node groups or Interconnect devices after you perform a nonstop upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group  
jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To verify that both Director devices are online, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric  
jinstall-qfabric-12.2X50-D10.3.rpm
```

(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node group, create an upgrade group and include the devices you want to reboot at the same time.



NOTE:

- Upgrade groups work best when you build in redundancy so that not all the Node devices within the Node group restart at the same time. We suggest using the upgrade group feature for Node groups that have at least 2 pairs of Node devices (4 or more members), such as the network Node group.
 - If you add Node devices that have links to the same link aggregation group (LAG), there might be traffic loss.
-
- Create the upgrade group by issuing the **set chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name* node-devices** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade1 node-devices
[ node1 node2 ]
user@qfabric# set chassis node-group NW-NG-0 nssu upgrade-group upgrade2 node-devices
[ node3 node4 ]
```

Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.



NOTE: You can upgrade multiple Node groups with this command. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To install software on a Node group:

- Issue the **request system software nonstop-upgrade node-group *node-group-name* *package-name*** command.

To perform a nonstop upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1  
jinstall-qfabric-12.2X50-D10.3.rpm
```

To perform a nonstop upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1  
nodegroup2 nodegroup3] jinstall-qfabric-12.2X50-D10.3.rpm
```

Related Documentation

- [Configuring Graceful Restart for QFabric Systems on page 1375](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 1410](#)
- [request system software nonstop-upgrade on page 410](#)
- [show system software upgrade status on page 1570](#)

Verifying Nonstop Software Upgrade for QFabric Systems

This topic discusses how you can monitor the progress of each of the three steps in a nonstop software upgrade. By identifying the key actions and events that define this process, you can track the status of the upgrade with confidence.



TIP: When performing a nonstop software upgrade, open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a

second session to verify that the QFabric system components respond to operational mode commands as expected.

- [Verifying a Director Group Nonstop Software Upgrade on page 2373](#)
- [Verifying a Fabric Nonstop Software Upgrade on page 2385](#)
- [Verifying a Redundant Server Node Group Nonstop Software Upgrade on page 2387](#)
- [Verifying a Network Node Group Nonstop Software Upgrade on page 2390](#)

Verifying a Director Group Nonstop Software Upgrade

Purpose During the Director group portion of a nonstop software upgrade, you should expect to see the Director device that hosts the CLI session selected as the master device. When mastership of all processes moves to the master, the QFabric system upgrades the backup Director device and this Director device reboots. After the backup Director device comes back online, the master Director device suspends CLI operations for 15 minutes, upgrades itself, and reboots. At this point, the backup becomes the new master Director device and you can issue CLI operational commands. Finally, the former master comes back online as a backup and both devices are operational once again. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In one SSH session to the QFabric CLI, verify the current status of the QFabric system by issuing the **show fabric administration inventory**, **show fabric administration inventory director-group status**, and **show fabric session-host** commands. In this case, Director device DG0 is the master device but DG1 hosts the CLI session.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			

DRE-0 Connected Configured

session1@qfabric> show fabric administration inventory director-group status

Director Group Status Tue Jun 5 15:11:26 UTC 2012

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

```

-----
Management Interface          up
Control Plane Bridge          up
Control Plane LAG              up
CP Link [0/2]                  up
CP Link [0/1]                  up
CP Link [0/0]                  up
CP Link [1/2]                  down
CP Link [1/1]                  down
CP Link [1/0]                  down
Crossover LAG                  up
CP Link [0/3]                  up
CP Link [1/3]                  up

```

```

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032  online  backup

```

Director Group Managed Services

```

-----
Shared File System          online
Network File System          online
Virtual Machine Server       online
Load Balancer/DHCP           online

```

Hard Drive Status

```

-----
Volume ID:8                  optimal
Physical ID:1                 online
Physical ID:0                 online
SCSI ID:1                     100%
SCSI ID:0                     100%

```

```

Size  Used Avail Used% Mounted on
-----

```

```

423G 5.5G 395G  2%  /
99M  16M  79M  17% /boot
93G  7.3G 86G   8% /pbdata

```

Director Group Processes

```

-----
Director Group Manager       online
Partition Manager             online
Software Mirroring            online
Shared File System master     online
Secure Shell Process          online
Network File System           online
DHCP Server master            online      backup

FTP Server                     online
Syslog                         online
Distributed Management         online
SNMP Trap Forwarder           online
SNMP Process                   online
Platform Management           online

```

Interface Link Status

```

-----
Management Interface          up

```

```

Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG            up
CP Link [0/3]             up
CP Link [1/3]             up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

- In a second SSH session to the QFabric CLI, issue the request for the Director group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm

```

- If the CLI session is being hosted by the master Director device, skip to step 4. However, if the CLI session is hosted by the backup Director device, the Director group mastership switches to the backup device after you issue the nonstop software upgrade command. In this example, mastership switches to Director device DG1.

```

session1@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 15:12:20 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.49.215.38	8%	31905924k	0	3 days, 21:16 hrs
dg1	online	master	10.49.215.39	6%	18010368k	3	3 days, 21:16 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	backup

Director Group Managed Services

Shared File System	offline
Network File System	offline
Virtual Machine Server	offline
Load Balancer/DHCP	offline

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         offline
DHCP Server master          offline    backup

FTP Server                  online
Syslog                      online
Distributed Management       offline
SNMP Trap Forwarder         offline
SNMP Process                offline
Platform Management         online

Interface Link Status
-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  master

Master Services
-----
Database Server            online
Load Balancer Director     online
QFabric Partition Address  online

Director Group Managed Services
-----
Shared File System         online
Network File System        online
Virtual Machine Server     online
Load Balancer/DHCP         online

Hard Drive Status
-----
Volume ID:8                optimal
Physical ID:1              online
Physical ID:0              online
SCSI ID:1                  100%
SCSI ID:0                  100%

Size  Used Avail Used% Mounted on
-----
423G  6.0G 395G   2%  /
99M   16M  79M  17%  /boot

```

93G 7.3G 86G 8% /pbdata

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master         online      master

FTP Server                  online
Syslog                      online
Distributed Management       online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge         up
Control Plane LAG           up
CP Link [0/2]                up
CP Link [0/1]                up
CP Link [0/0]                up
CP Link [1/2]                down
CP Link [1/1]                down
CP Link [1/0]                down
Crossover LAG               up
CP Link [0/3]                up
CP Link [1/3]                up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

4. The Director group nonstop software upgrade process continues by downloading and installing software for the fabric manager Routing Engines and the Director devices.

```

root@qfabric>
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

```

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

```
Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package
jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
```

5. When the system upgrades and reboots the backup Director device DG0, notice how this device is not displayed in the output of the **show fabric administration inventory director-group status** command. Because Director device DG1 appears, this means that the DG1 is operational and acts as the master device.



NOTE: If your second SSH session is being hosted by the rebooting Director device, your session terminates and you need to log back in to establish a new session running on the active Director device.

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 15:41:14 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg1	online	master	10.49.215.39	6%	8372272k	4	3 days, 21:25 hrs

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:8	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	6.0G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up


```

Control Plane LAG          up
CP Link [0/2]              up
CP Link [0/1]              up
CP Link [0/0]              up
CP Link [1/2]              down
CP Link [1/1]              down
CP Link [1/0]              down
Crossover LAG              up
CP Link [0/3]              up
CP Link [1/3]              up

```

6. The upgrade continues with master Director device DG1 suspending CLI services for 15 minutes, transferring mastership to Director device DG0, and then rebooting Director device DG1 (which terminates the CLI session).

```
root@qfabric>
```

```
[Peer Update Status]: Setting peer DG node as the master SFC
```

```
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
```

```
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
```

```
[Peer Update Status]: Check for VMs on dg0
```

```
Triggering Final Stage of Fabric Manager Upgrade:
```

```
Updating FM-0 to Junos version 12.2X50-D10.3
```

```
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
```

```
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
```

```
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
```

```
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
```

```
[DRE-0 2012-06-05 16:12:43]: Retrieving package
```

```
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
```

```
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
```

```
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
```

```
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
```

```
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
```

```
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
```

```
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
```

```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
```

```
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!
```

```
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by
peer
```

```
Connection to qfabric-partition0 closed.
```

7. Upon reopening the SSH session, notice that Director device DG0 is now the master device hosting the session and Director device DG1 does not appear in the QFabric system inventory while it is rebooting.

```
session1@qfabric> show fabric session-host
```

```
Identifier: 0281052011000001
```

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 16:21:23 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	13%	20739560k	3	36:29 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

```

Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG             up
CP Link [0/3]             up
CP Link [1/3]             up

```

8. When Director device DG1 comes back online, it returns to the QFabric system inventory as a backup Director device and hosts some of the Routing Engine processes (which should appear load balanced between the master and backup Director devices).

```
session1@qfabric> show fabric administration inventory director-group status
```

```
root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 16:41:02 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	15%	14759920k	6	56:09 mins
dg1	online	backup	10.49.215.39	8%	31486680k	0	07:51 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

```

Database Server           online
Load Balancer Director    online
QFabric Partition Address online

```

Director Group Managed Services

```

Shared File System        online
Network File System        online
Virtual Machine Server    online
Load Balancer/DHCP         online

```

Hard Drive Status

```

Volume ID:4               optimal
Physical ID:1             online
Physical ID:0             online
SCSI ID:1                 100%
SCSI ID:0                 100%

```

Size	Used	Avail	Used%	Mounted on
423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online      master

FTP Server                  online
Syslog                      online
Distributed Management       online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	backup

Director Group Managed Services

```

-----
Shared File System          online
Network File System         online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Volume ID:8                 optimal
Physical ID:1               online
Physical ID:0               online
SCSI ID:1                   100%
SCSI ID:0                   100%

```

Size Used Avail Used% Mounted on

```

-----
423G 5.3G 396G 2% /
99M 16M 79M 17% /boot
93G 7.4G 86G 8% /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online

```

```

Partition Manager          online
Software Mirroring         online
Shared File System master  online
Secure Shell Process       online
Network File System        online
DHCP Server master         online    backup

FTP Server                 online
Syslog                    online
Distributed Management     online
SNMP Trap Forwarder        online
SNMP Process               online
Platform Management        online

```

Interface Link Status

```

-----
Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG             up
CP Link [0/3]             up
CP Link [1/3]             up

```

session1@qfabric> show fabric administration inventory infrastructure

dg0:

Routing Engine Type CPU-Use(%)	Hostname	PID	
Fabric control	QFabric_default_FC-1_RE0	27906	2.5
Network Node group	QFabric_default_NW-NG-1_RE1	20421	1.8
Fabric manager	FM-0	4211	1.8
Debug Routing Engine	QFabric_DRE	1575	3.3

dg1:

Routing Engine Type CPU-Use(%)	Hostname	PID	
Fabric control	QFabric_default_FC-0_RE0	5686	2.3
Network Node group	QFabric_default_NW-NG-0_RE0	5866	1.9
Fabric manager	FM-1	572	1.6

Verifying a Fabric Nonstop Software Upgrade

Purpose During the fabric portion of a nonstop software upgrade, you should expect to see both fabric control Routing Engines upgrade first, followed by the upgrade of each Interconnect

device one at a time. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the fabric nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0    2012-06-05 16:48:53]: Retrieving package
[FC-1    2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0    2012-06-05 16:49:51]: ----- re0: -----
[FC-1    2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active
partition
[Status  2012-06-05 16:58:34]: Load completed with 0 errors...
[Status  2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status  2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status  2012-06-05 16:58:39]: Rebooting FC-0
[Status  2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status  2012-06-05 16:58:44]: Rebooting FC-1
[Status  2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status  2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

2. When the fabric components reboot, they appear as **Disconnected** in the output of the **show fabric administration inventory infrastructure fabric-controls** and **show fabric administration inventory interconnect-devices** commands.

```
session1@qfabric> show fabric administration inventory infrastructure fabric-controls
Item                               Identifier                               Connection                               Configuration
Fabric control
FC-0                               Disconnected
FC-1                               Disconnected

session1@qfabric> show fabric administration inventory interconnect-devices IC-F4912
Item                               Identifier                               Connection                               Configuration
Interconnect device
```

```

IC-F4912
F4912/RE0
Disconnected
Disconnected

```

3. When the fabric components return to full service, they appear as **Connected** in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

Verifying a Redundant Server Node Group Nonstop Software Upgrade

Purpose During the redundant server Node group portion of a nonstop software upgrade, you should expect to see the backup Node device upgrade first, followed by the upgrade of the master Node device. Server Node groups with a single device upgrade the device in the same way as a standalone switch. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the redundant server Node group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): RSNG

```

```

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system

```

```
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last
group :1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed
[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```


2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
                    RSNG in progress
```

3. During the redundant server Node group upgrade, the backup Node device (in this case, P1571-C) is upgraded first and appears in the **Disconnected** state in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
	NW-NG-0	Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Disconnected	
Interconnect device			
	IC-F4912	Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
	FM-0	Connected	Configured
Fabric control			
	FC-0	Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			
	DRE-0	Connected	Configured

4. After the backup Node device comes back online, the master Node device (in this case, P1550-C) appears in the **Disconnected** state in the output of the **show fabric administration inventory** command while the master Node device upgrades its software.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
	NW-NG-0	Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Disconnected	
	P1571-C	Connected	
Interconnect device			
	IC-F4912	Connected	Configured
	F4912/RE0	Connected	

Fabric manager FM-0	Connected	Configured
Fabric control FC-0	Connected	Configured
FC-1	Connected	Configured
Diagnostic routing engine DRE-0	Connected	Configured

5. After both Node devices in the redundant server Node group come back online, both Node devices appear as **Connected** to indicate the successful completion of the Node group nonstop software upgrade step.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Connected	
Interconnect device			
IC-F4912		Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

Verifying a Network Node Group Nonstop Software Upgrade

Purpose During the network Node group portion of a nonstop software upgrade, you should expect to see the backup network Node group Routing Engine upgrade first, followed by the Node devices within the network Node group upgrading one at a time, and ending with the upgrade of the master network Node group Routing Engine. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.



NOTE: If you configure an upgrade group for Node groups containing 2 or more Node devices, all Node devices within the upgrade group reboot at the same time.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the network Node group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade node-group NW-NG-0
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): NW-NG-0

[NW-NG-0 2012-06-01 09:45:06]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[NW-NG-0 2012-06-01 09:45:06]: Retrieving package
[NW-NG-0 2012-06-01 09:46:18]: Pushing bundle to fpc0
[NW-NG-0 2012-06-01 09:46:52]: fpc0: Validate package...
[NW-NG-0 2012-06-01 09:53:26]: ----- fpc0 -----
[NW-NG-0 2012-06-01 09:54:01]: Step 1 of 20 Creating temporary file system
[NW-NG-0 2012-06-01 09:54:01]: Step 2 of 20 Determining installation source
[NW-NG-0 2012-06-01 09:54:02]: Step 3 of 20 Processing format options
[NW-NG-0 2012-06-01 09:54:02]: Step 4 of 20 Determining installation slice
[NW-NG-0 2012-06-01 09:54:02]: Step 5 of 20 Creating and labeling new slices
[NW-NG-0 2012-06-01 09:54:03]: Step 6 of 20 Create and mount new file system
[NW-NG-0 2012-06-01 09:54:10]: Step 7 of 20 Getting OS bundles
[NW-NG-0 2012-06-01 09:54:10]: Step 8 of 20 Updating recovery media
[NW-NG-0 2012-06-01 09:54:31]: Step 9 of 20 Extracting incoming image
[NW-NG-0 2012-06-01 09:55:43]: Step 10 of 20 Unpacking OS packages
[NW-NG-0 2012-06-01 09:55:46]: Step 11 of 20 Mounting jbase package
[NW-NG-0 2012-06-01 09:56:09]: Step 12 of 20 Creating base OS symbolic links
[NW-NG-0 2012-06-01 09:57:05]: Step 13 of 20 Creating fstab
[NW-NG-0 2012-06-01 09:57:05]: Step 14 of 20 Creating new system files
[NW-NG-0 2012-06-01 09:57:05]: Step 15 of 20 Adding jbundle package
[NW-NG-0 2012-06-01 09:59:30]: Step 16 of 20 Backing up system data
[NW-NG-0 2012-06-01 09:59:44]: Step 17 of 20 Setting up shared partition data
[NW-NG-0 2012-06-01 09:59:44]: Step 18 of 20 Checking package sanity in
installation
[NW-NG-0 2012-06-01 09:59:44]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[NW-NG-0 2012-06-01 09:59:47]: Step 20 of 20 Setting da0s1 as new active
partition
[NW-NG-0 2012-06-01 09:59:55]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-dc-re.tgz
[NW-NG-0 2012-06-01 09:59:55]: Retrieving package
[NW-NG-0 2012-06-01 10:01:04]: Pushing bundle to re1
[NW-NG-0 2012-06-01 10:01:35]: re1: Validate package...
[NW-NG-0 2012-06-01 10:02:56]: re0: Validate package...
[NW-NG-0 2012-06-01 10:04:45]: Rebooting Backup RE
[NW-NG-0 2012-06-01 10:08:31]: Initiating Chassis In-Service-Upgrade
[NW-NG-0 2012-06-01 10:08:52]: Upgrading group: 0 fpc: 0
[NW-NG-0 2012-06-01 10:18:33]: Upgrade complete for group:0
[NW-NG-0 2012-06-01 10:18:33]: Finished processing all upgrade groups, last
group :0
[NW-NG-0 2012-06-01 10:18:37]: Preparing for Switchover
[NW-NG-0 2012-06-01 10:18:55]: Switchover Completed
[Status 2012-06-01 10:18:58]: Upgrade completed with 0 errors
Success

```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```

root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
NW-NG-0 in progress

```

3. Verify the progress of the upgrade by issuing the **show chassis nonstop-upgrade node-group**, **show fabric administration inventory**, **show fabric administration inventory infrastructure**, and **show fabric administration inventory node-groups NW-NG-0** commands. You should see the backup network Node group Routing Engine reboot first, followed by each Node device within the network Node group, and ending with the reboot of master network Node group Routing Engine. Restarting devices appear as **Disconnected** in the output of the **show fabric administration inventory** command and restarting Routing Engines do not appear in output of the **show fabric administration inventory infrastructure** command until they return to service.

**Related
Documentation**

- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [show chassis nonstop-upgrade node-group on page 817](#)
- [show fabric administration inventory on page 1543](#)
- [show fabric administration inventory director-group status on page 1548](#)
- [show fabric administration inventory infrastructure on page 1553](#)
- [show fabric administration inventory interconnect-devices on page 1556](#)
- [show fabric administration inventory node-groups on page 1560](#)
- [show system software upgrade status on page 1570](#)

request system software nonstop-upgrade

Syntax	request system software nonstop-upgrade <i>package-name</i> <fabric > <director-group> <node-group <i>name</i> >
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.
Options	<p><i>package-name</i>—Location from which the software is to be installed. For example:</p> <ul style="list-style-type: none"> • <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following: <ul style="list-style-type: none"> • ftp—File Transfer Protocol. Use <i>ftp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>ftp://<username>:<password>@hostname/pathname/package-name</i>. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed. • http—Hypertext Transfer Protocol. Use <i>http://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>http://<username>:<password>@hostname/pathname/package-name</i>. If a password is required and you omit it, you are prompted for it. • scp—Secure copy (available only for Canada and U.S. version). Use <i>scp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>scp://<username>:<password>@hostname/pathname/package-name</i>.



NOTE:

- The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

director-group—Install software package on the Director group and Fabric managers.

fabric—Install software package on the Interconnect devices and Fabric controls.

node-group *name* —Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Performing a Nonstop Software Upgrade on the QFabric System on page 105• Verifying Nonstop Software Upgrade for QFabric Systems on page 1410• show chassis nonstop-upgrade node-group on page 817
List of Sample Output	request system software nonstop-upgrade director-group on page 2394 request system software nonstop-upgrade fabric on page 2396 request system software nonstop-upgrade node-group (Redundant Server Node Group) on page 2396 request system software nonstop-upgrade node-group (Server Node Group) on page 2398
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software nonstop-upgrade director-group](#)

```
user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
```

```

[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
Setting peer DG node as the master SFC

```

```

Delaying start of local upgrade to allow peer services time to initialize [15
minutes]

```

```

Delaying start of local upgrade to allow peer services time to initialize [15
minutes]

```

```

Delaying start of local upgrade to allow peer services time to initialize [12
minutes]

```

```

Delaying start of local upgrade to allow peer services time to initialize [9
minutes]

```

```

Delaying start of local upgrade to allow peer services time to initialize [6
minutes]

```

```

Delaying start of local upgrade to allow peer services time to initialize [3
minutes]

```

```

[Peer Update Status]: Check for VMs on dg0

```

```

Triggering Final Stage of Fabric Manager Upgrade:

```

```

Updating FM-0 to Junos version 12.2X50-D10.3

```

```

[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started

```

```

[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG

```

```

[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch

```

```

[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3

```

```

[DRE-0 2012-06-05 16:12:43]: Retrieving package

```

```

[DRE-0 2012-06-05 16:13:46]: ----- re0: -----

```

```

[Status 2012-06-05 16:15:17]: Load completed with 0 errors...

```

```

[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...

```

```

[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...

```

```

[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...

```

```

[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0

```

```

[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...

```

```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...  
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!  
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer  
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric  
jinstall-qfabric-12.2X50-D10.3.rpm  
[FC-0 2012-06-05 16:48:53]: Retrieving package  
[FC-1 2012-06-05 16:48:53]: Retrieving package  
[IC-F4912 2012-06-05 16:48:59]: Retrieving package  
[FC-0 2012-06-05 16:49:51]: ----- re0: -----  
[FC-1 2012-06-05 16:49:52]: ----- re0: -----  
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----  
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system  
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source  
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options  
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice  
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices  
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system  
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles  
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media  
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image  
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages  
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package  
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links  
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab  
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files  
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package  
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data  
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data  
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in  
installation  
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary  
file systems  
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition  
[Status 2012-06-05 16:58:34]: Load completed with 0 errors...  
[Status 2012-06-05 16:58:34]: Reboot is required to complete upgrade ...  
[Status 2012-06-05 16:58:34]: Trying to Connect to Node: FC-0  
[Status 2012-06-05 16:58:39]: Rebooting FC-0  
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1  
[Status 2012-06-05 16:58:44]: Rebooting FC-1  
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912  
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912  
Success
```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group RSNG  
jinstall-qfabric-12.2X50-D10.3.rpm  
Upgrading target(s): RSNG  
  
[RSNG 2012-06-05 17:26:44]: Starting with package  
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
```



```

[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last group
:1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed

```

[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): P1507-C

[P1507-C 2012-06-26 14:02:44]: Retrieving package
[P1507-C 2012-06-26 14:03:21]: ----- P1507-C: -----
[P1507-C 2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C 2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C 2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C 2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C 2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C 2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C 2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C 2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
[P1507-C 2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C 2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C 2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C 2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C 2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C 2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C 2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C 2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C 2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C 2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in
installation
[P1507-C 2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[P1507-C 2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status 2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status 2012-06-26 14:10:32]: Rebooting P1507-C
[Status 2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```

show chassis nonstop-upgrade node-group

Syntax	show chassis nonstop-upgrade node-group <i>node-group-name</i>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display the status of the Node group after the most recent nonstop software upgrade (NSSU).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 105 • request system software nonstop-upgrade on page 410
List of Sample Output	show chassis nonstop-upgrade node-group on page 2399
Output Fields	Table 68 on page 817 lists the output fields for the show chassis nonstop-upgrade node-group command. Output fields are listed in the approximate order in which they appear.

Table 226: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```

user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
Item           Status           Reason
P1550-C        Online

```

Operational Mode Commands for VRRP

- [show vrrp](#)

show vrrp

Syntax	<pre>show vrrp <brief detail extensive summary> <interface <i>interface-name</i>> <track interfaces></pre>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information and status about VRRP groups.
Options	<p>none—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display information and status about the specified VRRP interface.</p> <p>track interfaces—(Optional) Display information and status about VRRP track interfaces.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Configuring VRRP for IPv6 (CLI Procedure)
List of Sample Output	show vrrp on page 2405 show vrrp brief on page 2405 show vrrp detail (IPv6) on page 2405 show vrrp detail (Route Track) on page 2406 show vrrp extensive on page 2406 show vrrp interface on page 2407 show vrrp summary on page 2408 show vrrp track detail on page 2408 show vrrp track summary on page 2409
Output Fields	Table 227 on page 2400 lists the output fields for the show vrrp command. Output fields are listed in the approximate order in which they appear.

Table 227: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	none, brief, extensive, summary
Interface index	Physical interface index number, which reflects its initialization sequence.	extensive
Groups	Total number of VRRP groups configured on the interface.	extensive

Table 227: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	extensive
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> • Advertisement sent—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • Advertisement received—Number of VRRP advertisement PDUs received by the interface. • Packets received—Number of VRRP packets received for VRRP groups on the interface. • No group match received—Number of VRRP packets received for VRRP groups that do not exist on the interface. 	extensive
Interface VRRP PDU error statistics	Errored statistics for the logical interface: <ul style="list-style-type: none"> • Invalid IPAH next type received—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • Invalid VRRP ttl value received—Number of packets received whose IP time-to-live (TTL) value is not 255. • Invalid VRRP version received—Number of packets received whose VRRP version is not 2. • Invalid VRRP pdu type received—Number of packets received whose VRRP PDU type is not 1. • Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5. • Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8. • Invalid VRRP checksum received—Number of packets received whose VRRP checksum does not match the calculated value. 	extensive
Physical interface	Name of the physical interface.	detail, extensive
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	none, brief, detail, extensive
Index	Physical interface index number, which reflects its initialization sequence.	detail, extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail, extensive
VRRP-Traps	Status of VRRP traps: Enabled or Disabled .	detail, extensive

Table 227: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> lcl—Configured local interface address. mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. vip—Configured virtual IP addresses. 	none, brief, summary
Interface state or Int state	State of the physical interface: <ul style="list-style-type: none"> down—The device is present and the link is unavailable. not present—The interface is configured, but no physical device is present. unknown—The VRRP process has not had time to query the kernel about the state of the interface. up—The device is present and the link is established. 	none, brief, extensive, summary
Group	VRRP group number.	none, brief, extensive, summary
State	VRRP state: <ul style="list-style-type: none"> backup—The interface is acting as the backup router interface. bringup—VRRP is just starting, and the physical device is not yet present. idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. initializing—VRRP is initializing. master—The interface is acting as the master router interface. transition—The interface is changing between being the backup and being the master router. 	extensive
Priority	Configured VRRP priority for the interface.	detail, extensive
Advertisement interval	Configured VRRP advertisement interval.	detail, extensive
Authentication type	Configured VRRP authentication type: none , simple , or md5 .	detail, extensive
Preempt	Whether preemption is allowed on the interface: yes or no .	detail, extensive
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: yes or no .	detail, extensive
VIP count	Number of virtual IP addresses that have been configured on the interface.	detail, extensive
VIP	List of virtual IP addresses configured on the interface.	detail, extensive
Advertisement timer	Time until the advertisement timer expires.	detail, extensive

Table 227: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is N/A .	detail, extensive
Virtual router uptime	Time that the virtual router has been up.	detail, extensive
Master router uptime	Time that the master router has been up.	detail, extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail, extensive
Tracking	Whether tracking is enabled or disabled .	detail, extensive
Current priority	Current operational priority for being the VRRP master.	detail, extensive
Configured priority	Configured base priority for being the VRRP master.	detail, extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. Disabled indicates no minimum interval.	detail, extensive
Remaining-time	(track option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface	Name of the tracked interface.	detail extensive
Int state/Interface state	Current operational state of the tracked interface: up or down .	detail, extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail, extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail, extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of down means that the corresponding priority cost is incurred when the interface is down.	detail, extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail, extensive
Route count	The number of routes being tracked.	detail, extensive

Table 227: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Route	The IP address of the route being tracked.	detail, extensive
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	detail, extensive
Route state	The state of the route being tracked: up , down , or unknown .	detail, extensive
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	detail, extensive
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	detail, extensive
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	extensive
Group VRRP PDU error statistics	Errored statistics for the VRRP group: <ul style="list-style-type: none"> • Bad authentication type received—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be none, simple, or md5 and must be the same for all routers in the VRRP group. • Bad password received—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group. • Bad MD5 digest received—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • Bad advertisement timer received—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • Bad VIP count received—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • Bad VIPADDR received—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance. 	extensive
Group state transition statistics	State transition statistics for the VRRP group: <ul style="list-style-type: none"> • Idle to master transitions—Number of times that the VRRP instance transitioned from the idle state to the master state. • Idle to backup transitions—Number of times that the VRRP instance transitioned from the idle state to the backup state. • Backup to master transitions—Number of times that the VRRP instance transitioned from the backup state to the master state. • Master to backup transitions—Number of times that the VRRP instance transitioned from the master state to the backup state. 	extensive
Vlan-id	ID of Vlan	detail

Table 227: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
VR state	VRRP information: <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • initializing—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. 	none, brief
Timer	VRRP timer information: <ul style="list-style-type: none"> • A—Time, in seconds, until the advertisement timer expires. • D—Time, in seconds, until the Master is Dead timer expires. 	none, brief

Sample Output

show vrrp

```

user@host> show vrrp
Interface      State      Group   VR state  Timer  Type  Address
ge-0/0/0.121   up         1       master    A 1.052 1c1   gec0::12:1:1:1
                                     vip    ge80::12:1:1:99
                                     vip    gec0::12:1:1:99
                                     vip    gec0::13:1:1:1
ge-0/0/2.131   up         1       master    A 0.364 1c1   gec0::13:1:1:1
                                     vip    ge80::13:1:1:99
                                     vip    gec0::13:1:1:99

```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 2405](#).

show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
 gec0::13:1:1:99
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled

show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
 Interface state: up, Group: 100, State: master
 Priority: 150, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
 Advertisement timer: 1.218s, Master router: 30.30.30.30
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
 Virtual MAC: 00:00:5e:00:01:64
 Tracking: enabled
 Current priority: 150, Configured priority: 150
 Priority hold-time: disabled
 Interface tracking: disabled
 Route tracking: enabled, Route count: 1

Route	VRF name	Route state	Priority cost
192.168.40.0/22	default	up	30

show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

Interface VRRP PDU statistics

Advertisement sent	:	188
Advertisement received	:	0
Packets received	:	0
No group match received	:	0

Interface VRRP PDU error statistics

Invalid IPAH next type received	:	0
Invalid VRRP TTL value received	:	0
Invalid VRRP version received	:	0
Invalid VRRP PDU type received	:	0
Invalid VRRP authentication type received	:	0
Invalid VRRP IP count received	:	0
Invalid VRRP checksum received	:	0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
 gec0::12:1:1:99
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled
 Group VRRP PDU statistics

```

    Advertisement sent           :          188
    Advertisement received       :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received           :         0
    Bad MD5 digest received         :         0
    Bad advertisement timer received:         0
    Bad VIP count received          :         0
    Bad VIPADDR received            :         0
Group state transition statistics
    Idle to master transitions       :         0
    Idle to backup transitions       :         1
    Backup to master transitions     :         1
    Master to backup transitions     :         0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
    Advertisement sent           :          186
    Advertisement received       :           0
    Packets received             :           0
    No group match received       :           0
Interface VRRP PDU error statistics
    Invalid IPAH next type received :         0
    Invalid VRRP TTL value received :         0
    Invalid VRRP version received   :         0
    Invalid VRRP PDU type received  :         0
    Invalid VRRP authentication type received:         0
    Invalid VRRP IP count received  :         0
    Invalid VRRP checksum received  :         0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent           :          186
    Advertisement received       :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received           :         0
    Bad MD5 digest received         :         0
    Bad advertisement timer received:         0
    Bad VIP count received          :         0
    Bad VIPADDR received            :         0
Group state transition statistics
    Idle to master transitions       :         0
    Idle to backup transitions       :         1
    Backup to master transitions     :         1
    Master to backup transitions     :         0

```

show vrrp interface

```
user@host> show vrrp interface
```

```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
  Packets received             :           0
  No group match received      :           0
Interface VRRP PDU error statistics
  Invalid IPAH next type received :          0
  Invalid VRRP TTL value received :          0
  Invalid VRRP version received  :          0
  Invalid VRRP PDU type received :          0
  Invalid VRRP authentication type received:          0
  Invalid VRRP IP count received :          0
  Invalid VRRP checksum received :          0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
Group VRRP PDU error statistics
  Bad authentication type received:          0
  Bad password received         :           0
  Bad MD5 digest received       :           0
  Bad advertisement timer received:          0
  Bad VIP count received        :           0
  Bad VIPADDR received         :           0
Group state transition statistics
  Idle to master transitions     :           0
  Idle to backup transitions     :           1
  Backup to master transitions   :           1
  Master to backup transitions   :           0

```

show vrrp summary

```

user@host> show vrrp summary

```

Interface	State	Group	VR state	Type	Address
ge-4/1/0.0	up	1	backup	lcl	10.57.0.2
				vip	10.57.0.100

show vrrp track detail

```

user@host> show vrrp track detail
Tracked interface: ae1.211
State: up, Speed: 400m
Incurred priority cost: 0

```

Threshold	Priority cost	Active
400m	10	
300m	60	
200m	110	
100m	160	
down	190	

```
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled,    Remaining-time: 50.351
```

show vrrp track summary

```
user@host> show vrrp track summary
```

Track if	State	Speed	VRRP if	Group	VR State	Current priority
ae1.211	up	400m	ae0.210	1	master	200

CHAPTER 27

Troubleshooting

- [Troubleshooting Procedures on page 2411](#)

Troubleshooting Procedures

- [Troubleshooting VRRP on page 2411](#)

Troubleshooting VRRP

Problem **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

Solution Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

Related Documentation • [failover-delay on page 2326](#)

PART 10

Interfaces

- [Overview on page 2415](#)
- [Configuration on page 2449](#)
- [Administration on page 2627](#)
- [Troubleshooting on page 2811](#)

CHAPTER 28

Overview

- [Interfaces Overview on page 2415](#)

Interfaces Overview

- [Interfaces Overview on page 2415](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2417](#)
- [Understanding Interface Naming Conventions on page 2419](#)
- [Understanding Interface Ranges on the QFX Series on page 2432](#)
- [Understanding Layer 3 Logical Interfaces on page 2434](#)
- [Understanding Management Interfaces on page 2434](#)
- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Overview of Uplink Failure Detection on page 2445](#)

Interfaces Overview

Juniper Networks QFX Series products have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *Junos® OS Network Interfaces*.

- [Network Interfaces on page 2415](#)
- [Special Interfaces on page 2416](#)

Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 228 on page 2415](#) lists the types of network interfaces supported on the QFX Series.

Table 228: Network Interface Types and Purposes

Type	Purpose
Aggregated Ethernet interfaces	You can group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.

Table 228: Network Interface Types and Purposes (*continued*)

Type	Purpose
Fibre Channel interfaces	You can use Fibre Channel interfaces to connect the switch to a Fibre Channel over Ethernet (FCoE) forwarder or a Fibre Channel switch in a storage area network (SAN). You can configure Fibre Channel interfaces only on ports 0 through 5 and 42 through 47 on QFX3500 devices. Fibre Channel interfaces do not forward Ethernet traffic.
LAN access interfaces	You can use these interfaces to connect to other servers, storage, and switches. When you power on a QFX Series product and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports.
Multichassis aggregated Ethernet (MC-AE) interfaces	You can group a LAG on one QFX3500 standalone switch with a LAG on another QFX3500 standalone switch to create a MC-AE. The MC-AE provides load balancing and redundancy across the two QFX3500 standalone switches.
Tagged-access mode interfaces	You can use tagged-access interfaces to connect a switch to an access layer device. Tagged-access interfaces can accept VLAN-tagged packets from multiple VLANs.
Trunk interfaces	You can use trunk interfaces to connect to other switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the switches or routers must also be configured for trunk mode. In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.

Special Interfaces

Table 229 on page 2416 lists the types of special interfaces supported on the QFX Series.

Table 229: Special Interface Types and Purposes

Type	Purpose
Console port	Each QFX Series product has a serial port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Ethernet Interfaces	You can configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches. You can configure 40-Gigabit data plane uplink ports to connect a Node device to an Interconnect device.
Loopback interface	All QFX Series products have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	The Juniper Networks Junos OS for the QFX Series includes management Ethernet interfaces. The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.
Routed VLAN interface (RVI)	<p>QFX Series products use a Layer 3 routed VLAN interface (RVI) named vlan to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The RVI functions as a logical router, eliminating the need for having both a switch and a router. The RVI (the vlan interface) must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.</p>

Related Documentation

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2417](#)
- [Understanding Interface Naming Conventions on page 2419](#)
- [Understanding Layer 3 Logical Interfaces on page 2434](#)
- [Understanding Management Interfaces on page 2434](#)
- [Understanding Routed VLAN Interfaces on page 1875](#)
- [Overview of Fibre Channel on the QFX Series on page 4956](#)

Understanding Aggregated Ethernet Interfaces and LACP

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a *link aggregation group (LAG)* or *bundle*.

Link aggregation is used to aggregate Ethernet interfaces between two devices. You can create a LAG between a QFX Series product and a router, switch, aggregation switch, server, or other devices. The aggregated Ethernet interfaces that participate in a LAG are called member links. Because a LAG is composed of multiple member links, even if one member link fails, the LAG continues to carry traffic over the remaining links.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard and is used as a discovery protocol.



NOTE: To ensure load balancing across the aggregated Ethernet (AE) interfaces on a redundant server Node group, the members of the AE must be equally distributed across the redundant server Node group.

- [Link Aggregation Group on page 2417](#)
- [Link Aggregation Control Protocol \(LACP\) on page 2418](#)

Link Aggregation Group

To create a LAG:

1. Create a logical aggregated Ethernet interface.
2. Define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and Link Aggregation Control Protocol (LACP).
3. Define the member links to be contained within the aggregated Ethernet interface—for example, two 10-Gigabit Ethernet interfaces.
4. Configure LACP for link detection.

Keep in mind these hardware and software guidelines:

- Up to 32 Ethernet interfaces can be grouped to form a LAG on a QFX3500 switch, a redundant server Node group, a server Node group, and a network Node group on a QFabric system.



NOTE: If you try to commit a configuration containing more than 32 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 32 has been exceeded, and the configuration checkout has failed.

- Up to 63 LAGs are supported on a QFX3500 switch.
- Up to 48 LAGs are supported on redundant server Node groups and server Node groups on a QFabric system, and up to 128 LAGs are supported on network Node groups on a QFabric system. You can configure LAGs across Node devices in redundant server Node groups, server Node groups, and network Node groups.
- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.



NOTE: Junos OS for the QFX Series assigns a unique ID and port priority to each port. The ID and priority are not configurable.

Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure Ethernet links to actively transmit protocol data units (PDUs), or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the LAG without user intervention.
- Link monitoring to check whether both ends of the bundle are connected to the correct group.

When a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server may not be able to exchange LACP PDUs. In such a situation you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when

the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When there are no received PDUs, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

Related Documentation

- [Configuring Link Aggregation on page 2537](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
- [Verifying the Status of a LAG Interface on page 2630](#)
- *Junos® OS Network Interfaces*

Understanding Interface Naming Conventions

The QFX Series use a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series.

This topic describes:

- [Physical Part of an Interface Name on page 2419](#)
- [Logical Part of an Interface Name on page 2431](#)
- [Wildcard Characters in Interface Names on page 2432](#)

Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

device-name:type-fpc/pic/port

QFX Series products apply this convention as follows:

- *device-name*—(QFabric systems only) The *device-name* is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
- *type*—The QFX Series interfaces use the following media types:
 - **fc**—Fibre Channel interface
 - **ge**—Gigabit Ethernet interface
 - **xe**—10-Gigabit Ethernet interface
 - **xle**—40-Gigabit Ethernet interface

- **fte**—40-Gigabit data plane uplink interface
- **me**—Management interface
- **fpc**—Flexible PIC Concentrator. QFX Series interfaces use the following convention for the FPC number in interface names:
 - On QFX3500 and QFX3600 devices, the FPC number is always 0.
 - The FPC number indicates the slot number of the line card that contains the physical interface.
- **pic**—QFX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
 - On a QFX3500 standalone switch, PIC 0 can support up to 48 ports, PIC 1 can support up to 15 10-Gigabit Ethernet ports, and PIC 2 can support up to 4 40-Gigabit Ethernet ports.
 - On a QFX3500 Node device, PIC 0 can support up to 48 ports and PIC 1 can support up to four 40-Gigabit data plane uplink ports.
 - On a QFX3600 Node device, PIC 0 can support up to 56 10-Gigabit Ethernet ports, and PIC 1 can support up to 8 40-Gigabit data plane uplink ports and up to 14 40-Gigabit Ethernet ports.
 - On a QFX3600 standalone switch, PIC 0 can support up to 63 10-Gigabit Ethernet ports, and PIC 1 can support up to 16 40-Gigabit Ethernet ports.
- **port**—QFX Series interfaces use the following convention for port numbers:
 - The QFX3500 device has 48 network access ports labeled 0 through 47 on PIC 0.
 - The QFX3500 device has 15 network access ports labeled 1 through 15 on PIC 1.
 - The QFX3500 device has four 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 2.
 - On a QFX3500 Node device, you can use the QSFP+ ports to connect the Node device to Interconnect devices.
 - On a QFX3500 standalone switch, by default, the 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. Optionally, you can choose to configure the QSFP+ ports as 40-Gigabit Ethernet ports (see [“Configuring the QSFP+ Port Type on QFX3500 Standalone Switches”](#) on page 2543).
 - The QFX3600 device has 16 40-Gbps QSFP+ ports labeled Q0 through Q15.
 - On a QFX3600 Node device, by default, four 40-Gbps QSFP+ ports (labeled Q0 through Q3) are configured for uplink connections between your Node device and your Interconnect devices, and twelve 40-Gbps QSFP+ ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support up to 48 10-Gigabit Ethernet ports for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between

your Node device and your Interconnect devices, and ports Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see [“Configuring the Port Type on QFX3600 Node Devices” on page 1367](#)).

- On a QFX3600 standalone switch, by default, all the QSFP+ ports are configured to operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the QSFP+ ports as 10-Gigabit Ethernet ports (see [“Configuring the Port Type on QFX3600 Standalone Switches” on page 2545](#)) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.



NOTE: Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.

port ranges— QFX Series devices can have different port ranges depending on the media type of the interface.

- On a QFX3500 device (see [Table 230 on page 2422](#) for physical port to logical port mappings):
 - The valid port range for a Fibre Channel (fc) interface is 0 through 5 and 42 through 47 on PIC 0, with a total of 12 available Fibre Channel ports.
 - The valid port range for a Gigabit Ethernet (ge) interface is 6 through 41 on PIC 0 because the ports between 0 and 5 and 42 and 47 are reserved as Fibre Channel ports. The total number of available Gigabit Ethernet ports is 36, because 12 of the remaining 48 ports are reserved for Fibre Channel and 10-Gigabit Ethernet interfaces. Fibre Channel ports cannot be configured as Gigabit Ethernet ports.
 - The valid port range for a 10-Gigabit Ethernet (xe) interface is 0 through 47 on PIC 0. The valid port range for a 10-Gigabit Ethernet (xe) interface is 1 through 15 on PIC 1. The total number of available 10-Gigabit Ethernet ports is 63.



NOTE: This port range is available on a QFX3500 standalone switch only.

- The valid port range for a 40-Gigabit Ethernet (xle) interface is 0 through 3 on PIC 2. There are four available ports.



NOTE: This port range is available on a QFX3500 standalone switch only.

- On a QFX3600 standalone switch (see [Table 231 on page 2426](#) for physical port to logical port mappings):

- The valid port range for a 10-Gigabit Ethernet (**xe**) interface is 1 through **63** on PIC **0**. There are 63 available ports.
- The valid port range for a 40-Gigabit Ethernet (**xle**) interface is **0** through **15** on PIC **1**. There are 16 available ports.
- On a QFX3600 Node device (see [Table 232 on page 2429](#) for physical port to logical port mappings):
 - The valid port range for a 10-Gigabit Ethernet (**xe**) interface is **8** through **63** on PIC **0**. There are 56 available ports.
 - The valid port range for a 40-Gigabit Ethernet (**xle**) interface is **2** through **15** on PIC **1**. There are 14 available ports.
 - The valid port range for a 40-Gigabit data plane uplink (**fte**) interface is **0** through **7** on PIC **1**. There are eight available ports.
- The valid port range for a management interface (**me**) is 0 through 6, with a total of seven available ports. However, you can only configure me0 and me1.
- On the QFX3500 standalone switch and QFX3500 Node device, the management interfaces are labeled **C0** and **C1**, and they correspond to me0 and me1.
- On the QFX3600 device, there are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**), and they correspond to me0 and me1.



NOTE: On a QFX3600 device, you can use either the RJ-45 or the SFP management ports, but not both at the same time.

Table 230: Valid Port Ranges on QFX3500 Devices

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
0	fc-0/0/0	Not supported on this port	xe-0/0/0	Not supported on this port
1	fc-0/0/1	Not supported on this port	xe-0/0/1	Not supported on this port
2	fc-0/0/2	Not supported on this port	xe-0/0/2	Not supported on this port
3	fc-0/0/3	Not supported on this port	xe-0/0/3	Not supported on this port
4	fc-0/0/4	Not supported on this port	xe-0/0/4	Not supported on this port

Table 230: Valid Port Ranges on QFX3500 Devices (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
5	fc-0/0/5	Not supported on this port	xe-0/0/5	Not supported on this port
6	Not supported on this port	ge-0/0/6	xe-0/0/6	Not supported on this port
7	Not supported on this port	ge-0/0/7	xe-0/0/7	Not supported on this port
8	Not supported on this port	ge-0/0/8	xe-0/0/8	Not supported on this port
9	Not supported on this port	ge-0/0/9	xe-0/0/9	Not supported on this port
10	Not supported on this port	ge-0/0/10	xe-0/0/10	Not supported on this port
11	Not supported on this port	ge-0/0/11	xe-0/0/11	Not supported on this port
12	Not supported on this port	ge-0/0/12	xe-0/0/12	Not supported on this port
13	Not supported on this port	ge-0/0/13	xe-0/0/13	Not supported on this port
14	Not supported on this port	ge-0/0/14	xe-0/0/14	Not supported on this port
15	Not supported on this port	ge-0/0/15	xe-0/0/15	Not supported on this port
16	Not supported on this port	ge-0/0/16	xe-0/0/16	Not supported on this port
17	Not supported on this port	ge-0/0/17	xe-0/0/17	Not supported on this port
18	Not supported on this port	ge-0/0/18	xe-0/0/18	Not supported on this port
19	Not supported on this port	ge-0/0/19	xe-0/0/19	Not supported on this port
20	Not supported on this port	ge-0/0/20	xe-0/0/20	Not supported on this port

Table 230: Valid Port Ranges on QFX3500 Devices (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
21	Not supported on this port	ge-0/0/21	xe-0/0/21	Not supported on this port
22	Not supported on this port	ge-0/0/22	xe-0/0/22	Not supported on this port
23	Not supported on this port	ge-0/0/23	xe-0/0/23	Not supported on this port
24	Not supported on this port	ge-0/0/24	xe-0/0/24	Not supported on this port
25	Not supported on this port	ge-0/0/25	xe-0/0/25	Not supported on this port
26	Not supported on this port	ge-0/0/26	xe-0/0/26	Not supported on this port
27	Not supported on this port	ge-0/0/27	xe-0/0/27	Not supported on this port
28	Not supported on this port	ge-0/0/28	xe-0/0/28	Not supported on this port
29	Not supported on this port	ge-0/0/29	xe-0/0/29	Not supported on this port
30	Not supported on this port	ge-0/0/30	xe-0/0/30	Not supported on this port
31	Not supported on this port	ge-0/0/31	xe-0/0/31	Not supported on this port
32	Not supported on this port	ge-0/0/32	xe-0/0/32	Not supported on this port
33	Not supported on this port	ge-0/0/33	xe-0/0/33	Not supported on this port
34	Not supported on this port	ge-0/0/34	xe-0/0/34	Not supported on this port
35	Not supported on this port	ge-0/0/35	xe-0/0/35	Not supported on this port
36	Not supported on this port	ge-0/0/36	xe-0/0/36	Not supported on this port

Table 230: Valid Port Ranges on QFX3500 Devices (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
37	Not supported on this port	ge-0/0/37	xe-0/0/37	Not supported on this port
38	Not supported on this port	ge-0/0/38	xe-0/0/38	Not supported on this port
39	Not supported on this port	ge-0/0/39	xe-0/0/39	Not supported on this port
40	Not supported on this port	ge-0/0/40	xe-0/0/40	Not supported on this port
41	Not supported on this port	ge-0/0/41	xe-0/0/41	Not supported on this port
42	fc-0/0/42	Not supported on this port	xe-0/0/42	Not supported on this port
43	fc-0/0/43	Not supported on this port	xe-0/0/43	Not supported on this port
44	fc-0/0/44	Not supported on this port	xe-0/0/44	Not supported on this port
45	fc-0/0/45	Not supported on this port	xe-0/0/45	Not supported on this port
46	fc-0/0/46	Not supported on this port	xe-0/0/46	Not supported on this port
47	fc-0/0/47	Not supported on this port	xe-0/0/47	Not supported on this port
Q0	Not supported on this port	Not supported on this port	xe-0/1/1 xe-0/1/2 xe-0/1/3 NOTE: Supported on QFX3500 standalone switch only.	xle-0/2/0 NOTE: Supported on QFX3500 standalone switch only.

Table 230: Valid Port Ranges on QFX3500 Devices (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
Q1	Not supported on this port	Not supported on this port	xe-0/1/4 xe-0/1/5 xe-0/1/6 xe-0/1/7 NOTE: Supported on QFX3500 standalone switch only.	xle-0/2/1 NOTE: Supported on QFX3500 standalone switch only.
Q2	Not supported on this port	Not supported on this port	xe-0/1/8 xe-0/1/9 xe-0/1/10 xe-0/1/11 NOTE: Supported on QFX3500 standalone switch only.	xle-0/2/2 NOTE: Supported on QFX3500 standalone switch only.
Q3	Not supported on this port	Not supported on this port	xe-0/1/12 xe-0/1/13 xe-0/1/14 xe-0/1/15 NOTE: Supported on QFX3500 standalone switch only.	xle-0/2/3 NOTE: Supported on QFX3500 standalone switch only.

Table 231: Valid Port Ranges on QFX3600 Standalone Switches

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	xe-0/0/1 xe-0/0/2 xe-0/0/3	xle-0/1/0

Table 231: Valid Port Ranges on QFX3600 Standalone Switches (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q1	xe-0/0/4 xe-0/0/5 xe-0/0/6 xe-0/0/7	xle-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	xle-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	xle-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	xle-0/1/4
Q5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23	xle-0/1/5
Q6	xe-0/0/24 xe-0/0/25 xe-0/0/26 xe-0/0/27	xle-0/1/6

Table 231: Valid Port Ranges on QFX3600 Standalone Switches (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q7	xe-0/0/28	xle-0/1/7
	xe-0/0/29	
	xe-0/0/30	
	xe-0/0/31	
Q8	xe-0/0/32	xle-0/1/8
	xe-0/0/33	
	xe-0/0/34	
	xe-0/0/35	
Q9	xe-0/0/36	xle-0/1/9
	xe-0/0/37	
	xe-0/0/38	
	xe-0/0/39	
Q10	xe-0/0/40	xle-0/1/10
	xe-0/0/41	
	xe-0/0/42	
	xe-0/0/43	
Q11	xe-0/0/44	xle-0/1/11
	xe-0/0/45	
	xe-0/0/46	
	xe-0/0/47	
Q12	xe-0/0/48	xle-0/1/12
	xe-0/0/49	
	xe-0/0/50	
	xe-0/0/51	

Table 231: Valid Port Ranges on QFX3600 Standalone Switches (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q13	xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55	xle-0/1/13
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	xle-0/1/14
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	xle-0/1/15

Table 232: Valid Port Ranges on QFX3600 Node Devices

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	Not supported on this port	fte-0/1/0	Not supported on this port
Q1	Not supported on this port	fte-0/1/1	Not supported on this port
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	fte-0/1/2	xle-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	fte-0/1/3	xle-0/1/3

Table 232: Valid Port Ranges on QFX3600 Node Devices (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q4	xe-0/0/16	fte-0/1/4	xle-0/1/4
	xe-0/0/17		
	xe-0/0/18		
	xe-0/0/19		
Q5	xe-0/0/20	fte-0/1/5	xle-0/1/5
	xe-0/0/21		
	xe-0/0/22		
	xe-0/0/23		
Q6	xe-0/0/24	fte-0/1/6	xle-0/1/6
	xe-0/0/25		
	xe-0/0/26		
	xe-0/0/27		
Q7	xe-0/0/28	fte-0/1/7	xle-0/1/7
	xe-0/0/29		
	xe-0/0/30		
	xe-0/0/31		
Q8	xe-0/0/32	Not supported on this port	xle-0/1/8
	xe-0/0/33		
	xe-0/0/34		
	xe-0/0/35		
Q9	xe-0/0/36	Not supported on this port	xle-0/1/9
	xe-0/0/37		
	xe-0/0/38		
	xe-0/0/39		

Table 232: Valid Port Ranges on QFX3600 Node Devices (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q10	xe-0/0/40	Not supported on this port	xle-0/1/10
	xe-0/0/41		
	xe-0/0/42		
	xe-0/0/43		
Q11	xe-0/0/44	Not supported on this port	xle-0/1/11
	xe-0/0/45		
	xe-0/0/46		
	xe-0/0/47		
Q12	xe-0/0/48	Not supported on this port	xle-0/1/12
	xe-0/0/49		
	xe-0/0/50		
	xe-0/0/51		
Q13	xe-0/0/52	Not supported on this port	xle-0/1/13
	xe-0/0/53		
	xe-0/0/54		
	xe-0/0/55		
Q14	xe-0/0/56	Not supported on this port	xle-0/1/14
	xe-0/0/57		
	xe-0/0/58		
	xe-0/0/59		
Q15	xe-0/0/60	Not supported on this port	xle-0/1/15
	xe-0/0/61		
	xe-0/0/62		
	xe-0/0/63		

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.)

separates the port and logical unit numbers: *device-name* (QFabric systems only):*type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
node-device1:xe-0/0/1.0	down	remote-analyzer	unblocked
node-device1:xe-0/0/2.0	down	default	unblocked
node-device1:xe-0/0/3.0	down	default	unblocked

When you configure aggregated Ethernet interfaces, you configure a logical interface, which is called a *bundle* or a LAG. Each LAG can include up to eight Ethernet interfaces, depending on the switch model.

Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (*) in quotation marks (" ").

Related Documentation

- [Interfaces Overview on page 2415](#)
- [Rear Panel of a QFX3500 Device](#)
- [Front Panel of a QFX3600 Device](#)
- [Junos® OS Network Interfaces](#)

Understanding Interface Ranges on the QFX Series

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces on Juniper Networks QFX Series products. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations is also a valid definition.



NOTE: The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** statement is used in the following configuration hierarchies:

- ethernet-switching-options analyzer *name* input egress interface
- ethernet-switching-options analyzer *name* input ingress interface
- ethernet-switching-options analyzer output interface
- ethernet-switching-options bpdu-block interface
- ethernet-switching-options interfaces
- ethernet-switching-options redundant-trunk-group group-name interface
- ethernet-switching-options secure-access-port interface
- ethernet-switching-options voip interface
- protocols igmp-snooping vlan *vlan-name* interface
- protocols isis interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management te-link *name* interface
- protocols lldp interface
- protocols mstp interface
- protocols mstp msti-*id* interface
- protocols mstp msti-*id* vlan *vlan-id* interface
- protocols sflow interfaces
- protocols stp interface
- protocols vstp vlan *vlan-id* interface
- vlans *vlan-name* interface

**Related
Documentation**

- [Interfaces Overview on page 2415](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533](#)
- [Configuring Link Aggregation on page 2537](#)
- [Configuring a Layer 3 Logical Interface on page 2537](#)
- [Junos® OS Network Interfaces](#)
- [interface-range on page 2582](#)

Understanding Layer 3 Logical Interfaces

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks QFX3500 Switch to a Layer 2 switch. Only one physical connection is required between the switches. You can also use Layer 3 logical interfaces to provide alternative gateway addresses for smart DHCP relay.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series systems support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. Double-tagging, which is assigning more than one VLAN ID in the frame header, is not supported. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.

Related Documentation

- [Interfaces Overview on page 2415](#)
- [Configuring a Layer 3 Logical Interface on page 2537](#)
- [Configuring DHCP and BOOTP Relay on page 4907](#)
- *Junos® OS Network Interfaces*

Understanding Management Interfaces

You use management interfaces to access devices remotely. Typically, a management interface is not connected to the in-band network, but is connected to a device in the internal network. Through a management interface, you can access the device over the network using utilities such as **ssh** and **telnet** and configure it from anywhere, regardless of its physical location. As a security feature, users cannot log in as **root** through a management interface. To access the device as **root**, you must use the console port.



NOTE: Juniper Networks does not support configuring two management interfaces in the same subnet.

Management interface port ranges vary based on device type:

- QFX3500 devices:

The valid port range for a management interface (**me**) on a QFX3500 device is between 0 and 6, with a total of seven available ports. However, you can only configure **me0** and **me1** as management interfaces. The management interfaces are labeled **C0** and **C1**, and they correspond to **me0** and **me1**.

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**.



NOTE: On a QFX3600 device, you can use either the RJ-45 or the SFP management interfaces, but not both at the same time.



NOTE: Before you can use the management interfaces on either the QFX3500 device or QFX3600 device, you must configure the logical interfaces with valid IP addresses. Juniper Networks does not support configuring two management interfaces in the same subnet.

- QFabric system:

On a QFabric system, there are management interfaces on the Node devices, Interconnect devices, and Director devices. However, you cannot access the management interfaces on the Node devices or Interconnect devices directly. You can only manage and configure these devices using the Director device. You can connect to the management interface over the network using utilities such as SSH.

Related Documentation

- [Interfaces Overview on page 2415](#)

Understanding Multichassis Link Aggregation

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network are necessary. Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and although it scales to very large network environments, it still only provides one active path from one device to another, regardless of how many actual connections might exist in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur. The Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (QFX3500 and QFX3600 devices). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Interchassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.

See [Table 233 on page 2437](#) for information about ICCP failure scenarios.

The following sections provide an overview of the terms and features associated with MC-LAG:

- [Active-Active Mode on page 2437](#)
- [ICCP and ICL-PL on page 2437](#)
- [Failure Handling on page 2437](#)
- [Multichassis Link Protection on page 2438](#)
- [MC-LAG Packet Forwarding on page 2438](#)
- [Layer 3 Routing on page 2438](#)
- [Spanning Tree Protocol \(STP\) Guidelines on page 2439](#)
- [MC-LAG Upgrade Guidelines on page 2439](#)
- [Layer 2 Unicast Features Supported on page 2440](#)
- [Layer 2 Multicast Features Supported on page 2440](#)
- [IGMP Snooping on an Active-Active MC-LAG on page 2440](#)
- [Layer 3 Unicast Features Supported on page 2441](#)
- [VRRP Active-Standby Support on page 2441](#)
- [Routed VLAN Interface \(RVI\) MAC Address Synchronization on page 2442](#)
- [Address Resolution Protocol \(ARP\) on page 2442](#)

- [DHCP Relay with Option 82 on page 2442](#)
- [Private VLAN \(PVLAN\) on page 2443](#)
- [Layer 3 Multicast on page 2443](#)

Active-Active Mode

In active-active mode, all member links are active on the MC-LAG. In this mode, MAC addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is the only mode supported at this time.

ICCP and ICL-PL

ICCP replicates control traffic and forwarding states across the MC-LAG peers and communicates the operational state of the MC-LAG members. Because ICCP uses TCP/IP to communicate between the peers, the two peers must be connected to each other. ICCP messages exchange MC-LAG configuration parameters and ensure that both peers use the correct LACP parameters.

The interchassis link-protection link (ICL-PL) provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL-PL can be either a 10-Gigabit Ethernet interface or an aggregated Ethernet interface. You can configure only one ICL-PL between the two peers, although you can configure multiple MC-LAGs between them.

Failure Handling

Configuring ICCP adjacency over aggregated links mitigates the possibility of a split-brain state. A split brain state occurs when the ICL-PL configured between the MC-LAG peers goes down. To work around this problem, enable backup liveness detection. With backup liveness enabled, the MC-LAG peers can communicate through the keepalive link.

During a split-brain state, the standby peer brings down local members in the MC-LAG links by changing the LACP system ID. When the ICCP connection is active, both of the MC-LAG peers use the configured LACP system ID. If the LACP system ID is changed during failures, the server that is connected over the MC-LAG removes these links from the aggregated Ethernet bundle.

When the ICL-PL is operationally down and the ICCP connection is active, the LACP state of the links with status control configured as standby is set to the standby state. When the LACP state of the links is changed to standby, the server that is connected over the MC-LAG makes these links inactive and does not use them for sending data.

[Table 233 on page 2437](#) describes the different ICCP failure scenarios. The dash means that the item is not applicable.

Table 233: ICCP Failure Scenarios

ICCP Connection Status	ICL-PL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet (MC-AE) Interface with Status Set to Standby
Down	Down or Up	Not configured	LACP system ID is changed to default value.

Table 233: ICCP Failure Scenarios (*continued*)

ICCP Connection Status	ICL-PL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet (MC-AE) Interface with Status Set to Standby
Down	Down or Up	Active	LACP system ID is changed to default value.
Down	Down or Up	Inactive	No change in LACP system ID.
Up	Down	–	LACP state is set to standby. MUX state moves to waiting state.

Split-brain states bring down the MC-LAG link completely if the primary peer members are also down for other reasons. Recovery from the split-brain state occurs automatically when the ICCP adjacency comes up between the MC-LAG peers.

Multichassis Link Protection

Multichassis link protection provides link protection between the two MC-LAG peers hosting an MC-LAG. If the ICCP connection is up and the ICL-PL comes up, the peer configured as standby brings up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on each MC-LAG peer that is hosting an MC-LAG.

MC-LAG Packet Forwarding

To prevent the server from receiving multiple copies from both of the MC-LAG peers, a block mask is used to prevent forwarding of traffic received on the ICL-PL toward the MC-AE interface. Preventing forwarding of traffic received on the ICL-PL interface toward the MC-AE interface ensures that traffic received on MC-LAG links is not forwarded back to the same link on the other peer. The forwarding block mask for a given MC-LAG link is cleared if all of the local members of the MC-LAG link go down on the peer. To achieve faster convergence, if all local members of the MC-LAG link are down, outbound traffic on the MC-LAG is redirected to the ICL-PL interface on the data plane.

Layer 3 Routing

To provide Layer 3 routing functions to downstream clients, configure the same gateway address on both MC-LAG network peers. To upstream routers, the MC-LAG network peers could be viewed as either equal-cost multipath (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using Virtual Router Redundancy Protocol (VRRP) over routed VLAN interfaces (RVIs). Junos OS also supports active-active MC-LAGs by using RVI MAC address synchronization. You must configure the RVI using the same IP address across MC-LAG peers.

Spanning Tree Protocol (STP) Guidelines

- Enable STP globally.
STP might detect local miswiring loops within the peer or across MC-LAG peers.
STP might not detect network loops introduced by MC-LAG peers.
- Disable STP on ICL-PL links; otherwise, it might block ICL-PL ports and disable protection.
- Disable STP on interfaces that are connected to aggregation switches.
- Do not enable bridge protocol data unit (BPDU) block on interfaces connected to aggregation switches.

For more information about BPDU block, see [“Understanding BPDU Protection for STP, RSTP, and MSTP” on page 1903](#).

MC-LAG Upgrade Guidelines

Upgrade the MC-LAG peers according to the following guidelines. See [“Upgrading Software on QFX3500 and QFX3600 Standalone Switches” on page 132](#) for exact details about how to perform a software upgrade.



NOTE: After a reboot, the MC-AE interfaces come up immediately and might start receiving packets from the server. If routing protocols are enabled, and the routing adjacencies have not been formed, packets might be dropped.

To prevent this scenario, issue the `set interfaces interface-name aggregated-ether-options mc-ae init-delay-time time` to set a time by which the routing adjacencies are formed.

1. Make sure that both of the MC-LAG peers (node1 and node2) are in the active-active state using the following command on any one of the MC-LAG peers:

```
user@switch> show interfaces mc-ae id 1
Member Link           : ae0
Current State Machine's State: mcae active state
Local Status          : active<<<<<<<
Local State           : up
Peer Status           : active<<<<<<<
Peer State            : up
  Logical Interface    : ae0.0
  Topology Type        : bridge
  Local State          : up
  Peer State           : up
  Peer Ip/MCP/State    : 20.1.1.2 ae2.0 up
```

2. Upgrade node1 of the MC-LAG.

When node1 is upgraded it is rebooted, and all traffic is sent across the available LAG interfaces of node2, which is still up. The amount of traffic lost depends on how quickly the neighbor devices detect the link loss and rehash the flows of the LAG.

3. Verify that node1 is running the software you just installed. Issue the **show version** command.
4. Make sure that both nodes of the MC-LAG (node1 and node2) are in the active-active state after the reboot of node1.
5. Upgrade node2 of the MC-LAG.
Repeat step 1 through step 3 to upgrade node2.

Layer 2 Unicast Features Supported

The following Layer 2 unicast features are supported:

- L2 unicast: learning and aging
 - Learned MAC addresses are propagated across MC-LAG peers for all of the VLANs that are spawned across the peers.
 - Aging of MAC addresses occurs when the MAC address is not seen on both of the peers.
 - MAC learning is disabled on the ICL-PL automatically.
 - MAC addresses learned on single-homed links are propagated across all of the VLANs that have MC-LAG links as members.

Layer 2 Multicast Features Supported

The following Layer 2 multicast features are supported:

- L2 multicast: unknown unicast and IGMP snooping
 - Flooding happens on all links across peers if both peers have virtual LAN membership. Only one of the peers forwards traffic on a given MC-LAG link.
 - Known and unknown multicast packets are forwarded across the peers by adding the ICL-PL port as a multicast router port.
 - IGMP membership learned on MC-LAG links is propagated across peers.
 - During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synced with the peer.

IGMP Snooping on an Active-Active MC-LAG

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts. IGMP uses Protocol Independent Multicast (PIM) to route the multicast traffic. PIM uses distribution trees to determine which traffic is forwarded.

In an active-active MC-LAG configuration, IGMP snooping replicates the Layer 2 multicast routes so that each MC-LAG peer has the same routes. If a device is connected to an MC-LAG peer by way of a single-homed interface, IGMP snooping replicates join message to its IGMP snooping peer. If a multicast source is connected to an MC-LAG by way of a Layer 3 device, the Layer 3 device passes this information to the RVI that is configured on the MC-LAG. The first hop DR is responsible for sending the register and register-stop messages for the multicast group. The last hop DR is responsible for sending PIM join and leave messages toward the rendezvous point and source for the multicast group. The routing device with the smallest preference metric forwards traffic on transit LANs.

Configure the ICL-PL interface as a router-facing interface. For the scenario in which traffic arrives by way of a Layer 3 interface, PIM and IGMP must be enabled on the RVI interface configured on the MC-LAG peers.

Layer 3 Unicast Features Supported

The following Layer 3 unicast features are supported:

- VRRP active-standby support enables Layer 3 routing over MC-AE interfaces.
- Routed VLAN interface (RVI) MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI MAC address or its peer's RVI MAC address.
- Address Resolution Protocol (ARP) synchronization enables ARP resolution on both of the MC-LAG peers.
- DHCP Relay with option 82 enables option 82 on the MC-LAG peers. Option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client.

VRRP Active-Standby Support

VRRP in active-standby mode enables Layer 3 routing over the MC-AE interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The virtual routers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address, known as a routed VLAN interface (RVI), maps to either of the VRRP MAC addresses or the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP routers is the master (active), and the other is a backup (standby). Both VRRP active and VRRP backup routers forward Layer 3 traffic arriving on the MC-AE interface. If the master router fails, all the traffic shifts to the MC-AE link on the backup router.



NOTE: You must configure VRRP on both MC-LAG peers in order for both the active and standby members to accept and route packets. Additionally, configure the VRRP backup router to send and receive ARP requests.

Routing protocols run on the primary IP address of the RVI, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address

of the RVI and the RVI MAC address to communicate with the MC-LAG peers. The RVI MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL-PL.

Routed VLAN Interface (RVI) MAC Address Synchronization

Routed VLAN interface (RVI) MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI MAC address or its peer's RVI MAC address. Each MC-LAG peer installs its own RVI MAC address as well as the peer's RVI MAC address in the hardware. Each MC-LAG peer treats the packet as if it were its own packet. If RVI MAC address synchronization is not enabled, the RVI MAC address is installed on the MC-LAG peer as if it was learned on the ICL-PL.



NOTE: If you need routing capability, configure both VRRP and routing protocols on each MC-LAG peer.

Control packets destined for a particular MC-LAG peer that arrive on an MC-AE interface of its MC-LAG peer are not forwarded on the ICL-PL interface. Additionally, using the gateway IP address as a source address when you issue either a ping, traceroute, telnet, or FTP request is not supported.

To enable RVI MAC address synchronization, issue the **set vlan *vlan-name* l3_interface *rvi-name* mcae-mac-synchronize** on each MC-LAG peer. Configure the same IP address on both MC-LAG peers. This IP address is used as the default gateway for the MC-LAG servers or hosts.

Address Resolution Protocol (ARP)

Address Resolution Protocol (ARP) maps IP addresses to MAC addresses. Without synchronization, if one MC-LAG peer sends an ARP request, and the other MC-LAG peer receives the response, ARP resolution is not successful. With synchronization, the MC-LAG peers synchronize the ARP resolutions by sniffing the packet at the MC-LAG peer receiving the ARP response and replicating this to the other MC-LAG peer. This ensures that the entries in ARP tables on the MC-LAG peers are consistent.

When one of the MC-LAG peers restarts, the ARP destinations on its MC-LAG peer are synchronized. Because the ARP destinations are already resolved, its MC-LAG peer can forward Layer 3 packets out of the MC-AE interface.

DHCP Relay with Option 82

DHCP relay with option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client. With DHCP relay enabled, DHCP request packets might take the path to the DHCP server through either of the MC-LAG peers. Because the MC-LAG peers have different host names, chassis MAC addresses, and interface names, you need to observe these requirements when you configure DHCP relay with option 82:

- Use the interface description instead of the interface name.
- Do not use the hostname as part of the circuit ID or remote ID strings.

- Do not use the chassis MAC address as part of the remote ID string.
- Do not enable the vendor ID.
- If the ICL-PL interface receives DHCP request packets, the packets are dropped to avoid duplicate packets in the network.

A counter called *Due to received on ICL interface* has been added to the **show helper statistics** command, which tracks the packets that the ICL-PL interface drops.

An example of the CLI output follows:

```
user@switch> show helper statistics
BOOTP:
  Received packets: 6
  Forwarded packets: 0
  Dropped packets: 6
    Due to no interface in DHCP Relay database: 0
    Due to no matching routing instance: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
    Due to invalid server address: 0
    Due to no valid local address: 0
    Due to no route to server/client: 0
    Due to received on ICL interface: 6
```

The output shows that six packets received on the ICL-PL interface have been dropped.

Private VLAN (PVLAN)

Private VLANs allow you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside of a VLAN. A PVLAN can span multiple peers on an MC-LAG.

When configuring a PVLAN, you must configure the ICL-PL interface as the PVLAN trunk interface for the PVLAN. This is essential for traffic to be switched to the required primary and secondary ports of the PVLAN across the MC-LAG peers.

Layer 3 Multicast

- [PIM Operation With Normal Mode DR Election on page 2443](#)
- [PIM Operation with Dual-DR Mode on page 2444](#)
- [Configuration Guidelines and Caveats on page 2444](#)

Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) provide support for Layer 3 multicast. In addition to the standard mode of PIM operation, there is a special mode called PIM dual DR (designated router). PIM dual DR minimizes traffic loss in case of failures.

PIM Operation With Normal Mode DR Election

In normal mode DR election, the RVI interfaces on both of the MC-LAG peers are configured with PIM enabled. In this mode, one of the MC-LAG peers becomes the DR through the PIM DR election mechanism. The elected DR maintains the rendezvous-point tree (RPT) and shortest-path tree (SPT) so it can receive data from the source device. The elected DR participates in periodic PIM join and prune activities toward the rendezvous point (RP) or the source.

The trigger for initiating these join and prune activities is the IGMP membership reports that are received from interested receivers. IGMP reports received over MC-AE interfaces (potentially hashing on either of the MC-LAG peers) and single-homed links are synchronized to the MC-LAG peer through ICCP.

Both MC-LAG peers receive traffic on their incoming interface (IIF). The non-DR receives traffic by way of the ICL-PL interface, which acts as a multicast router (mrouter) interface.

If the DR fails, the non-DR has to build the entire forwarding tree (RPT and SPT), which can cause multicast traffic loss.

PIM Operation with Dual-DR Mode

In this mode, both of MC-LAG peers act as DRs (active and backup) and send periodic join and prune messages upstream towards the RP, or source, and eventually join the RPT or SPT.

The primary MC-LAG peer forwards the multicast traffic to the receiver devices even if the standby MC-LAG peer has a smaller preference metric.

The standby MC-LAG peer also joins the forwarding tree and receives the multicast data. The standby MC-LAG peer drops the data because it has an empty outgoing interface list (OIL). When the standby MC-LAG peer detects the primary MC-LAG peer failure, it adds the receiver VLAN to the OIL, and starts to forward the multicast traffic

To enable a multicast dual DR, issue the **set protocols pim interface interface-name dual-dr** command on the VLAN interfaces of each MC-LAG peer.

Configuration Guidelines and Caveats

- Configure the IP address on the active MC-LAG peer with a high IP address or a high DR priority. To ensure that the active MC-LAG peer retains the DR membership designation if PIM neighborship with the peer goes down.
- Using Bidirectional Forwarding Detection (BFD) and RVI MAC synchronization together is not supported because ARP fails.
- When using RVI MAC synchronization, make sure that you configure the primary IP address on both MC-LAG peers. Doing this ensures that both MC-LAG peers cannot become assert winners.
- The number of BFD sessions on RVIs with PIM enabled is restricted to 100. Also, If you have more than 100 RVIs configured, do not configure BFD, and make sure that the hello interval is 2 seconds.

Related Documentation

- [Configuring Link Aggregation on page 2537](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2477](#)

- *Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization*
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2514](#)

Overview of Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate this information to the downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure that the traffic of the failed link is not dropped.

This topic describes:


- [Uplink Failure Detection Configuration on page 2445](#)
- [Failure Detection Pair on page 2446](#)

Uplink Failure Detection Configuration

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces bound to the uplink interface. A server that is connected to the disabled downlink interface triggers a network adapter failover to a secondary link to avoid any traffic loss.

[Figure 63 on page 2446](#) illustrates a typical setup for uplink failure detection.

Figure 63: Uplink Failure Detection Configuration on Switches



For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.

The switch can monitor both physical interface links and logical interface links for uplink failures, but you must put the two types of interfaces into separate groups.



NOTE: For logical interfaces, the server must send keepalives between the switch and the server to detect failure of logical links.

Failure Detection Pair

Uplink failure detection requires that you create pairs of uplink and downlink interfaces in a group. Each pair includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplinks the switch monitors. You can configure a maximum of eight uplink interfaces as link-to-monitor interfaces for a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlinks the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlinks to disable in the group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

- Related Documentation**
- [Configuring Interfaces for Uplink Failure Detection on page 2536](#)
 - [Example: Configuring Interfaces for Uplink Failure Detection on page 2449](#)

CHAPTER 29

Configuration

- [Configuration Examples on page 2449](#)
- [Configuration Tasks on page 2532](#)
- [Configuration Statements on page 2546](#)

Configuration Examples

- [Example: Configuring Interfaces for Uplink Failure Detection on page 2449](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2477](#)
- [Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization on page 2500](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2514](#)

Example: Configuring Interfaces for Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate the failure information to the downlink interfaces. All of the network interface cards (NICs) on a server are configured as being either the primary link or the secondary link and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link to ensure that the traffic on the failed link is not dropped.

This example describes:

- [Requirements on page 2450](#)
- [Overview and Topology on page 2450](#)
- [Configuring Uplink Failure Detection on Both Switches on page 2451](#)
- [Verification on page 2452](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.1 or later for the QFX Series
- Two QFX3500 switches
- Two aggregation switches
- One dual-homed server

Overview and Topology

The topology in this example illustrates how to configure uplink failure detection on Switch A and Switch B. Switch A and Switch B are both configured with a link-to-monitor interface (the uplink interface to the aggregation switch) and a link-to-disable interface (the downlink interface to the server). For simplicity, only one group of link-to-monitor interfaces and link-to-disable interfaces is configured for each switch. The server is dual-homed to both Switch A and Switch B. In this scenario, if the link-to-monitor interface to Switch A is disabled, the server uses the link-to-monitor interface to Switch B instead.



NOTE: This example does not describe how to configure the dual-homed server or the aggregation switches. Please refer to the documentation for each of these devices for more information.

Figure 63 on page 2446 illustrates a typical setup for uplink failure detection.

Figure 64: Uplink Failure Detection Configuration on Switches




Table 234 on page 2451 lists uplink failure settings for each QFX3500 switch.

Table 234: Settings for Uplink Failure Protection Example

Switch A	Switch B
<ul style="list-style-type: none"> Group name: Group1 Link-to-monitor interface: xe-0/0/0 Link-to-disable interface: xe-0/0/1 	<ul style="list-style-type: none"> Group name: Group2 Link-to-monitor interface: xe-0/0/0 Link-to-disable interface: xe-0/0/1

Configuring Uplink Failure Detection on Both Switches

To configure uplink failure detection on both switches, perform these tasks:

CLI Quick Configuration

To quickly configure uplink failure protection on Switch A and Switch B, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set uplink-failure-detection group group1
set uplink-failure-detection group group2
set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
set uplink-failure-detection group group1 link-to-disable xe-0/0/1
set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

Step-by-Step Procedure

To configure uplink failure protection on both switches:

- Specify a name for the uplink failure detection group on Switch A:

```
[edit protocols]
user@switch# set uplink-failure-detection group group1
```
- Add an uplink interface to the group on Switch A:

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
```
- Add a downlink interface to the group on Switch A:

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-disable xe-0/0/1
```
- Specify a name for the uplink failure detection group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2
```
- Add an uplink interface to the group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
```
- Add a downlink interface to the group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

Results Display the results of the configuration:

```
uplink-failure-detection {
  group {
    group1 {
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
  }
}
```

```
    }  
  }  
  group2 {  
    link-to-monitor {  
      xe-0/0/0;  
    }  
    link-to-disable {  
      xe-0/0/1;  
    }  
  }  
}
```

Verification

To verify that uplink failure detection is working correctly, perform the following tasks on Switch A and Switch B:

- [Verifying That Uplink Failure Detection is Working Correctly on page 2452](#)

Verifying That Uplink Failure Detection is Working Correctly

Purpose Verify that the switch disables the downlink interface when it detects an uplink failure.

Action 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection  
Group           : group1  
Uplink          : xe-0/0/0*  
Downlink       : xe-0/0/1*  
Failure Action  : Inactive
```



NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]  
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection  
Group           : group1  
Uplink          : xe-0/0/0  
Downlink       : xe-0/0/1  
Failure Action  : Active
```

Meaning The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 2445](#)
 - [Configuring Interfaces for Uplink Failure Detection on page 2536](#)
 - [Verifying That Uplink Failure Detection Is Working Correctly on page 2632](#)

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

A QFX Series product allows you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. You can configure LAGs to connect a QFX Series product to other switches, like aggregation switches, servers, or routers. This example describes how to configure LAGs to connect a QFX3500 switch to an aggregation switch.

- [Requirements on page 2454](#)
- [Overview and Topology on page 2454](#)
- [Configuration on page 2455](#)
- [Verification on page 2457](#)
- [Troubleshooting on page 2457](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX Series
- One QFX3500 switch

Overview and Topology

In this example, the QFX3500 switch has one LAG comprising two 10-Gigabit Ethernet interfaces. This LAG is configured in port mode so that the QFX3500 switch and the VLAN to which it has been assigned can send and receive traffic.

Configuring the Ethernet interfaces as LAGs has the following advantages:

- If one physical port is lost for any reason (a cable is unplugged or a switch port fails), the logical port transparently continues to function over the remaining physical port.
- Link Aggregation Control Protocol (LACP) can optionally be configured for link monitoring and automatic addition and deletion of individual links without user intervention.



NOTE: If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

The topology used in this example consists of one QFX3500 switch with a LAG configured between two of its 10-Gigabit Ethernet interfaces. The QFX3500 switch is connected to an aggregation switch.




Table 235 on page 2455 details the topology used in this configuration example.

Table 235: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch

Hostname	Base Hardware	Trunk Port
switch	QFX3500 switch	ae0 is configured as a trunk port and combines the following two interfaces: xe-0/0/2 and xe-0/0/3

Configuration

To configure a LAG between two 10-Gigabit Ethernet interfaces:

CLI Quick Configuration

To quickly configure a LAG between two 10-Gigabit Ethernet interfaces on a QFX3500 switch, copy the following commands and paste them into the QFX3500 switch terminal window:

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family ethernet-switching vlan members green
set interfaces xe-0/0/2 ether-options 802.ad ae0
set interfaces xe-0/0/3 ether-options 802.ad ae0
```

Step-by-Step Procedure

To configure a LAG between a QFX Series product and an aggregation switch:

- Specify the number of LAGs to be created on the switch:


```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 1
```
- Specify the number of links that need to be present for the ae0 LAG interface to be up:


```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```
- Specify the media speed of the ae0 link:


```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```
- Assign the LAG to a VLAN:


```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members vlan200
```

5. (Optional): Designate one side of the LAG as active for LACP:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
```

6. (Optional): Designate the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp periodic fast
```

Results

Display the results of the configuration:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
green {
  vlan-id 200;
}
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      description yellow;
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members green;
        }
      }
    }
  }
  xe-0/0/2 {
    ether-options {
      802.ad ae0;
    }
  }
  xe-0/0/3 {
    ether-options {
      802.ad ae0;
    }
  }
}
```

Verification

To verify that switching is operational and one LAG has been created, perform these tasks:

- [Verifying That LAG ae0.0 Has Been Created on page 2457](#)
- [Verifying That LAG ae0 Has Been Created on page 2457](#)

Verifying That LAG ae0.0 Has Been Created

Purpose Verify that LAG ae0.0 has been created on the switch.

Action `show interfaces ae0 terse`

Interface	Admin	Link	Proto	Local	Remote
ae0	up	up			
ae0.0	up	up	eth-switch		

Meaning The output confirms that the ae0.0 link is up and shows the **family** and IP address assigned to this link.

Verifying That LAG ae0 Has Been Created

Purpose Verify that LAG ae0 has been created on the switch

Action `show interfaces ae0 terse`

Interface	Admin	Link	Proto	Local	Remote
ae0	up	down			
ae0.0	up	down	eth-switch		

Meaning The output shows that the ae0.0 link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The `show interfaces terse` command shows that the LAG is **down**.

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.

Related Documentation

- [Configuring Link Aggregation on page 2537](#)
- [Verifying the Status of a LAG Interface on page 2630](#)

- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
- [show lacp statistics interfaces \(View\) on page 2806](#)

Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch

QFX Series products allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. On a QFX3500 standalone switch, you can group up to 32 Ethernet interfaces to form a LAG. On a QFabric system, you can group up to 8 Ethernet interfaces to form a LAG. QFX Series products allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 2454](#):

- [Requirements on page 2458](#)
- [Overview and Topology on page 2458](#)
- [Configuring LACP for the LAG on the QFX Series on page 2459](#)
- [Verification on page 2459](#)
- [Troubleshooting on page 2461](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX Series
- One QFX3500 switch

Before you configure LACP, be sure you have:

- Configured the ports on the switches as trunk ports.
- Configured the LAG.

Overview and Topology

The topology in this example is exactly the same as the topology used in the [Configuring a LAG Between a QFX Switch and an Aggregation Switch](#) example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.



NOTE: If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

The interval can be fast (every second) or slow (every 30 seconds).

Configuring LACP for the LAG on the QFX Series

To configure LACP for a QFX Series LAG, perform these tasks:

CLI Quick Configuration To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure To configure LACP for LAG ae0 :

1. Specify the aggregated Ethernet options for the LAG:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active periodic fast
```

Results Display the results of the configuration:

```
[edit interfaces]
user@switch# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Verification

To verify that LACP packets are being exchanged, perform the following tasks:

- [Verifying the LACP Settings on page 2459](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 2460](#)

Verifying the LACP Settings

Purpose Verify that LACP has been set up correctly.

Action Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@switch> show lacp interfaces xe-0/0/2
```

```
Aggregated interface: ae0
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/2	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/0/2	Partner	No	Yes	No	No	No	Yes	Fast	Passive
LACP protocol:	Receive State	Transmit State				Mux State			
xe-0/0/2	Defaulted	Fast periodic				Detached			

Meaning The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged.

Action Use the **show interfaces aex statistics** command to display LACP information.

```
user@switch> show interfaces ae0 statistics
```

```
Physical interface: ae0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 30
Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :          0          0          0          0
  Output:          0          0          0          0
Protocol inet
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged.

Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- [Troubleshooting a Nonworking LACP Link on page 2461](#)

Troubleshooting a Nonworking LACP Link

Problem The LACP link is not working.

Solution Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.
- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface lag-member detail** command.

**Related
Documentation**

- [Configuring Link Aggregation on page 2537](#)
- [Verifying the Status of a LAG Interface on page 2630](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
- [show lacp statistics interfaces \(View\) on page 2806](#)

Example: Configuring Multichassis Link Aggregation

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

The peers in an MC-LAG use an interchassis control link-protection link (ICL-PL) to replicate forwarding information across the peers. The Interchassis Control Protocol (ICCP) exchanges the control information between two MC-LAG switches. Additionally, ICCP propagates the operational state of MC-LAG members through the ICL-PL.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG switches. Each of the switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.

- [Requirements on page 2462](#)
- [Overview on page 2462](#)
- [Configuration on page 2463](#)
- [Verification on page 2473](#)
- [Troubleshooting on page 2476](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2 or later for the QFX Series
- Two switches

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 2454](#).
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch” on page 2458](#).


Overview

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. [Figure 65 on page 2463](#) shows the topology of this example.

Figure 65: Configuring a Multichassis LAG Between Switch A and Switch B



[Table 235 on page 2455](#) details the topology used in this configuration example.

Table 236: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A	QFX3500 switch or QFX3600 switch	ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of ae0 : xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. ae1 is configured as an MC-LAG, and the following two interfaces are part of ae1 : xe-0/0/44 on Switch A and xe-0/0/46 on Switch B.
Switch B	QFX3500 switch or QFX3600 switch	

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network

configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
```

```

set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0

```

Configuring MC-LAG on Two Switches

Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.

```

[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2

```
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```

[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1

```
3. Configure a trunk interface between Switch A and Switch B.

```

[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk

```
4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```

[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0

```

Switch B:

```

[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

```

**Step-by-Step
Procedure**

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



NOTE: Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
```

Switch B:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```

6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

```
[edit vlans]
```

```
user@switch# set v500 vlan-id 500
```

```
[edit vlans]
```

```
user@switch# set v500 l3-interface vlan.500
```

```
[edit interfaces]
```

```
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members v500
```

Step-by-Step Procedure

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.



NOTE: At least one end needs to be active. The other end can be either active or passive.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp active
```

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

3. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

4. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



NOTE: Only active-active mode is supported at this time.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

5. Specify the status control for MC-LAG on Switch A and Switch B.



NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

6. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



NOTE: The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

7. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

8. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

9. Enable a VLAN on the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
```

```
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```


- ```

user@switch# set ae1 unit 0 family ethernet-switching vlan members v100

```
10. (Optional) Enable a private VLAN on the MC-LAG on Switch A and Switch B.
 

```

[edit]
user@switch# set vlans v100 pvlan isolation-vlan-id 200
extend-secondary-vlan-id
[edit]
user@switch# set vlans v100 interface ae0.0 pvlan-trunk

```

### Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.
 

```

[edit]
user@switch# set protocols rstp interface all mode point-to-point

```
2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:
 

```

[edit]
user@switch# set protocols rstp interface ae0.0 disable

```
3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



**NOTE:** The ae1 interface is a downstream interface. This is why RSTP and bpd-block-on-edge need to be configured.

- ```

[edit]
user@switch# set protocols rstp interface ae1.0 edge

```
4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpd-block-on-edge need to be configured.

```

[edit]
user@switch# set protocols rstp bpd-block-on-edge

```

Results

Display the results of the configuration on Switch A.

```

chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {

```

```
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/13 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/44 {
        ether-options {
            802.3ad ae1;
        }
    }
    ae0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members v500;
                }
            }
        }
    }
    ae1 {
        aggregated-ether-options {
            lacp {
                active;
                system-id 00:01:02:03:04:05;
                admin-key 3;
            }
            mc-ae {
                mc-ae-id 3;
                chassis-id 0;
                mode active-active;
                status-control active;
                init-delay-time 240
            }
        }
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members v100;
                }
            }
        }
    }
    vlan {
        unit 500 {
            family inet {
                address 3.3.3.2/24;
            }
        }
    }
}
```

```

protocols {
  iccp {
    local-ip-addr 3.3.3.2;
    peer 3.3.3.1 {
      session-establishment-hold-time 50;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 60;
        transmit-interval {
          minimum-interval 60;
        }
      }
    }
  }
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface all {
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
}
multi-chassis {
  multi-chassis-protection 3.3.3.1 {
    interface ae0;
  }
}
}
vllans {
  v100 {
    vlan-id 100;
  }
  v500 {
    vlan-id 500;
    l3-interface vlan.500;
  }
}
}

```

Display the results of the configuration on Switch B.

```

chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
}
interfaces {
  xe-0/0/12 {
    ether-options {

```

```
        802.3ad ae0;
    }
}
xe-0/0/13 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/46 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members v500;
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            chassis-id 1;
            mode active-active;
            status-control standby;
            init-delay-time 240
        }
    }
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members v100;
            }
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}
protocols {
```

```

iccp {
  local-ip-addr 3.3.3.1;
  peer 3.3.3.2 {
    session-establishment-hold-time 50;
    backup-liveness-detection {
      backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
      minimum-receive-interval 60;
      transmit-interval {
        minimum-interval 60;
      }
    }
  }
}
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface all {
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
}
multi-chassis {
  multi-chassis-protection 3.3.3.2 {
    interface ae0;
  }
}
}
vllans {
  v100 {
    vlan-id 100;
  }
  v500 {
    vlan-id 500;
    l3-interface vllan.500;
  }
}
}

```

Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2474](#)
- [Verifying That ICCP Is Working on Switch B on page 2474](#)
- [Verifying That LACP Is Active on Switch A on page 2474](#)
- [Verifying That LACP Is Active on Switch B on page 2475](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2475](#)

- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2475](#)
- [Verifying that MAC Learning Is Occurring on Switch A on page 2476](#)
- [Verifying that MAC Learning Is Occurring on Switch B on page 2476](#)

Verifying That ICCP Is Working on Switch A

Purpose Verify that ICCP is running on Switch A.

Action [edit]
user@switch# **show iccp**
Redundancy Group Information for peer 3.3.3.1
TCP Connection : Established
Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That ICCP Is Working on Switch B

Purpose Verify that ICCP is running on Switch B.

Action **show iccp**

[edit]
user@switch# **show iccp**
Redundancy Group Information for peer 3.3.3.2
TCP Connection : Established
Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That LACP Is Active on Switch A

Purpose Verify that LACP is active on Switch A.

Action [edit]
user@switch# **show lacp interfaces**
Aggregated interface: ae1
LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity
xe-0/0/46 Actor No No Yes Yes Yes Yes Fast Active
xe-0/0/46 Partner No No Yes Yes Yes Yes Fast Active
LACP protocol: Receive State Transmit State Mux State
xe-0/0/46 Current Fast periodic Collecting distributing

Meaning This output shows that Switch A is participating in LACP negotiation.

Verifying That LACP Is Active on Switch B

Purpose Verify that LACP is active on Switch B

Action [edit]

```
user@switch# show lacp interfaces
```

```
Aggregated interface: ae1
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/44	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/44	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active

LACP protocol:	Receive State	Transmit State	Mux State
xe-0/0/44	Current	Fast periodic	Collecting distributing

Meaning This output shows that Switch B is participating in LACP negotiation.

Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A

Purpose Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

Action [edit]

```
user@switch# show interfaces mc-ae
```

```
Member Link           : ae1
Current State Machine's State: mcae active state
Local Status           : active
Local State             : up
Peer Status             : active
Peer State              : up
Logical Interface       : ae1.0
Topology Type           : bridge
Local State             : up
Peer State              : up
Peer Ip/MCP/State       : 3.3.3.1 ae0.0 up
```

Meaning This output shows that the MC-AE interface on Switch A is up and active.

Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B

Purpose Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

Action [edit]
user@switch# **show interfaces mc-ae**
Member Link : ae1
Current State Machine's State: mcae active state
Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae1.0
Topology Type : bridge
Local State : up
Peer State : up
Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

Meaning This output shows that the MC-AE interface on Switch B is up and active.

Verifying that MAC Learning Is Occurring on Switch A

Purpose Verify that MAC learning is working on Switch A.

Action [edit]
user@switch# **show ethernet-switching table**
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries
VLAN MAC address Type Age Interfaces
V100 * Flood - All-members
V100 00:10:94:00:00:05 Learn(L) 33 ae0.0 (MCAE)

Meaning The output shows four learned MAC addresses entries.

Verifying that MAC Learning Is Occurring on Switch B

Purpose Verify that MAC learning is working on Switch B.

Action [edit]
user@switch# **show ethernet-switching table**
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries
VLAN MAC address Type Age Interfaces
V100 * Flood - All-members
V100 00:10:94:00:00:05 Learn(L) 33 ae0.0 (MCAE)

Meaning The output shows four learned MAC addresses entries.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.

- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related
Documentation**

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

- [Requirements on page 2477](#)
- [Overview on page 2477](#)
- [Configuration on page 2479](#)
- [Verification on page 2499](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two QFX3500 or QFX3600 switches

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 2454](#).
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch” on page 2458](#).

Overview

In this example, you configure two MC-LAGs across two switches, consisting of two aggregated Ethernet interfaces (ae1 and ae2). To support the MC-LAG, create a third aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.



NOTE: Layer 3 connectivity is required for ICCP.


To complete the configuration, enable VRRP by completing the following steps for each MC-LAG:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group
- Configure Layer 3 connectivity between the VRRP groups

Topology

The topology used in this example consists of two switches hosting two MC-LAGs—ae1 and ae2. The two switches are connected to a multicast source (Server 1) over the MC-LAG ae1, and a multicast receiver (Server 2) over the MC-LAG ae2. [Figure 66 on page 2478](#) shows the topology of this example.

Figure 66: Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP



[Table 237 on page 2479](#) details the topology used in this configuration example.

Table 237: Components of the Topology for Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A	QFX3500 or QFX3600 switch	<ul style="list-style-type: none"> ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following two interfaces are part of ae0: xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. ae1 is configured as an MC-LAG for the multicast source (Server 1), and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B. ae2 is configured as an MC-LAG for the multicast receiver (Server 2), and the following two interfaces are part of ae2: xe-0/0/45 on Switch A and xe-0/0/47 on Switch B.
Switch B		

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.

```

set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces xe-0/0/45 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk

```

```
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 virtual-address
  10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority 200
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 200
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 600
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
```

set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces xe-0/0/47 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-aether-options mc-aether-id 3
set interfaces ae1 aggregated-ether-options mc-aether-options mc-aether-chassis-id 1
set interfaces ae1 aggregated-ether-options mc-aether-options mc-aether-mode active-active
set interfaces ae1 aggregated-ether-options mc-aether-options mc-aether-status-control standby
set interfaces ae1 aggregated-ether-options mc-aether-options mc-aether-init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-aether-options mc-aether-id 4
set interfaces ae2 aggregated-ether-options mc-aether-options mc-aether-chassis-id 1
set interfaces ae2 aggregated-ether-options mc-aether-options mc-aether-mode active-active
set interfaces ae2 aggregated-ether-options mc-aether-options mc-aether-status-control active
set interfaces ae2 aggregated-ether-options mc-aether-options mc-aether-init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 virtual-address
10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority 150
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 virtual-address
10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority 150
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
set protocols igmp-snooping vlan all

```

```
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 100
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 500
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Configuring MC-LAG for Layer 3 Multicast Using VRRP on Two Switches

Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 3
```
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

Switch A and Switch B

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
```

Switch A

```
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
user@switch# set xe-0/0/45 ether-options 802.3ad ae2
```

Switch B

```
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
user@switch# set xe-0/0/47 ether-options 802.3ad ae2
```

3. Configure ae0 as the trunk interface between Switch A and Switch B.

Switch A and Switch B

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

4. Configure ae0 as the multichassis protection link between Switch A and Switch B.

Switch A

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

Switch B

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address, minimum receive interval, and minimum transmit interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
```

```
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval  
minimum-interval 60
```

Switch B

```
[edit protocols]  
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval  
60  
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval  
minimum-interval 60
```

3. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



NOTE: Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

Switch A

```
[edit protocols]  
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B

```
[edit protocols]  
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

4. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A

```
[edit protocols]  
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip  
10.207.64.233
```

Switch B

```
[edit protocols]  
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip  
10.207.64.232
```

5. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

Switch A and Switch B

```
[edit vlans]
```



```

user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface vlan.500

[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v500

```

Switch A

```

[edit interfaces]
user@switch# set vlan unit 500 family inet address 3.3.3.2/24

```

Switch B

```

[edit interfaces]
user@switch# set vlan unit 500 family inet address 3.3.3.1/24

```

Step-by-Step Procedure

To enable the ae1 and ae2 MC-LAG interfaces:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interfaces on Switch A and Switch B.



NOTE: At least one end needs to be active. The other end can be either active or passive.

```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active
user@switch# set ae2 aggregated-ether-options lacp active

```

2. Specify the same multichassis aggregated Ethernet (MC-AE) identification number for each MC-LAG peer on Switch A and Switch B.

```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 4

```

3. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A

```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 0

```

Switch B

```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 1

```

4. Specify the operating mode of the MC-LAGs on both Switch A and Switch B.



NOTE: Only active-active mode is supported at this time.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

```
user@switch# set ae2 aggregated-ether-options mc-ae mode active-active
```

5. Specify the status control for the MC-LAGs on Switch A and Switch B.



NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAGs. If one peer is in active mode, the other must be in standby mode.

Switch A

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

```
user@switch# set ae2 aggregated-ether-options mc-ae status-control active
```

Switch B

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

```
user@switch# set ae2 aggregated-ether-options mc-ae status-control standby
```

6. Specify the number of seconds by which the bring-up of the MC-LAG interfaces should be deferred after you reboot Switch A or Switch B.



NOTE: The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 420
```

```
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 420
```

7. Specify the same LACP system ID for each MC-LAG on Switch A and Switch B.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

```
user@switch# set ae2 aggregated-ether-options lacp system-ID 00:01:02:03:04:06
```

8. Specify the same LACP administration key on both Switch A and Switch B.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

```
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
```

9. Enable a VLAN for each MC-LAG on Switch A and Switch B.

[edit vlans]

```

user@switch# set v100 vlan-id 100
user@switch# set v200 vlan-id 200

[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
user@switch# set ae2 unit 0 family ethernet-switching vlan members v200

```

10. Configure ae1 and ae2 as trunk interfaces between Switch A and Switch B.

```

[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching port-mode trunk

```

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on Switch A and Switch B:

1. Create a routed VLAN interface (RVI) for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP groups, and assign an individual IP address for each switch in the VRRP groups.

Switch A

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2
virtual-address 10.1.1.2

```

Switch B

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
virtual-address 10.1.1.2

```

2. Assign the priority for each switch in the VRRP groups:



NOTE: The switch configured with the highest priority is the master.

Switch A

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority
200
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority
200

```

Switch B

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority
150
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority
150

```

3. Enable the switch to accept all packets destined for the virtual IP address if it is the master in a VRRP group:

Switch A

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2
accept-data
```

Switch B

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1
accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
accept-data
```

4. Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set v100 l3-interface vlan.100
user@switch# set v200 l3-interface vlan.200
```

**Step-by-Step
Procedure**

To enable IGMP snooping:

1. Enable IGMP snooping for all VLANs on Switch A and Switch B.

```
[edit protocols]
user@switch# set igmp-snooping vlan all
```

**Step-by-Step
Procedure**

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on Switch A and Switch B.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100
user@switch# set interface vlan.200
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
threshold 500
```

```

user@switch# set interface vlan.200 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
threshold 500

```

Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address on Switch A and Switch B.

```

[edit protocols pim]
user@switch# set rp static address 1.0.0.3

```
2. Configure the address ranges of the multicast groups for which Switch A and Switch B can be a rendezvous point (RP).

```

[edit protocols pim rp static address 1.0.0.3]
user@switch# set group-ranges 239.0.0.0/8

```
3. Enable PIM on the VLAN interfaces for the MC-LAGs on Switch A and Switch B.

```

[edit protocols pim]
user@switch# set interface vlan.100 dual-dr
user@switch# set interface vlan.200 dual-dr

```
4. Configure each PIM interface's priority for being selected as the designated router (DR).

An interface with a higher priority value has a higher probability of being selected as the DR.

Switch A

```

[edit protocols pim]
user@switch# set interface vlan.100 priority 200
user@switch# set interface vlan.200 priority 600

```

Switch B

```

[edit protocols pim]
user@switch# set interface vlan.100 priority 100
user@switch# set interface vlan.200 priority 500

```

5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on Switch A and Switch B.

```

[edit protocols pim]
user@switch# set interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
threshold 500
user@switch# set interface vlan.200 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
minimum-interval 350

```

```
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
threshold 500
```

Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit protocols rstp]
user@switch# set interface all mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B.

```
[edit protocols rstp]
user@switch# set interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



NOTE: The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpdv-block-on-edge need to be configured.

```
[edit protocols rstp]
user@switch# set interface ae1.0 edge
user@switch# set interface ae2.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



NOTE: The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpdv-block-on-edge need to be configured.

```
[edit protocols rstp]
user@switch# set bpdv-block-on-edge
```

Results

From configuration mode on Switch A, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 3;
    }
  }
}
```

```
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
  xe-0/0/45 {
    ether-options {
      802.3ad ae2;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
  ae2 {
```

```
aggregated-ether-options {
  lacp {
    active;
    system-id 00:01:02:03:04:06;
    admin-key 3;
  }
  mc-ae {
    mc-ae-id 4;
    chassis-id 0;
    mode active-active;
    status-control active;
    init-delay-time 240;
  }
}
unit 0 {
  family ethernet-switching {
    port-mode trunk;
    vlan {
      members v200;
    }
  }
}
vlan {
  unit 100 {
    family inet {
      address 10.1.1.11/24 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.21/24 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.2/24;
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
```



```

interface vlan.100 {
  bfd-liveness-detection {
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  bfd-liveness-detection {
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 200;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  priority 600;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
}
iccp {
  local-ip-addr 3.3.3.2;
  peer 3.3.3.1 {

```

```
    session-establishment-hold-time 50;
    backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
        minimum-receive-interval 60;
        transmit-interval {
            minimum-interval 60;
        }
    }
}
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface ae2.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.1 {
        interface ae0;
    }
}
vlands {
    v100 {
        vlan-id 100;
        l3-interface vlan.100;
    }
    v200 {
        vlan-id 200;
        l3-interface vlan.200;
    }
    v500 {
        vlan-id 500;
        l3-interface vlan.500;
    }
}
```

From configuration mode on Switch B, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

chassis {
  aggregated-devices {
    ethernet {
      device-count 3;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/46 {
    ether-options {
      802.3ad ae1;
    }
  }
  xe-0/0/47 {
    ether-options {
      802.3ad ae2;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 1;
      mode active-active;
      status-control standby;
    }
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {

```

```
        members v100;
      }
    }
  }
}
ae2 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:06;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 4;
      chassis-id 1;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v200;
      }
    }
  }
}
vlan {
  unit 100 {
    family inet {
      address 10.1.1.10/24 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 150;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.20/24 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 150;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.1/24;
    }
  }
}
```

```

    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface vlan.100 {
        bfd-liveness-detection {
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 100;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  priority 500;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}

```

```
    }  
  }  
}  
iccp {  
  local-ip-addr 3.3.3.1;  
  peer 3.3.3.2 {  
    session-establishment-hold-time 50;  
    backup-liveness-detection {  
      backup-peer-ip 10.207.64.233;  
    }  
    liveness-detection {  
      minimum-receive-interval 60;  
      transmit-interval {  
        minimum-interval 60;  
      }  
    }  
  }  
}  
}  
igmp-snooping {  
  vlan all;  
}  
rstp {  
  interface ae0.0 {  
    disable;  
  }  
  interface ae1.0 {  
    edge;  
  }  
  interface ae2.0 {  
    edge;  
  }  
  interface all {  
    mode point-to-point;  
  }  
  bpdu-block-on-edge;  
}  
}  
multi-chassis {  
  multi-chassis-protection 3.3.3.2 {  
    interface ae0;  
  }  
}  
vllans {  
  v100 {  
    vlan-id 100;  
    l3-interface vlan.100;  
  }  
  v200 {  
    vlan-id 200;  
    l3-interface vlan.200;  
  }  
  v500 {  
    vlan-id 500;  
    l3-interface vlan.500;  
  }  
}
```

Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That Switch A is the Master Designated Router on page 2499](#)
- [Verifying That Switch B is the Backup Designated Router on page 2499](#)

Verifying That Switch A is the Master Designated Router

Purpose Verify that Switch A is the master designated router (DR).

Action From operational mode, enter the [show pim interfaces](#) command.

```
user@switch> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
```

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR	address
p1me.32769	Down	S	4	2	P2P,NotCap	0	0/0		
vlan.100	Up	S	4	2	DDR-DR,NotCap	1	0/0		10.1.1.11
vlan.200	Up	S	4	2	DDR-DR,NotCap	2	0/0		10.1.1.21

Meaning The DDR-DR state of the VLAN interfaces in the output shows that Switch A is the master designated router.

Verifying That Switch B is the Backup Designated Router

Purpose Verify that Switch B is the backup designated router (BDR).

Action From operational mode, enter the [show pim interfaces](#) command.

```
user@switch> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
```

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR	address
p1me.32769	Down	S	4	2	P2P,NotCap	0	0/0		
vlan.100	Up	S	4	2	DDR-BDR,NotCap	1	0/0		10.1.1.11
vlan.200	Up	S	4	2	DDR-BDR,NotCap	2	0/0		10.1.1.21

Meaning The DDR-BDR state of the VLAN interfaces in the output shows that Switch B is the backup designated router.

Related Documentation

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)

Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization

There are 2 methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to synchronize the MAC addresses between the switches for the participating MC-LAG interfaces, or you can configure Virtual Router Redundancy Protocol (VRRP). The procedure to configure MAC address synchronization is included in this example. For more information on configuring VRRP for use in a Layer 3 unicast MC-LAG, see [“Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\)”](#) on page 2514.

- [Requirements on page 2500](#)
- [Overview on page 2500](#)
- [Configuration on page 2501](#)
- [Verification on page 2511](#)
- [Troubleshooting on page 2514](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two switches

Before you configure an MC-LAG for Layer 3 unicast, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch”](#) on page 2454.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch”](#) on page 2458.


Overview

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. [Figure 65 on page 2463](#) shows the topology of this example.

Figure 67: Configuring a Multichassis LAG Between Switch A and Switch B



g041294

Table 235 on page 2455 details the topology used in this configuration example.

Table 238: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A	QFX3500 switch or QFX3600 switch	ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of ae0 : xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. ae1 is configured as an MC-LAG, and the following two interfaces are part of ae1 : xe-0/0/44 on Switch A and xe-0/0/46 on Switch B.
Switch B	QFX3500 switch or QFX3600 switch	

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500

```

```

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0

```

Configuring MC-LAG on Two Switches

Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2
```
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```
3. Configure a trunk interface between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```
4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Switch B:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```
2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



NOTE: Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
```

Switch B:

```
[edit protocols]
```

- ```

user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232

```
6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.
 

```

[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface vlan.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members v500

```

### Step-by-Step Procedure

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.



**NOTE:** At least one end needs to be active. The other end can be either active or passive.

- ```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active

```
2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

- ```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3

```
3. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0

```

Switch B:

- ```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1

```
4. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



NOTE: Only active-active mode is supported at this time.

- ```

[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active

```
5. Specify the status control for MC-LAG on Switch A and Switch B.



**NOTE:** You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

- Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



**NOTE:** The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

- Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

- Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

- Enable a VLAN on the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
[edit]
user@switch# set vlans v100 vlan-id 100
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

### Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

- Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

- Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.  

```
[edit]
user@switch# set protocols rstp interface ae1.0 edge
```
4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.  

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

### Results

Display the results of the configuration on Switch A.

```
chassis {
 aggregated-devices {
 ethernet {
 device-count 2;
 }
 }
}
interfaces {
 xe-0/0/12 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/13 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/44 {
 ether-options {
 802.3ad ae1;
 }
 }
}
ae0 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v500;
 }
 }
 }
}
ae1 {
 aggregated-ether-options {
 lacp {
 active;
 system-id 00:01:02:03:04:05;
 admin-key 3;
 }
 }
 mc-ae {
 mc-ae-id 3;
 chassis-id 0;
 }
}
```

```
 mode active-active;
 status-control active;
 init-delay-time 240
 }
}
unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v100;
 }
 }
}
vlan {
 unit 500 {
 family inet {
 address 3.3.3.2/24;
 }
 }
}
}
protocols {
 iccp {
 local-ip-addr 3.3.3.2;
 peer 3.3.3.1 {
 session-establishment-hold-time 50;
 backup-liveness-detection {
 backup-peer-ip 10.207.64.233;
 }
 liveness-detection {
 minimum-receive-interval 60;
 transmit-interval {
 minimum-interval 60;
 }
 }
 }
 }
}
rstp {
 interface ae0.0 {
 disable;
 }
 interface ae1.0 {
 edge;
 }
 interface all {
 mode point-to-point;
 }
 bpdu-block-on-edge;
}
}
multi-chassis {
 multi-chassis-protection 3.3.3.1 {
 interface ae0;
 }
}
}
```



```

vlangs {
 v100 {
 vlan-id 100;
 }
 v500 {
 vlan-id 500;
 l3-interface vlan.500;
 }
}

```

Display the results of the configuration on Switch B.

```

chassis {
 aggregated-devices {
 ethernet {
 device-count 2;
 }
 }
}
interfaces {
 xe-0/0/12 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/13 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/46 {
 ether-options {
 802.3ad ae1;
 }
 }
 ae0 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v500;
 }
 }
 }
 }
 ae1 {
 aggregated-ether-options {
 lacp {
 active;
 system-id 00:01:02:03:04:05;
 admin-key 3;
 }
 mc-ae {
 mc-ae-id 3;
 chassis-id 1;
 mode active-active;
 }
 }
 }
}

```

```
 status-control standby;
 init-delay-time 240
 }
}
unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v100;
 }
 }
}
vlan {
 unit 500 {
 family inet {
 address 3.3.3.1/24;
 }
 }
}
}
protocols {
 iccp {
 local-ip-addr 3.3.3.1;
 peer 3.3.3.2 {
 session-establishment-hold-time 50;
 backup-liveness-detection {
 backup-peer-ip 10.207.64.233;
 }
 liveness-detection {
 minimum-receive-interval 60;
 transmit-interval {
 minimum-interval 60;
 }
 }
 }
 }
}
rstp {
 interface ae0.0 {
 disable;
 }
 interface ae1.0 {
 edge;
 }
 interface all {
 mode point-to-point;
 }
 bpdu-block-on-edge;
}
}
multi-chassis {
 multi-chassis-protection 3.3.3.2 {
 interface ae0;
 }
}
}
vlangs {
```

```

v100 {
 vlan-id 100;
}
v500 {
 vlan-id 500;
 l3-interface vlan.500;
}
}

```

### Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2511](#)
- [Verifying That ICCP Is Working on Switch B on page 2511](#)
- [Verifying That LACP Is Active on Switch A on page 2512](#)
- [Verifying That LACP Is Active on Switch B on page 2512](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2512](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2513](#)
- [Verifying that MAC Learning Is Occurring on Switch A and Switch B on page 2513](#)

#### *Verifying That ICCP Is Working on Switch A*

**Purpose** Verify that ICCP is running on Switch A.

**Action** [edit]  
 user@switch# **show iccp**  
 Redundancy Group Information for peer 3.3.3.1  
   TCP Connection : Established  
   Liveliness Detection : Up  
  
 Client Application: MCSNOOPD  
  
 Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

#### *Verifying That ICCP Is Working on Switch B*

**Purpose** Verify that ICCP is running on Switch B.

**Action** **show iccp**  
  
 [edit]  
 user@switch# **show iccp**  
 Redundancy Group Information for peer 3.3.3.2  
   TCP Connection : Established  
   Liveliness Detection : Up  
  
 Client Application: MCSNOOPD

Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

#### *Verifying That LACP Is Active on Switch A*

**Purpose** Verify that LACP is active on Switch A.

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

| LACP state: | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| xe-0/0/46   | Actor   | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |
| xe-0/0/46   | Partner | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |

LACP protocol:                      Receive State      Transmit State                      Mux State  
xe-0/0/46                              Current      Fast periodic Collecting distributing

**Meaning** This output shows that Switch A is participating in LACP negotiation.

#### *Verifying That LACP Is Active on Switch B*

**Purpose** Verify that LACP is active on Switch B

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

| LACP state: | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| xe-0/0/44   | Actor   | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |
| xe-0/0/44   | Partner | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |

LACP protocol:                      Receive State      Transmit State                      Mux State  
xe-0/0/44                              Current      Fast periodic Collecting distributing

**Meaning** This output shows that Switch B is participating in LACP negotiation.

#### *Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A*

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

**Action** [edit]  
 user@switch# **show interfaces mc-ae**  
 Member Link : ae1  
 Current State Machine's State: mcae active state  
 Local Status : active  
 Local State : up  
 Peer Status : active  
 Peer State : up  
 Logical Interface : ae1.0  
 Topology Type : bridge  
 Local State : up  
 Peer State : up  
 Peer Ip/MCP/State : 3.3.3.1 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch A is up and active.

#### *Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B*

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

**Action** [edit]  
 user@switch# **show interfaces mc-ae**  
 Member Link : ae1  
 Current State Machine's State: mcae active state  
 Local Status : active  
 Local State : up  
 Peer Status : active  
 Peer State : up  
 Logical Interface : ae1.0  
 Topology Type : bridge  
 Local State : up  
 Peer State : up  
 Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch B is up and active.

#### *Verifying that MAC Learning Is Occurring on Switch A and Switch B*

**Purpose** Verify that MAC learning is working on Switch A and B.

**Action** [edit]  
 user@switch# **show ethernet-switching table**  
 Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries

| VLAN | MAC address       | Type     | Age | Interfaces   |
|------|-------------------|----------|-----|--------------|
| v222 | *                 | Flood    | -   | All-members  |
| v222 | 00:00:5e:00:01:01 | Static   | -   | Router       |
| v222 | 00:10:94:00:00:05 | Learn(L) | 33  | ae0.0 (MCAE) |
| v222 | 84:18:88:df:ac:ae | Learn(R) | 0   | ae2.0        |

**Meaning** The output shows four learned MAC addresses entries.

## Troubleshooting

---

### *Troubleshooting a LAG That Is Down*

**Problem** The `show interfaces terse` command shows that the MC-LAG is **down**

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2477](#)

### Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP)

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example. For more information on configuring MAC address synchronization, see *Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization*.

- [Requirements on page 2514](#)
- [Overview on page 2515](#)
- [Configuration on page 2516](#)
- [Verification on page 2528](#)
- [Troubleshooting on page 2532](#)

## Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two QFX3500 or QFX3600 switches

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch”](#) on page 2454.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch”](#) on page 2458.
- Configure Virtual Router Redundancy Protocol (VRRP) on a switch. See [“Configuring Basic VRRP Support”](#) on page 2296.

## Overview

In this example, you configure an MC-LAG across two switches by including interfaces from both switches in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.



**NOTE:** Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group
- Configure Layer 3 connectivity between the VRRP groups

## Topology

The topology used in this example consists of two switches hosting an MC-LAGs. The two switches are connected to a server. [Figure 68 on page 2516](#) shows the topology of this example.

Figure 68: Configuring a Multichassis LAG Between Switch A and Switch B




Table 235 on page 2455 details the topology used in this configuration example.

Table 239: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

| Hostname | Base Hardware                    | Multichassis Link Aggregation Group                                                                                                                                                                                                                                                                                                                                                                                            |
|----------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch A | QFX3500 switch or QFX3600 switch | <p><b>ae0</b> is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of <b>ae0</b>: <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch A and <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch B.</p> <p><b>ae1</b> is configured as an MC-LAG, and the following two interfaces are part of <b>ae1</b>: <b>xe-0/0/44</b> on Switch A and <b>xe-0/0/46</b> on Switch B.</p> |
| Switch B |                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                |

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500

```



```

set interfaces ae1 aggregated-ether-options lcp active
set interfaces ae1 aggregated-ether-options lcp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lcp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lcp active
set interfaces ae1 aggregated-ether-options lcp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lcp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50

```

```
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

### *Configuring MC-LAG on Two Switches*

#### **Step-by-Step Procedure**

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.  
  
[edit chassis]  
user@switch# set aggregated-devices ethernet device-count 2
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```

3. Configure a trunk interface between Switch A and Switch B.  
  
[edit interfaces]  
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Switch B:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

**Step-by-Step  
Procedure**

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 60
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



**NOTE:** Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
```

- user@switch# **set iccp peer 3.3.3.2 session-establishment-hold-time 50**
- (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



**NOTE:** By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
Switch B:
```

- ```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```
- Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface vlan.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members v500
```

Step-by-Step Procedure

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

- Enable LACP on the MC-LAG interface on Switch A and Switch B.



NOTE: At least one end needs to be active. The other end can be either active or passive.

- ```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active
```
- Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

- ```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```
- Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

4. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



NOTE: Only active-active mode is supported at this time.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

5. Specify the status control for MC-LAG on Switch A and Switch B.



NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

6. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



NOTE: The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

7. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

8. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

9. Enable a VLAN on the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
```

```
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

10. Enable VRRP on the MC-LAG on Switch A and Switch B:

- Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each switch in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1
virtual-address 100.1.1.1
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1
virtual-address 100.1.1.1
```

- Assign the priority for each switch in the VRRP group:



NOTE: The switch configured with the highest priority is the master.

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
```

- Enable the switch to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
```

- Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set vlans v100 l3-interface vlan.100
```

**Step-by-Step
Procedure**

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```
2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```
3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

- ```
[edit]
user@switch# set protocols rstp interface ae1.0 edge
```
4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



**NOTE:** The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

**Results**

Display the results of the configuration on Switch A.

```
chassis {
 aggregated-devices {
 ethernet {
 device-count 2;
 }
 }
}
interfaces {
 xe-0/0/12 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/13 {
 ether-options {
 802.3ad ae0;
 }
 }
}
```

```
}
xe-0/0/44 {
 ether-options {
 802.3ad ae1;
 }
}
ae0 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v500;
 }
 }
 }
}
ae1 {
 aggregated-ether-options {
 lacp {
 active;
 system-id 00:01:02:03:04:05;
 admin-key 3;
 }
 mc-ae {
 mc-ae-id 3;
 chassis-id 0;
 mode active-active;
 status-control active;
 init-delay-time 240;
 }
 }
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v100;
 }
 }
 }
}
vlan {
 unit 100 {
 family inet {
 address 100.1.1.1/24 {
 vrrp-group 1 {
 virtual-address 100.1.1.1;
 priority 200;
 accept-data;
 }
 }
 }
 }
 unit 500 {
 family inet {
 address 3.3.3.2/24;
 }
 }
}
```



```

 }
 }
}
protocols {
 iccp {
 local-ip-addr 3.3.3.2;
 peer 3.3.3.1 {
 session-establishment-hold-time 50;
 backup-liveness-detection {
 backup-peer-ip 10.207.64.233;
 }
 liveness-detection {
 minimum-receive-interval 60;
 transmit-interval {
 minimum-interval 60;
 }
 }
 }
 }
}
rstp {
 interface ae0.0 {
 disable;
 }
 interface ae1.0 {
 edge;
 }
 interface all {
 mode point-to-point;
 }
 bpdu-block-on-edge;
}
}
multi-chassis {
 multi-chassis-protection 3.3.3.1 {
 interface ae0;
 }
}
vllans {
 v100 {
 vlan-id 100;
 l3-interface vlan.100;
 }
 v500 {
 vlan-id 500;
 l3-interface vlan.500;
 }
}
}

```

Display the results of the configuration on Switch B.

```

chassis {
 aggregated-devices {
 ethernet {
 device-count 2;
 }
 }
}

```

```
}
interfaces {
 xe-0/0/12 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/13 {
 ether-options {
 802.3ad ae0;
 }
 }
 xe-0/0/44 {
 ether-options {
 802.3ad ae1;
 }
 }
}
ae0 {
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v500;
 }
 }
 }
}
ae1 {
 aggregated-ether-options {
 lacp {
 active;
 system-id 00:01:02:03:04:05;
 admin-key 3;
 }
 mc-ae {
 mc-ae-id 3;
 chassis-id 1;
 mode active-active;
 status-control active;
 init-delay-time 240;
 }
 }
 unit 0 {
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members v100;
 }
 }
 }
}
vlan {
 unit 100 {
 family inet {
 address 100.1.1.10/24 {
 vrrp-group 1 {
```

```

 virtual-address 100.1.1.1;
 priority 200;
 accept-data;
 }
}
}
unit 500 {
 family inet {
 address 3.3.3.1/24;
 }
}
}
}
protocols {
 iccp {
 local-ip-addr 3.3.3.1;
 peer 3.3.3.2 {
 session-establishment-hold-time 50;
 backup-liveness-detection {
 backup-peer-ip 10.207.64.233;
 }
 liveness-detection {
 minimum-receive-interval 60;
 transmit-interval {
 minimum-interval 60;
 }
 }
 }
 }
}
}
rstp {
 interface ae0.0 {
 disable;
 }
 interface ae1.0 {
 edge;
 }
 interface all {
 mode point-to-point;
 }
 bpdu-block-on-edge;
}
}
multi-chassis {
 multi-chassis-protection 3.3.3.2 {
 interface ae0;
 }
}
}
vllans {
 v100 {
 vlan-id 100;
 l3-interface vllan.100;
 }
 v500 {
 vlan-id 500;
 l3-interface vllan.500;
 }
}
}

```

```
}
}
```

## Verification

---

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2528](#)
- [Verifying That ICCP Is Working on Switch B on page 2528](#)
- [Verifying That LACP Is Active on Switch A on page 2529](#)
- [Verifying That LACP Is Active on Switch B on page 2529](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2529](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2530](#)
- [Verifying that MAC Learning Is Occurring on Switch A on page 2530](#)
- [Verifying that MAC Learning Is Occurring on Switch B on page 2531](#)
- [Verifying that Switch A is the Master in the VRRP Group on page 2531](#)
- [Verifying that Switch B is the Backup Member in the VRRP Group on page 2531](#)
- [Verifying that the Virtual IP Address is Attached to an Individual Address on Switch A on page 2532](#)
- [Verifying that the Virtual IP Address is Attached to an Individual Address on Switch B on page 2532](#)

### *Verifying That ICCP Is Working on Switch A*

**Purpose** Verify that ICCP is running on Switch A.

**Action** [edit]  
user@switch# **show iccp**  
Redundancy Group Information for peer 3.3.3.1  
TCP Connection : Established  
Liveliness Detection : Up  
  
Client Application: MCSNOOPD  
  
Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

### *Verifying That ICCP Is Working on Switch B*

**Purpose** Verify that ICCP is running on Switch B.

**Action** **show iccp**  
  
[edit]  
user@switch# **show iccp**  
Redundancy Group Information for peer 3.3.3.2

```
TCP Connection : Established
Liveliness Detection : Up
```

```
Client Application: MCSNOOPD
```

```
Client Application: eswd
```

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveliness detection is up, and MCSNOOPD and ESWD client applications are running.

### *Verifying That LACP Is Active on Switch A*

**Purpose** Verify that LACP is active on Switch A.

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

| LACP state: | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| xe-0/0/46   | Actor   | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |
| xe-0/0/46   | Partner | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |

LACP protocol:

|           | Receive State | Transmit State | Mux State               |
|-----------|---------------|----------------|-------------------------|
| xe-0/0/46 | Current       | Fast periodic  | Collecting distributing |

**Meaning** This output shows that Switch A is participating in LACP negotiation.

### *Verifying That LACP Is Active on Switch B*

**Purpose** Verify that LACP is active on Switch B

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

| LACP state: | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| xe-0/0/44   | Actor   | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |
| xe-0/0/44   | Partner | No  | No  | Yes  | Yes | Yes | Yes  | Fast    | Active   |

LACP protocol:

|           | Receive State | Transmit State | Mux State               |
|-----------|---------------|----------------|-------------------------|
| xe-0/0/44 | Current       | Fast periodic  | Collecting distributing |

**Meaning** This output shows that Switch B is participating in LACP negotiation.

### *Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A*

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.1 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch A is up and active.

***Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B***

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch B is up and active.

***Verifying that MAC Learning Is Occurring on Switch A***

**Purpose** Verify that MAC learning is working on Switch A.

**Action** [edit]  
user@switch# **show ethernet-switching table**  
Ethernet-switching table: 6 entries, 1 learned, 0 persistent entriesC

| VLAN | MAC address       | Type     | Age | Interfaces    |
|------|-------------------|----------|-----|---------------|
| v100 | *                 | Flood    |     | - All-members |
| v100 | 00:00:5e:00:01:01 | Static   |     | - Router      |
| v100 | 78:fe:3d:5a:07:42 | Static   |     | - Router      |
| v100 | 78:fe:3d:5b:ad:c2 | Learn(R) | 0   | ae0.0         |
| v500 | *                 | Flood    |     | - All-members |
| v500 | 78:fe:3d:5a:07:42 | Static   |     | - Router      |

**Meaning** The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the

VRRP master member learned the VLAN v100 Learn (R) MAC address of the VRRP backup member.

### *Verifying that MAC Learning Is Occurring on Switch B*

**Purpose** Verify that MAC learning is working on Switch B.

**Action** [edit]  
 user@switch# **show ethernet-switching table**  
 Ethernet-switching table: 7 entries, 1 learned, 0 persistent entries

| VLAN | MAC address       | Type     | Age | Interfaces    |
|------|-------------------|----------|-----|---------------|
| v100 | *                 | Flood    |     | - All-members |
| v100 | 00:00:5e:00:01:01 | Static   |     | - Router      |
| v100 | 78:fe:3d:5a:07:42 | Learn(R) | 0   | ae0.0         |
| v100 | 78:fe:3d:5b:ad:c2 | Static   |     | - Router      |
| v200 | 78:fe:3d:5b:ad:c2 | Static   |     | - Router      |
| v500 | *                 | Flood    |     | - All-members |
| v500 | 78:fe:3d:5b:ad:c2 | Static   |     | - Router      |

**Meaning** The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the VRRP backup member learned the VLAN v100 Learn (R) MAC address of the VRRP master member.

### *Verifying that Switch A is the Master in the VRRP Group*

**Purpose** Verify that Switch A is the master member in the VRRP group.

**Action** [edit]  
 user@switch# **show vrrp**

| Interface | State | Group | VR state | VR Mode | Timer   | Type | Address    |
|-----------|-------|-------|----------|---------|---------|------|------------|
| vlan.100  | up    | 1     | master   | Active  | A 0.605 | lcl  | 100.1.1.11 |
|           |       |       |          |         |         | vip  | 100.1.1.1  |

**Meaning** The output shows that Switch A is the master member in the VRRP group.

### *Verifying that Switch B is the Backup Member in the VRRP Group*

**Purpose** Verify that Switch B is the backup member in the VRRP group.

**Action** [edit]  
 user@switch# **show vrrp**

| Interface | State | Group | VR state | VR Mode | Timer   | Type | Address    |
|-----------|-------|-------|----------|---------|---------|------|------------|
| vlan.100  | up    | 1     | backup   | Active  | A 0.605 | lcl  | 100.1.1.10 |
|           |       |       |          |         |         | vip  | 100.1.1.1  |

**Meaning** The output shows that Switch B is the backup member in the VRRP group.

***Verifying that the Virtual IP Address is Attached to an Individual Address on Switch A***

**Action** [edit]  
user@switch# run show interfaces terse vlan

| Interface | Admin | Link | Proto | Local                         | Remote |
|-----------|-------|------|-------|-------------------------------|--------|
| vlan      | up    | up   |       |                               |        |
| vlan.100  | up    | up   | inet  | 100.1.1.1/24<br>100.1.1.11/24 |        |
| vlan.500  | up    | up   | inet  | 3.3.3.2/24                    |        |

**Meaning** The output shows that the virtual IP address (100.1.1.1/24) is bound to the individual IP address (100.1.1.11/24) on Switch A.

***Verifying that the Virtual IP Address is Attached to an Individual Address on Switch B***

**Action** [edit]  
user@switch# run show interfaces terse vlan

| Interface | Admin | Link | Proto | Local                         | Remote |
|-----------|-------|------|-------|-------------------------------|--------|
| vlan      | up    | up   |       |                               |        |
| vlan.100  | up    | up   | inet  | 100.1.1.1/24<br>100.1.1.10/24 |        |
| vlan.500  | up    | up   | inet  | 3.3.3.1/24                    |        |

**Meaning** The output shows that the virtual IP address (100.1.1.1/24) is bound to the individual IP address (100.1.1.10/24) on Switch B.

---

**Troubleshooting*****Troubleshooting a LAG That Is Down***

**Problem** The `show interfaces terse` command shows that the MC-LAG is **down**

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)

---

**Configuration Tasks**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533](#)
- [Configuring Aggregated Ethernet LACP on page 2534](#)



- [Configuring Ethernet Loopback Capability on page 2535](#)
- [Configuring Interfaces for Uplink Failure Detection on page 2536](#)
- [Configuring a Layer 3 Logical Interface on page 2537](#)
- [Configuring Link Aggregation on page 2537](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)
- [Configuring the QSFP+ Port Type on QFX3500 Standalone Switches on page 2543](#)
- [Configuring the Port Type on QFX3600 Standalone Switches on page 2545](#)

## Configuring Gigabit and 10-Gigabit Ethernet Interfaces

QFX Series products include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces
- Provides basic Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP) configuration

This topic describes:

- [Configuring Port Mode on page 2533](#)
- [Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces on page 2533](#)
- [Configuring the IP Options on page 2534](#)

### Configuring Port Mode

If you are connecting a switch to other switches and to routers on the LAN, you need to assign the interface to a logical port and you need to configure the logical port as a trunk port.

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode:

[edit]

```
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode trunk
```

### Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces

QFX Series products include a factory default configuration that enables Gigabit Ethernet and 10-Gigabit Ethernet and interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- The speed for Gigabit Ethernet interfaces is set to 1000 Mbps by default.

- Gigabit Ethernet interfaces operate in full-duplex mode by default, and autonegotiation is supported. If you want to disable autonegotiation, you need to manually set the speed to 1g.

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default. The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default. Autonegotiation is not supported.

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- **autonegotiation**—Enable or disable autonegotiation of flow control, link mode, and speed for Gigabit Ethernet interfaces.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic** for Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

To set **ether-options** for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces:

```
[edit]
user@switch# set interfaces interface-name ether-options
```

---

### Configuring the IP Options

To specify an IP address for the logical unit:

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

#### Related Documentation

- [Monitoring Interface Status and Traffic on page 317](#)
- [show interfaces xe on page 2217](#)
- [show interfaces ge-](#)
- [Understanding Interface Naming Conventions on page 2419](#)

## Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces on the QFX Series, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link.

You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

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

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

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

For example, to specify the interval as fast, execute the following command:

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

#### Related Documentation

- [Configuring Link Aggregation on page 2537](#)
- [Verifying the Status of a LAG Interface on page 2630](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
- [show lacp statistics interfaces \(View\) on page 2806](#)

## Configuring Ethernet Loopback Capability

To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@switch# delete interfaces interface-name ether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

**no-loopback;**

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**

**Related  
Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533](#)

## Configuring Interfaces for Uplink Failure Detection

You can configure uplink failure detection on the QFX Series to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure information to downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Follow these configuration guidelines:

- Configure an interface in only one group.
- Configure a maximum of eight groups for each switch.
- Configure a maximum of eight uplinks to monitor and a maximum of 48 downlinks to disable in each group.
- Configure physical links and logical links in separate groups.

To configure uplink failure detection on a switch:

1. Specify a name for an uplink failure detection group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for each uplink interface you add to the group.

4. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for each downlink interface you add to the group.



**NOTE:** After you have configured an uplink failure detection group, use the **show uplink-failure-detection group (*Uplink Failure Detection*) *group-name*** command to verify that all interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.

---

**Related  
Documentation**

- [Overview of Uplink Failure Detection on page 2445](#)

- [Example: Configuring Interfaces for Uplink Failure Detection on page 2449](#)
- [Verifying That Uplink Failure Detection Is Working Correctly on page 2632](#)

## Configuring a Layer 3 Logical Interface

QFX Series systems use Layer 3 logical interfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. Layer 3 logical interfaces route traffic between subnets.

To configure Layer 3 logical interfaces, enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See [“Configuring VLANs” on page 2048](#).

To configure Layer 3 logical interfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id vlan-id-number
```

### Related Documentation

- [Understanding Layer 3 Logical Interfaces on page 2434](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 2630](#)

## Configuring Link Aggregation

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



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.

1. [Creating an Aggregated Ethernet Interface on page 2538](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 2538](#)
3. [Configuring Aggregated Ethernet LACP on page 2538](#)

## Creating an Aggregated Ethernet Interface

---

To create an aggregated Ethernet interface:

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

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count device-count
```

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count
```

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



**NOTE:** By default only one link must be up for the bundle to be labeled “up”.

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links minimum-links
```

For example, to specify 5:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links 5
```

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

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed link-speed
```

For example, to specify 10g:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed 10g
```

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

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

## Configuring the VLAN Name and VLAN ID Number

---

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

For example, 100.

## Configuring Aggregated Ethernet LACP

---

For aggregated Ethernet interfaces on the QFX Series, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

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



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

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

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

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

For example, to specify the interval as fast, execute the following command:

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

#### Related Documentation

- [Understanding Interface Naming Conventions on page 2419](#)
- [Verifying the Status of a LAG Interface on page 2630](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)

## Configuring Multichassis Link Aggregation

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

The MC-LAG switches use the Interchassis Control Protocol (ICCP) to exchange the control information between two MC-LAG switches.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of MC-LAG are two MC-LAG switches. Each of the switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregated Ethernet interface or multichassis aggregated Ethernet interface group.

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae mc-ae-id number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae chassis-id number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.



**NOTE:** Only active-active mode is supported at this time.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```



4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.



**NOTE:** You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

5. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

6. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

7. Configure ICCP by doing the following on each switch hosting the MC-LAG:

- a. Configure the local IP address to be used by all switches hosting the MC-LAG.

```
[edit protocols]
user@switch# set iccp local-ip-addr local-ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

- b. (Optional) Configure the IP address of the switch and the time during which an ICCP connection must succeed between the switches hosting the MC-LAG.

Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address session-establishment-hold-time seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

- c. (Optional) Configure the IP address to be used for backup liveness detection:



**NOTE:** By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip
ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip
10.207.64.232
```

- d. Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.



**NOTE:** Configuring the minimum receive interval is required to enable BFD.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
```

- e. Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
60
```

8. Configure a multichassis protection link between the switches.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit protocols]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

9. Enable RSTP globally on all interfaces.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

10. Disable RSTP on the ICL-PL interfaces on both switches.

```
[edit]
user@switch# set protocols rstp interface interface-name disable
```

For example:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

11. Configure the MC-LAG interfaces as edge ports on both switches.

```
set protocols rstp interface interface-name
```

For example:

```
[edit]
user@switch# set protocols rstp interface ae1.0
```

12. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both switches.

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

For example:

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

#### Related Documentation

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2514](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2477](#)

## Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

By default, the four 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet (xe) ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. You can, however, configure the four 40-Gbps QSFP+ ports to operate as 40-Gigabit Ethernet (xle) ports.



**NOTE:** Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports, because one port is reserved.



**CAUTION:** The Packet Forwarding Engine on the QFX3500 standalone switch is restarted when you commit port type configuration changes (for example, configuring or deleting an xle port). As a result, you might experience packet loss on the device.

The following steps describe how to configure either a block of ports or an individual port to operate as 40-Gigabit Ethernet (xle) ports, as well as how to delete a 40-Gigabit Ethernet (xle) configuration.



**NOTE:** When you delete an xle block of ports or individual port, the ports return to operating as 10-Gigabit Ethernet ports.

1. To configure a block of ports to operate as 40-Gigabit Ethernet (xle) ports, specify a port range:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port-range port-range-low port-range-high
```

For example, to configure ports Q0 through Q3 to operate as 40-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port-range 0 3
```

2. To configure an individual port to operate as a 40-Gigabit Ethernet (xle) port, specify a port number:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port port-number
```

For example, to configure port Q2 to operate as a 40-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port 2
```

3. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

4. To delete a block of ports configured as 40-Gigabit Ethernet (xle) ports (and return to the default 10-Gigabit Ethernet configuration), specify a port range:

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port-range port-range-low port-range-high
```

For example, to delete the 40-Gigabit Ethernet (xle) port configuration for ports Q0 through Q3 (and return to the default 10-Gigabit Ethernet configuration):

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port-range 0 3
```

5. To delete an individual port configured as a 40-Gigabit Ethernet (xle) port (and return to the default 10-Gigabit Ethernet configuration), specify an individual port:

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port port-number
```

For example, to delete the 40-Gigabit Ethernet (xle) port configuration for port Q2 (and return to the default 10-Gigabit Ethernet configuration):

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port 2
```

- Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

**Related  
Documentation**

- [Understanding Interface Naming Conventions on page 2419](#)
- [pic on page 2612](#)

## Configuring the Port Type on QFX3600 Standalone Switches

The QFX3600 standalone switch provides 16 40-Gbps QSFP+ ports. By default, all 16 ports operate as 40-Gigabit Ethernet (xle) ports. Optionally, you can choose to configure the 40-Gbps ports to operate as four 10-Gigabit Ethernet (xe) ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.



**NOTE:** Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports, because one port is reserved. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.

This topic explains how to configure the port type on QFX3600 standalone switches.



**CAUTION:** The Packet Forwarding Engine on the QFX3600 standalone switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

The following message may be displayed in the system log file when the Packet Forwarding Engine is restarted. You can ignore this message.

Pipe write error: Broken pipe

flush operation failed

The following steps describe how to configure either a block of ports or an individual port to operate as 10-Gigabit Ethernet (xe) ports, as well as how to delete a 10-Gigabit Ethernet (xe) port configuration.



**NOTE:** When you delete the xe port type configuration for an individual port or a block of ports, the ports return to operating as 40-Gigabit Ethernet (xle) ports.

- To configure a block of ports to operate as 10-Gigabit Ethernet (xe) ports, specify a port range:

```
[edit chassis fpc 0 pic 0]
```

```
user@switch# set xe port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port-range 4 7
```

2. To configure an individual port to operate as a 10-Gigabit Ethernet (xe) port, specify a port number:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port port-number
```

For example, to configure port Q4 to operate as a 10-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port 4
```

3. Review your configuration and issue the **commit** command.

```
[edit chassis fpc 0 pic 0]
user@switch# commit
commit complete
```

4. To delete the 10-Gigabit Ethernet (xe) port configuration for a block of ports (and return to the default 40-Gigabit Ethernet configuration), specify a port range:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port-range port-range-low port-range-high
```

For example, to delete the 10-Gigabit Ethernet port configuration for ports Q4 through Q7:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port-range 4 7
```

5. To delete the 10-Gigabit Ethernet (xe) port configuration for an individual port (and return to the default 40-Gigabit Ethernet configuration), specify a port number:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port port-number
```

For example, to delete the 10-Gigabit Ethernet port configuration for port Q4:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port 4
```

- Related Documentation**
- [Understanding Interface Naming Conventions on page 2419](#)
  - [pic on page 2612](#)

---

## Configuration Statements

- [802.3ad on page 2549](#)
- [address on page 2550](#)
- [aggregated-devices on page 2552](#)
- [aggregated-ether-options on page 2553](#)

- [alarm \(chassis\) on page 2554](#)
- [authentication-key \(ICCP\) on page 2555](#)
- [backup-liveness-detection on page 2555](#)
- [backup-peer-ip on page 2556](#)
- [chassis on page 2557](#)
- [chassis \(QFabric System\) on page 2558](#)
- [chassis-id on page 2559](#)
- [container-devices on page 2560](#)
- [craft-lockout on page 2561](#)
- [description \(Interfaces\) on page 2562](#)
- [detection-time \(Liveness Detection\) on page 2563](#)
- [device-count on page 2563](#)
- [disk-failure-action on page 2564](#)
- [ethernet on page 2564](#)
- [ethernet \(Alarm\) on page 2565](#)
- [ethernet-switching on page 2566](#)
- [ether-options on page 2567](#)
- [eui-64 on page 2568](#)
- [family on page 2569](#)
- [fibre-channel \(Alarm\) on page 2571](#)
- [flow-control on page 2572](#)
- [force-up on page 2573](#)
- [fpc on page 2574](#)
- [fpc \(Interconnect Device\) on page 2575](#)
- [gratuitous-arp-reply on page 2575](#)
- [group on page 2576](#)
- [hold-time \(Physical Interface\) on page 2577](#)
- [iccp on page 2578](#)
- [inet \(interfaces\) on page 2579](#)
- [inet6 \(interfaces\) on page 2580](#)
- [interconnect-device \(Chassis\) on page 2581](#)
- [interface-range on page 2582](#)
- [interface \(Multichassis Protection\) on page 2583](#)
- [interfaces on page 2584](#)
- [lacp \(802.3ad\) on page 2590](#)
- [lacp \(Aggregated Ethernet\) on page 2591](#)
- [link-to-disable on page 2591](#)

- [link-to-monitor](#) on page 2592
- [link-down](#) on page 2593
- [link-mode](#) on page 2594
- [link-speed](#) on page 2595
- [liveness-detection](#) on page 2596
- [local-ip-addr \(ICCP\)](#) on page 2596
- [loopback \(Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet\)](#) on page 2597
- [management-ethernet \(Alarm\)](#) on page 2597
- [mc-ae](#) on page 2598
- [mc-ae-id](#) on page 2599
- [member](#) on page 2599
- [member-range](#) on page 2600
- [minimum-interval \(Liveness Detection\)](#) on page 2600
- [minimum-receive-interval \(Liveness Detection\)](#) on page 2601
- [minimum-links](#) on page 2601
- [mode \(QFX Series\)](#) on page 2602
- [mtu](#) on page 2603
- [multi-chassis](#) on page 2604
- [multi-chassis-protection](#) on page 2604
- [multiplier \(Liveness Detection\)](#) on page 2605
- [no-adaptation \(Liveness Detection\)](#) on page 2605
- [no-gratuitous-arp-request](#) on page 2606
- [node-device \(Chassis\)](#) on page 2607
- [node-group \(Chassis\)](#) on page 2608
- [on-disk-failure](#) on page 2609
- [peer \(ICCP\)](#) on page 2610
- [peer \(Multichassis\)](#) on page 2611
- [periodic](#) on page 2611
- [pic](#) on page 2612
- [port-mode](#) on page 2613
- [reflective-relay](#) on page 2614
- [routing-engine](#) on page 2614
- [session-establishment-hold-time](#) on page 2615
- [speed](#) on page 2616
- [status-control](#) on page 2616
- [targeted-broadcast](#) on page 2617
- [threshold \(Detection Time\)](#) on page 2617



- [traceoptions \(Individual Interfaces\) on page 2618](#)
- [traceoptions \(ICCP\) on page 2619](#)
- [transmit-interval \(Liveness Detection\) on page 2620](#)
- [traps on page 2621](#)
- [unit on page 2622](#)
- [uplink-failure-detection on page 2623](#)
- [version \(Liveness Detection\) on page 2623](#)
- [vlan-id on page 2624](#)
- [vlan-tagging on page 2624](#)
- [xe \(Port\) on page 2625](#)
- [xle \(Port\) on page 2626](#)

## 802.3ad

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>802.3ad aex;   lacp {     force-up;     (primary   backup);   } }</pre>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify the aggregated Ethernet logical interface number.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>aex</b> —Aggregated Ethernet logical interface number.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2534</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul> |

## address

```

Syntax address address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 destination address;
 destination-profile name;
 eui-64;
 master-only;
 multipoint-destination address dlcid dlcid-identifier;
 multipoint-destination address {
 epd-threshold cells;
 inverse-arp;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 shaping {
 (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
 length);
 queue-length number;
 }
 vci vpi-identifier.vci-identifier;
 }
 primary;
 preferred;
 (vrrp-group | vrrp-inet6-group) group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-type authentication;
 authentication-key key;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority-number number;
 track {
 priority-cost seconds;
 priority-hold-time interface-name {
 interface priority;
 bandwidth-threshold bits-per-second {
 priority;
 }
 }
 }
 route ip-address/mask routing-instance instance-name priority-cost cost;
 }
 virtual-address [addresses];
 }
}

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the interface address.

**Options** *address*—Address of the interface.

The remaining statements are explained separately.



**NOTE:** The `edit logical-systems` hierarchy is not available on QFabric systems.

**Required Privilege** interface—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring the Protocol Family*
  - *negotiate-address*
  - *unnumbered-address (Ethernet)*
  - *Junos OS System Basics Configuration Guide*
  - *family*

## aggregated-devices

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aggregated-devices {<br/>    ethernet {<br/>        device-count <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <i>chassis</i> ],<br>[edit <i>chassis node-group name</i> ]                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure properties for aggregated devices on the switch.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | Aggregated devices are disabled.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li><li>• <a href="#">Configuring Link Aggregation on page 2537</a></li><li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454</a></li><li>• <i>Junos® OS Network Interfaces</i></li></ul> |

## aggregated-ether-options

```
Syntax aggregated-ether-options {
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 ethernet-switch-profile {
 tag-protocol-id;
 (flow-control | no-flow-control);
 lacp mode {
 admin-key key;
 periodic interval;
 system-id mac-address;
 }
 }
 (link-protection | no-link-protection);
 link-speed speed;
 (loopback | no-loopback);
 mc-ae {
 chassis-id chassis-id;
 mc-ae-id mc-ae-id;
 mode (active-active);
 status-control (active | standby);
 }
 minimum-links number;
 rebalance-periodic;
 source-address-filter filter;
 (source-filtering | no-source-filtering);
 }
```

**Hierarchy Level** [edit [interfaces aex](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure properties specific to a specific aggregated Ethernet interface.

The statements are explained separately.

**Default** Options are not enabled.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2417](#)
- [Configuring Aggregated Ethernet LACP on page 2534](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
- [Junos® OS Network Interfaces](#)

## alarm (chassis)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>alarm {<br/>    interface-type {<br/>        alarm-name (ignore   red   yellow);<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit chassis],<br>[edit chassis <b>interconnect-device</b> name],<br>[edit chassis <b>node-group</b> name]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for the ACX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description              | <p>Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the <b>RED ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the <b>YELLOW ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.</p> <p>To configure more than one alarm, include multiple <b>alarm-name</b> lines.</p> |
| Options                  | <p><b>alarm-name</b>—Alarm condition. For a list of conditions, see <i>System-Wide Alarms and Alarms for Each Interface Type</i>.</p> <p><b>ignore</b>—The specified alarm condition does not set off any alarm.</p> <p><b>interface-type</b>—Type of interface on which you are configuring the alarm: <b>atm</b>, <b>ethernet</b>, <b>sonet</b>, or <b>t3</b>.</p> <p><b>red</b>—The specified alarm condition sets off a red alarm.</p> <p><b>yellow</b>—The specified alarm condition sets off a yellow alarm.</p>                                                                                                                             |
| Required Privilege Level | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 6487</a></li><li>• <a href="#">Chassis Conditions That Trigger Alarms</a></li><li>• <a href="#">Chassis Alarm Messages on a QFX3500 Device on page 6488</a></li><li>• <a href="#">Interface Alarm Messages on page 6491</a></li></ul>                                                                                                                                                                                                                                                                                                                             |

## authentication-key (ICCP)

|                                 |                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key key;</code>                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit protocols <b>iccp</b> peer &lt;peer-IP-address&gt;],</code><br><code>[edit protocols <b>iccp</b>]</code>                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify the authentication key (password). The QFX3500 device uses this password to verify the authenticity of packets sent from the peers hosting a multichassis link aggregation group (MC-LAG). Peer-level authentication takes precedence over global-level authentication.<br><br>Interchassis Control Protocol (ICCP) uses MD5 authentication. |
| <b>Options</b>                  | <b>key</b> —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").                                                |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                  |

## backup-liveness-detection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>backup-liveness-detection {</code><br><code>    <b>backup-peer-ip</b> ip4-address</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit protocols <b>iccp</b> peer]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Backup liveness detection determines the peer status (whether it is up or down) by exchanging keep alive messages (UDP-based packets) over the management link between the two Interchassis Control Protocol (ICCP) peers. When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG. |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## backup-peer-ip

---

|                                 |                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | backup-peer-ip <i>ip4-address</i>                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer backup-liveness-detection</a> ]                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                             |
| <b>Description</b>              | Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.           |



## chassis

```
Syntax chassis {
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 aggregated-devices {
 ethernet {
 device-count number;
 }
 alarm {
 interface-type {
 alarm-name (red | yellow | ignore);
 }
 }
 }
 fpc slot {
 pic pic-number {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 xe {
 (port port-number | port-range port-range-low port-range-high);
 }
 xle {
 (port port-number | port-range port-range-low port-range-high);
 }
 }
 }
 }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation on page 2537](#)

## chassis (QFabric System)

---

```
Syntax chassis {
 interconnect-device {
 alarm {
 (ethernet | management-ethernet) {
 link-down (red | yellow | ignore);
 }
 }
 container-devices {
 device-count number;
 }
 craft-lockout {
 alarm {
 interface-type {
 link-down (red | yellow | ignore);
 }
 }
 }
 container-devices {
 device-count number;
 }
 fpc slot {
 power (on | off);
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }
 fpc slot {
 power (on | off);
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }
 node-group name {
 aggregated-devices {
 ethernet {
 device-count number;
 }
 }
 alarm {
 interface-type {
 link-down (ignore | red | yellow);
 }
 }
 container-devices {
 device-count number;
 }
 node-device name {
```

```

 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 pic pic-number {
 fte {
 port port-number;
 port-range port-range-low port-range-high;
 }
 xe {
 port port-number;
 port-range port-range-low port-range-high;
 }
 }
}
routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
}
}
}

```

**Hierarchy Level** [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure chassis-specific properties for the switch.  
  
The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

## chassis-id

**Syntax** `chassis-id chassis-id;`

**Hierarchy Level** [\[edit\]](#)[interfaces](#) [aggregated-ether-options](#) [mc-ae](#)

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

## container-devices

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>container-devices {<br/>    device-count <i>number</i>;<br/>}</pre>                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit chassis]<br/>[edit chassis <i>interconnect-device name</i>]<br/>[edit chassis <i>node-group name</i>]</pre>                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                                                                                             |
| <b>Description</b>              | Specify the container devices configuration. The <b>number</b> option specifies the number of sequentially numbered container interfaces, from <b>ci0</b> to <b>ci127</b> maximum. |
| <b>Options</b>                  | <b>number</b> —Number of container devices.<br><b>Range:</b> 1 through 128                                                                                                         |
| <b>Required Privilege Level</b> | <b>chassis</b> —To view this statement in the configuration.<br><b>chassis-control</b> —To add this statement to the configuration.                                                |

## craft-lockout

```

Syntax craft-lockout {
 alarm {
 interface-type {
 link-down (red | yellow | ignore);
 }
 }
 container-devices {
 device-count number;
 }
 fpc slot {
 pic pic-number {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 }
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }

```

**Hierarchy Level** [edit [chassis interconnect-device](#)]

**Release Information** Statement introduced in Junos Release 11.3 for the QFX Series.

**Description** Disable the physical operation of the craft interface front panel.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Junos OS to Disable the Physical Operation of the Craft Interface*

## description (Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ],<br>[edit <a href="#">interfaces</a> <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit <a href="#">logical-systems</a> <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                              |
| <b>Description</b>              | <p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>                   |
| <b>Options</b>                  | <b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Adding an Interface Description to the Configuration</i></li><li>• <i>Adding a Logical Unit Description to the Configuration</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <i>Enabling and Disabling Insertion of Option 82 Information</i></li><li>• <i>Junos® OS Network Interfaces</i></li></ul> |

## detection-time (Liveness Detection)

---

|                                 |                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>detection-time {<br/>    <i>milliseconds</i>;<br/>}</code>                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                           |
| <b>Description</b>              | <p>The Bidirectional Forwarding Detection (BFD) timers are adaptive and can be adjusted to be faster or slower.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                              |

## device-count

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>device-count <i>number</i>;</code>                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit <a href="#">chassis aggregated-devices ethernet</a>],</p> <p>[edit <a href="#">chassis node-group <i>name</i> aggregated-devices ethernet</a>]</p>                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure the number of aggregated Ethernet logical devices available to the switch.                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454</a></li> </ul> |

## disk-failure-action

---

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disk-failure-action (halt   reboot);                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit chassis routing-engine on-disk-failure],<br>[edit chassis <a href="#">node-group name</a> routing-engine on-disk-failure],<br>[edit chassis <a href="#">interconnect-device name</a> routing-engine on-disk-failure] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Halt or reboot when the Routing Engine hard disk fails.                                                                                                                                                                    |
| <b>Options</b>                  | <b>halt</b> —Stop the Routing Engine.<br><br><b>reboot</b> —Reboot the Routing Engine.                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors</a></li></ul>                                                                        |

## ethernet

---

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ethernet {<br><a href="#">device-count</a> <i>number</i> ;<br>}                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices</a> ],<br>[edit chassis <a href="#">node-group aggregated-devices</a> ]                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                  |
| <b>Description</b>              | Configure properties for aggregated Ethernet devices on the switch.<br><br>The remaining statement is explained separately.                                        |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 2537</a></li><li>• <a href="#">Junos® OS Network Interfaces</a></li></ul> |



---

## ethernet (Alarm)

---

|                                 |                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ethernet {<br><a href="#">link-down</a> (red   yellow   ignore);<br>}                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">alarm</a> ],<br>[edit chassis <a href="#">interconnect-device</a> <i>name</i> <a href="#">alarm</a> ],<br>[edit chassis <a href="#">node-group</a> <i>name</i> <a href="#">alarm</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                |
| <b>Description</b>              | Configure alarms for an Ethernet interface.                                                                                                                                                                      |
| <b>Options</b>                  | The remaining statement is explained separately.—                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 6487</a></li><li>• <a href="#">Interface Alarm Messages on page 6491</a></li></ul>                                              |

## ethernet-switching

---

**Syntax** ethernet-switching {  
    `filter` input *filter-name*;  
    `filter` output *filter-name*;  
    `native-vlan-id` *vlan-id*;  
    `port-mode` *mode*;  
    `reflective-relay`;  
    vlan {  
        `members` [ (all | *names* | *vlan-ids*) ];  
    }  
}

**Hierarchy Level** [edit `interfaces` *ge-chassis/slot/port unit logical-unit-number*] family

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure Ethernet switching protocol family information for the logical interface.  
  
The remaining statements are explained separately.

**Default** You must configure a logical interface to be able to use the physical device.

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533](#)
- [JUNOS Software Network Interfaces Configuration Guide](#)

## ether-options

|                                 |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ether-options {   802.3ad aex {     lacp {       force-up;       (primary   backup);     }   }   (auto-negotiation   no-auto-negotiation);   configured-flow-control {     rx-buffers (on   off);     tx-buffers (on   off);   }   (flow-control   no-flow-control);   link-mode mode;   (loopback   no-loopback);   speed (auto-negotiation   no-auto-negotiation); }</pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure <b>ether-options</b> properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.</p> <p>The statements are explained separately.</p>                                                                                                                                                                                                                         |
| <b>Default</b>                  | Enabled.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul>                                                                                                                                                                                            |

## eui-64

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | eui-64;                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i> ]                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>              | For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Address on page 2040</a></li></ul>                                                                               |

## family

```
Syntax family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 fibre-channel {
 port-mode (f-port | np-port);
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 targeted-broadcast;
 }
 inet6 {
 address address {
 eui-64
 preferred
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 }
 }
```

**Hierarchy Level** [edit [interfaces interface-name unit logical-unit-number](#)],  
[edit [interfaces interface-range interface-name unit logical-unit-number family](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure protocol family information for the logical interface on the QFX Series product.

**Default** Access interfaces on the QFX Series are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.

You must configure a logical interface to be able to use the physical device.

**Options** See [Table 240 on page 2570](#) for protocol families available on the QFX Series interfaces. Different protocol families support different subsets of the interface types on the QFX Series.

Interface types on the switch are:

- Aggregated Ethernet (**ae**)

- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Table 240: Protocol Families and Supported Interface Types**

| Family             | Description                        | Supported Interface Types |     |     |      |    |    |
|--------------------|------------------------------------|---------------------------|-----|-----|------|----|----|
|                    |                                    | ae                        | lo0 | me0 | vlan | ge | xe |
| ethernet-switching | Ethernet switching protocol family | ✓                         |     | ✓   | ✓    | ✓  | ✓  |
| fibre-channel      | Fibre Channel protocol family      | ✓                         |     |     | ✓    |    | ✓  |
| inet               | IPv4 protocol family               | ✓                         | ✓   | ✓   | ✓    | ✓  | ✓  |

The remaining statements are explained separately.

|                                 |                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.                                |
|                                 | interface-control—To add this statement to the configuration.                         |
| <b>Related Documentation</b>    | • <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a> |
|                                 | • <a href="#">Configuring Link Aggregation on page 2537</a>                           |
|                                 | • <a href="#">Configuring Routed VLAN Interfaces on page 2046</a>                     |
|                                 | • <i>Junos® OS Network Interfaces</i>                                                 |


---

## fibre-channel (Alarm)

---

|                                 |                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | fibre-channel {<br><a href="#">link-down</a> (red   yellow   ignore);<br>}                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">alarm</a> ],<br>[edit chassis <a href="#">interconnect-device</a> <i>name</i> <a href="#">alarm</a> ],<br>[edit chassis <a href="#">node-group</a> <i>name</i> <a href="#">alarm</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                |
| <b>Description</b>              | Configure alarms for a Fibre Channel interface.                                                                                                                                                                  |
| <b>Options</b>                  | The remaining statement is explained separately.—                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 6487</a></li><li>• <a href="#">Interface Alarm Messages on page 6491</a></li></ul>                                              |

## flow-control

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> |
|                                 | <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>                                                                                                                                                                                                          |
|                                 | <ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | Flow control is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 5843</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



---

## force-up

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | force-up;                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> 802.3ad lacp]                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Set the state of the interface as up when the peer has limited LACP capability.                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2534</a></li><li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458</a></li><li>• <i>Junos® OS Network Interfaces</i></li></ul> |

## fpc

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>fpc slot {<br/>  pic <i>pic-number</i> {<br/>    fibre-channel {<br/>      port-range {<br/>        <i>port-range-low</i> <i>port-range-high</i>;<br/>      }<br/>    }<br/>    xe {<br/>      (<i>port</i> <i>port-number</i>   <i>port-range</i> <i>port-range-low</i> <i>port-range-high</i>);<br/>    }<br/>    xle {<br/>      (<i>port</i> <i>port-number</i>   <i>port-range</i> <i>port-range-low</i> <i>port-range-high</i>);<br/>    }<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit chassis],<br>[edit chassis <i>interconnect-device</i> <i>name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><i>slot</i>—Number of the FPC slot. For QFX3500 and QFX3600 devices, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show chassis fpc on page 641</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                            |

## fpc (Interconnect Device)

---

|                                 |                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>fpc slot {<br/>    power (on   off);<br/>}</code>                                                                                                       |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">interconnect-device</a> ]                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                             |
| <b>Description</b>              | Configure the FPC slot number.                                                                                                                                |
| <b>Options</b>                  | <p><code>slot</code>—Number of the FPC slot. For QFX3500 switches, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                            |
| <b>Related Documentation</b>    |                                                                                                                                                               |

## gratuitous-arp-reply

---


|                                 |                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>(gratuitous-arp-reply   no-gratuitous-arp-reply);</code>                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i>],</p> <p>[edit interfaces <a href="#">interface-range</a> <i>interface-range-name</i>]</p>    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                       |
| <b>Description</b>              | Enable processing of ARP updates received via gratuitous ARP reply messages.                                                            |
| <b>Default</b>                  | Updating of the ARP cache is disabled on all Ethernet interfaces.                                                                       |
| <b>Options</b>                  | <p><code>gratuitous-arp-reply</code>—Update the ARP cache.</p> <p><code>no-gratuitous-arp-reply</code>—Do not update the ARP cache.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>      |

## group

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group group-name {<br/>    link-to-monitor interface-name;<br/>    link-to-disable interface-name;<br/>}</pre>                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection]                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure a group of uplink and downlink interfaces for uplink failure detection.                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><i>group-name</i>—Name of the uplink failure detection group.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2445</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2536</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2449</a></li></ul> |

## hold-time (Physical Interface)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <code>hold-time up <i>milliseconds</i> down <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <code>[edit <b>interfaces</b> <i>interface-name</i>],</code><br><code>[edit <b>interfaces</b> <b>interface-range</b> <i>interface-range-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Statement introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 10.4R5 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Specify the hold-time value to use to damp interface transitions. When an interface goes from up to down, it is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly, an interface is not advertised as being up until it has remained up for the hold-time period.                                                                                                                                                                                                                                                                                                                                          |
| <div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>We recommend that you configure the hold-time value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for hold-time depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.</li> <li>The hold-time option is not available for controller interfaces.</li> </ul> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Interface transitions are not damped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p><b>down <i>milliseconds</i></b>—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295 milliseconds</p> <p><b>Default:</b> 0 milliseconds (interface transitions are not damped)</p> <p><b>up <i>milliseconds</i></b>—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295 milliseconds</p> <p><b>Default:</b> 0 milliseconds (interface transitions are not damped)</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><code>advertise-interval</code></li> <li><code>interfaces</code> (for EX Series switches)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## iccp

```
Syntax iccp {
 authentication-key string;
 local-ip-addr local-ip-addr;
 peer ip-address{
 authentication-key string;
 backup-liveness-detection {
 backup-peer-ip ip-address;;
 }
 liveness-detection {
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 local-ip-addr ipv4-address;
 session-establishment-hold-time seconds;
 }
 session-establishment-hold-time seconds;
 traceoptions {
 file <filename> <files number> <match regular-expression> <microsecond-stamp>
 <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## inet (interfaces)

---

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> inet {     address <i>address</i> {         primary;         filter input <i>filter-name</i>;         filter output <i>filter-name</i>;         targeted-broadcast;     } } </pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name unit logical-unit-number</a> family],<br>[edit <a href="#">interfaces interface-range interface-name unit logical-unit-number</a> family]    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                        |
| <b>Description</b>              | Configure the primary IP address for the logical interface.                                                                                                                              |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.                                                                                                            |
| <b>Options</b>                  | The remaining statements are explained separately.—                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> </ul>                                                  |

## inet6 (interfaces)

---

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>inet6 {<br/>    address address {<br/>        eui-64<br/>        preferred<br/>        primary;<br/>        filter input filter-name;<br/>        filter output filter-name;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name unit logical-unit-number</a> family],<br>[edit <a href="#">interfaces interface-range interface-name unit logical-unit-number</a> family]                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                       |
| <b>Description</b>              | Configure the primary IP address for the logical interface.                                                                                                                                             |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.                                                                                                                           |
| <b>Options</b>                  | The remaining statements are explained separately.—                                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.interface-control—To add this statement to the configuration.                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li></ul>                                                                   |



## interconnect-device (Chassis)

```
Syntax interconnect-device {
 alarm {
 (ethernet | management-ethernet) {
 link-down (red | yellow | ignore);
 }
 }
 container-devices {
 device-count number;
 }
 craft-lockout {
 alarm {
 interface-type {
 link-down (red | yellow | ignore);
 }
 }
 container-devices {
 device-count number;
 }
 fpc slot {
 power (on | off);
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }
 fpc slot {
 power (on | off);
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure properties specific to a QFabric system Interconnect device.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Interconnect Devices on page 1183](#)

## interface-range

**Syntax** `interface-range interface-range-name {`  
     `disable;`  
     `description text;`  
     `ether-options {`  
         `802.3ad aex {`  
             `lacp {`  
                 `force-up;`  
             `}`  
         `}`  
     `(auto-negotiation| no-auto-negotiation);`  
     `(flow-control | no-flow-control);`  
     `link-mode mode;`  
     `speed (auto-negotiation | speed);`  
     `}`  
     `hold-time milliseconds down milliseconds;`  
     `member interface-name;`  
     `member-range starting-interface-name to ending-interface-name;`  
     `mtu bytes;`  
     `unit logical-unit-number {`  
         `description text;`  
         `disable;`  
         `family family-name {...}`  
         `(traps | no traps);`  
         `vlan-id vlan-id-number;`  
     `}`  
     `}`

**Hierarchy Level** [edit [interfaces](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX series.

**Description** Group interfaces that share a common configuration profile.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces.

**Options** `interface-range-name`—Name of the interface range.



**NOTE:** You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

**Required Privilege Level** `interface`—To view this statement in the configuration.  
     `interface-control`—To add this statement to the configuration.

- Related Documentation**
- [Understanding Interface Ranges on the QFX Series on page 2432](#)
  - [Interfaces Overview on page 2415](#)
  - [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533](#)
  - *Junos® OS Network Interfaces*

## interface (Multichassis Protection)

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | interface <i>interface-name</i> ;                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">multi-chassis multi-chassis-protection peer</a> ]                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two QFX3500 devices hosting a multichassis link aggregation group (MC-LAG) use this link to pass Interchassis Control Protocol (ICCP) and data traffic. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                            |

## interfaces

---

```
Syntax interfaces {
 aex {
 disable;
 aggregated-ether-options {
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 }
 (flow-control | no-flow-control);
 lacp mode {
 admin-key key;
 periodic interval;
 system-id mac-address;
 }
 link-speed speed;
 loopback;
 no-loopback;
 minimum-links number;
 }
 mc-ae {
 chassis-id chassis-id;
 mc-ae-id mc-ae-id;
 mode (active-active);
 status-control (active | standby);
 }
 description text;
 gratuitous-arp-reply | no-gratuitous-arp-reply
 hold-time down milliseconds up milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
 disable;
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 }
 }
 }
```

```

 primary;
 targeted-broadcast;
 }
 (traps | no traps);
 vlan-id vlan-id-number;
}
vlan-tagging;
}
interface-range interface-range-name {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 }
 }
 }
 (auto-negotiation | no-auto-negotiation);
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (flow-control | no-flow-control);
 link-mode mode;
 speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
 disable;
 description text;
 family family-name {...}
 (traps | no traps);
 vlan-id vlan-id-number;
}
}
lo0 {
 disable;
 description text;
 hold-time milliseconds down milliseconds;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
 disable;
 description text;
 family {
 inet {
 address address {
 primary;
 }
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 }
}

```

```
 }
 (traps | no traps);
 }
}
mex {
 disable;
 description text;
 hold-time milliseconds down milliseconds;
 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 no-gratuitous-arp-request;
 traceoptions;
 traps;
 unit logical-unit-number {
 disable;
 }
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 }
 inet {
 address address {
 primary;
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 }
 traps;
 vlan-id vlan-id-number;
}
vlan-tagging;
vlan {
 disable;
 description text;
 (gratuitous-arp-reply | no-gratuitous-arp-reply);
 hold-time milliseconds down milliseconds;
 mtu bytes;
 no-gratuitous-arp-request;
 traceoptions;
 (traps | no traps);
 unit logical-unit-number {
 description text;
 disable;
 family {
 inet {
 address address {
 primary;
 }
 }
 }
 }
}
```

```

 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
}
}
fc-0/0/port {
 fibrechannel-options {
 bb-sc-n;
 (loopback | no-loopback);
 speed (auto-negotiation | 2g | 4g | 8g);
 }
 unit logical-unit-number {
 disable;
 description text;
 family {
 fibre-channel {
 port-mode np-port;
 }
 }
 (traps | no traps);
 }
}
ge-0/0/port {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 primary;
 }
 }
 }
 (auto-negotiation | no-auto-negotiation);
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
 (flow-control | no-flow-control);
 link-mode mode;
 loopback;
 no-loopback;
 speed (auto-negotiation | speed);
}
gratuitous-arp-reply| no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
 description text;
 disable;
 family {
 ethernet-switching {
 filter input filter-name;

```

```
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
}
inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
}
(traps | no traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bits-per-second priority-cost priority;
 priority-cost priority;
 }
 priority-hold-time seconds;
 route prefix/prefix-length routing-instance instance-name priority-cost priority;
 }
}
virtual-address [addresses];
}
xe-0/0/port {
 disable;
 description text;
 ether-options {
 802.3ad aex {
 lacp {
 force-up;
 (primary | backup);
 }
 }
 }
 configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
 }
}
```



```

}
(flow-control | no-flow-control);
loopback;
no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
 disable;
 description text;
 family {
 ethernet-switching {
 filter input filter-name;
 filter output filter-name;
 native-vlan-id vlan-id;
 port-mode mode;
 reflective-relay;
 vlan {
 members [(all | names | vlan-ids)];
 }
 }
 fibre-channel {
 port-mode (f-port | np-port);
 }
 inet {
 address address {
 primary;
 }
 filter input filter-name;
 filter output filter-name;
 primary;
 targeted-broadcast;
 }
 (traps | no traps);
 vlan-id vlan-id-number;
 }
 vlan-tagging;
}
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the interfaces on the QFX Series.

Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b>aex</b>—Configure an aggregated Ethernet interface.</p> <p><b>xe-0/0/</b><i>port</i>/<b>—</b>Configure a 10-Gigabit Ethernet interface.</p> <p><b>ge-0/0/</b><i>port</i>/<b>—</b>Configure a Gigabit Ethernet interface.</p> <p><b>fc-0/0/</b><i>port</i>/<b>—</b>Configure a Fibre Channel interface.</p> <p><b>meX/</b>—Configure a management interface.</p> <p><b>mc-ae</b>—Configure a multichassis aggregated Ethernet (MC-AE) interface.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Interfaces Overview on page 2415</a></li><li>• <a href="#">Understanding Interface Ranges on the QFX Series on page 2432</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <a href="#">Configuring Link Aggregation on page 2537</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2537</a></li></ul>                                                                         |

---

## lACP (802.3ad)

---

|                                 |                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>lACP {<br/>    force-up;<br/>    (primary   backup);<br/>}</pre>                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> 802.3ad]                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 2537</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2534</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li></ul> |

## lacp (Aggregated Ethernet)

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | lacp (active   passive) {<br>admin-key <i>key</i> ;<br>periodic (fast   slow);<br>system-ID <i>mac-address</i> ;<br>}                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name</a> <a href="#">aggregated-ether-options</a> ]                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2534</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li> </ul> |

## link-to-disable

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-to-disable <i>interface-name</i> ;                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection group <i>group-name</i> ]                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of eight downlink interfaces in a group.                                                                                                                                                  |
| <b>Options</b>                  | <i>interface-name</i> —Name of the downlink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.                                                                                                                                                          |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Uplink Failure Detection on page 2445</a></li> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2536</a></li> <li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2449</a></li> </ul> |

## link-to-monitor

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-to-monitor <i>interface-name</i> ;                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection group <i>group-name</i> ]                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of eight uplink interfaces in a group.                                                                                                                                                              |
| <b>Options</b>                  | <i>interface-name</i> —Name of the uplink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.                                                                                                                                                        |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2445</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2536</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2449</a></li></ul> |

## link-down

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-down (red   yellow   ignore);                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">alarm ethernet</a> ],<br>[edit chassis <a href="#">alarm fibre-channel</a> ],<br>[edit chassis <a href="#">interconnect-device</a> <i>name</i> <a href="#">alarm ethernet</a> ],<br>[edit chassis <a href="#">node-group</a> <i>name</i> <a href="#">alarm fibre-channel</a> ]                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify either red, yellow, or ignore to display when the link is down.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>red</b>—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.</p> <p><b>yellow</b>—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.</p> <p><b>ignore</b>—Suppresses or ignores the alarm.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## link-mode

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>link-mode mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Set the device's link-connection characteristic.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | The <b>full-duplex</b> mode is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>mode</b> —Link characteristic:</p> <ul style="list-style-type: none"><li>• <b>full-duplex</b>—Connection is full duplex.</li><li>• <b>half-duplex</b>—Connection is half duplex.</li><li>• <b>automatic</b>—Link mode is negotiated.</li></ul> <p>If <b>no-auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select only <b>full-duplex</b> or <b>half-duplex</b>. If <i>auto-negotiation</i> is specified in the <b>ether-options</b> option, you can select any mode.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <i>Junos® OS Network Interfaces</i></li></ul>                                                                                                                                                                                                                                                                                                                           |

## link-speed

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-speed <i>speed</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces aex <a href="#">aggregated-ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | For aggregated Ethernet interfaces only, set the required link speed.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>speed</b>—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p>Aggregated Ethernet links on the QFX Series can have one of the following speed values:</p> <ul style="list-style-type: none"> <li>• <b>1g</b>—Links are 1 Gbps.</li> <li>• <b>10g</b>—Links are 10 Gbps.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                       |

## liveness-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>liveness-detection {<br/>  detection-time {<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  minimum-interval <i>milliseconds</i>;<br/>  minimum-receive-interval <i>milliseconds</i>;<br/>  multiplier <i>number</i>;<br/>  no-adaptation;<br/>  transmit-interval {<br/>    minimum-interval <i>milliseconds</i>;<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  version (1   automatic);<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp</b> <b>peer</b> ]                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Include the <b>liveness-detection</b> statement to enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.</p> <p>The remaining statement are explained separately.</p>                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                    |

## local-ip-addr (ICCP)

---


|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>local-ip-addr <i>local-ip-address</i>;</pre>                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp</b> ],<br>[edit protocols <b>iccp</b> <b>peer</b> <i>peer-IP-address</i> ],                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                        |
| <b>Description</b>              | Specify the local IP address of the interchassis link (ICL) interface that Interchassis Control Protocol (ICCP) uses to communicate to the peers hosting a multichassis link aggregation group (MC-LAG). |
| <b>Options</b>                  | <b>local-ip-address</b> —Default local IP address to be used by all peers.                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                      |




## loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)

|                                 |                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (loopback   no-loopback);                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>aggregated-ether-options</b> ],<br>[edit interfaces <i>interface-name</i> <b>ether-options</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                          |
| <b>Description</b>              | For aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet Loopback Capability on page 2535</a></li> </ul>                  |

## management-ethernet (Alarm)

|                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                            | management-ethernet {<br><b>link-down</b> (red   yellow   ignore);<br>}                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                   | [edit chassis <b>alarm</b> ],<br>[edit chassis <b>interconnect-device</b> <i>name</i> <b>alarm</b> ],<br>[edit chassis <b>node-group</b> <i>name</i> <b>alarm</b> ]    |
| <b>Release Information</b>                                                                                                                                                                                                                                                                               | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                       | Configure alarms for a management Ethernet interface.                                                                                                                  |
| <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.</p> </div> </div> |                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                                                                                                                           | The remaining statement is explained separately.—                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                          | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Alarms on page 6487</a></li> <li>• <a href="#">Interface Alarm Messages on page 6491</a></li> </ul> |

## mc-ae

|                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                        | <pre>mc-ae {   chassis-id chassis-id;   mc-ae-id mc-ae-id;   mode (active-active);   status-control (active   standby); }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                               | [edit <a href="#">interfaces aggregated-ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                           | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                   | Specify the multichassis aggregated Ethernet interface configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                       | <p><b>chassis-id</b>—Specify the chassis ID for LACP to calculate the port number of the MC-LAG physical member links.</p> <p><b>mc-ae-id</b>—Specify the identification number of MC-LAG device. The two MC-LAG QFX3500 devices that manage a given MC-LAG must have the same <b>mc-lag-id</b>.</p> <p><b>mode (active   active)</b>—Specify that the MC-LAG is in active-active mode. In this mode, if a member interface of the MC-LAG goes down, traffic can still be forwarded to the QFX3500 devices hosting the MC-LAG using the interchassis link-protection link (ICL-PL). The links from the client-device connected to both of the QFX3500 devices will remain active. Only active-active mode is supported at this time.</p> <p><b>status-control (active   standby)</b>—Specify if a peer is in active or standby mode. In active mode, the peer is considered the primary device, and in standby mode it is considered the secondary device. If the ICL-PL goes down, the peer in standby mode will bring its member links to standby state. If the Interchassis Control Protocol (ICCP) connection goes down, the peer in standby mode will change the LACP system ID to the default value on its member links.</p> |
| <div style="display: flex; align-items: center;">  <p><b>NOTE:</b> You cannot have both peers hosting an MC-LAG be in active or standby mode. One peer must be in active mode, and the other peer must be in standby mode.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                      | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## mc-ae-id

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mc-ae-id <i>mc-ae-id</i>;</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options mc-ae</a> ]                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID. |
| <b>Options</b>                  | <b>Range:</b> 1 through 65535.                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |

## member

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member <i>interface-name</i>;</code>                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-range interface-range-name</a> ]                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Specify the name of the member interface belonging to an interface range on the QFX Series switch.                                                                                                                                                 |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <a href="#">Interfaces Overview on page 2415</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul> |

## member-range

---

|                                 |                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member-range <i>starting-interface-name ending-interface-name</i>;</code>                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <code>interface-range interface-range-name</code> ]                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>starting interface-name ending interface-name</i> —Name of the first member and the name of the last member in the interface sequence.                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Interface Ranges on the QFX Series on page 2432</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <a href="#">Interfaces Overview on page 2415</a></li><li>• <i>Junos® OS Network Interfaces</i></li></ul> |

## minimum-interval (Liveness Detection)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the minimum intervals at which the peer transmits liveness detection requests and then expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. This value represents the minimum interval at which the peer transmits liveness detection requests as well as the minimum interval that the peer expects to receive a reply from a peer with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <b>transmit-interval</b> <b>minimal-interval</b> and <b>minimum-receive-interval</b> statements. |
| <b>Options</b>                  | <i>milliseconds</i> —Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## minimum-receive-interval (Liveness Detection)

---

|                                 |                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-receive-interval <i>milliseconds</i> ;                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp</b> peer <b>liveness-detection</b> ]                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                   |
| <b>Description</b>              | Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. |
| <b>Options</b>                  | <i>milliseconds</i> —Specify the minimum interval value.<br><b>Range:</b> 1 through 255,000                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                 |

## minimum-links

---

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-links <i>number</i> ;                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces aex <b>aggregated-ether-options</b> ]                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                       |
| <b>Description</b>              | For an aggregated Ethernet interface, set the minimum number of links that must be up for the bundle to be labeled up.  |
| <b>Options</b>                  | <i>number</i> —Number of links.<br><b>Range:</b> 1 through 8<br><b>Default:</b> 1                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> </ul>           |

## mode (QFX Series)

---

|                                 |                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode active-active ;                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options mc-ae</a> ]                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both QFX3500 devices. Only active-active mode is supported at this time. |
| <b>Options</b>                  | <b>active-active</b> —Specify that all of the members of the MC-LAG will be active on both QFX3500 devices.                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                    |
| <b>Related Documentation</b>    |                                                                                                                                                                                                                                            |

## mtu

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mtu bytes;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name</a> ],<br>[edit <a href="#">interfaces interface-range interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On a QFX3500 switch, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. |
| <b>Options</b>                  | <p><b>bytes</b> —MTU size.</p> <p><b>Range:</b> 64 through 9216 bytes</p> <p><b>Default:</b> 1514 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## multi-chassis

---

|                                 |                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multi-chassis {<br/>    multi-chassis-protection peer-ip-address {<br/>        interface interface-name;<br/>    }<br/>}</pre>                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure an interchassis link-protection link (ICL-PL) between the two peers. You can configure either an Aggregated Ethernet interface or a 10-Gigabit Ethernet interface to be an ICL-PL.</p> <p>The remaining statements are explained separately.</p> |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —Specify the logical interface name of the peer.                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                       |
| <b>Related Documentation</b>    |                                                                                                                                                                                                                                                               |

## multi-chassis-protection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multi-chassis-protection peer-ip-address {<br/>    interface interface-name;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">multi-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Mutlichassis protection provides link protection between the two peers hosting a multichassis link aggregation group (MC-LAG). If the Interchassis Control Protocol (ICCP) connection is up and the interchassis link (ICL) comes up, the peer configured as standby brings up the multichass aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on one interface for each peer.</p> <p>The remaining statements are explained separately.</p> |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —Specify the logical interface name of the peer.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                |



## multiplier (Liveness Detection)

---

|                                 |                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>multiplier <i>number</i>;</code>                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                       |
| <b>Description</b>              | Configure the number of liveness detection requests not received by the peer before Bidirectional Forwarding Detection (BFD) declares the peer is down. |
| <b>Options</b>                  | <b>number</b> —Maximum allowable number of liveness detection requests missed by the peer.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                     |
| <b>Related Documentation</b>    |                                                                                                                                                         |

## no-adaptation (Liveness Detection)

---

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>no-adaptation;</code>                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                   |
| <b>Description</b>              | Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    |                                                                                                                     |

## no-gratuitous-arp-request

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-gratuitous-arp-request;                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ],<br>[edit <a href="#">interfaces</a> <i>interface-range</i> <i>interface-name</i> ]                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Routed VLAN Interfaces on page 2046</a></li></ul>                                                                                     |

## node-device (Chassis)

```
Syntax node-device name {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 pic pic-number {
 fte {
 port port-number;
 port-range port-range-low port-range-high;
 }
 xe {
 port port-number;
 port-range port-range-low port-range-high;
 }
 }
 }
```

**Hierarchy Level** [edit chassis [node-group](#)]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure properties specific to a Node device in a QFabric system.  
The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation on page 2537](#)

## node-group (Chassis)

```
Syntax node-group name {
 aggregated-devices {
 ethernet {
 device-count number;
 }
 }
 alarm {
 interface-type {
 link-down (ignore | red | yellow);
 }
 }
 container-devices {
 device-count number;
 }
 node-device name {
 fibre-channel {
 port-range {
 port-range-low port-range-high;
 }
 }
 }
 pic pic-number {
 fte {
 port port-number;
 port-range port-range-low port-range-high;
 }
 xe {
 port port-number;
 port-range port-range-low port-range-high;
 }
 }
 routing-engine {
 on-disk-failure {
 disk-failure-action (halt | reboot);
 }
 }
 }
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure properties specific to a Node group.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation on page 2537](#)

---

## on-disk-failure

---

|                                 |                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | on-disk-failure {<br>disk-failure-action (halt   reboot);<br>}                                                                             |
| <b>Hierarchy Level</b>          | [edit chassis routing-engine],<br>[edit chassis node-group name routing-engine],<br>[edit chassis interconnect-device name routing-engine] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                          |
| <b>Description</b>              | Halt or reboot the switch if it detects hard disk errors on the Routing Engine.                                                            |
| <b>Options</b>                  | The remaining statement is explained separately.                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors</i></li></ul> |

## peer (ICCP)

---

**Syntax**    `peer ip-address {  
              authentication-key string;  
              backup-liveness-detection {  
                  backup-peer-ip ip-address;  
              }  
              liveness-detection {  
                  detection-time {  
                      threshold milliseconds;  
                  }  
              minimum-interval milliseconds;  
              minimum-receive-interval milliseconds;  
              multiplier number;  
              no-adaptation;  
              transmit-interval {  
                  minimum-interval milliseconds;  
                  threshold milliseconds;  
              }  
              version (1 | automatic);  
          }  
          local-ip-addr ipv4-address;  
          session-establishment-hold-time seconds;  
          }`

**Hierarchy Level**    [edit protocols [iccp](#)]

**Release Information**    Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description**    Configure the peers hosting a multichassis link aggregation group (MC-LAG). You must configure Interchassis Control Protocol (ICCP) for both peers hosting an MC-LAG.

The remaining statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

## peer (Multichassis)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>peer ip-address {<br/>    interface interface-name;<br/>}</code>                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">multi-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the IP address of the peer that is part of the interchassis link-protection link (ICL-PL). If the Interchassis Control Connection Protocol (ICCP) is up and the interchassis link (ICL) comes up, the peer configured as standby will bring up the MC-AE interfaces shared with the active peer specified by the <b>peer</b> statement. You must specify the physical interface of the peer. |
| <b>Options</b>                  | <b>interface interface-name</b> —Specify the logical interface name of the peer.                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                |

## periodic

|                                 |                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>periodic (fast   slow);</code>                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aex aggregated-ether-options lacp</a> ]                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                   |
| <b>Description</b>              | Configure the interval for periodic transmission of LACP packets.                                                                                                                                                                                   |
| <b>Default</b>                  | <b>fast</b>                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>interval</b> —Interval at which to periodically transmit LACP packets: <ul style="list-style-type: none"> <li><b>fast</b>—Transmit packets every second. This is the default.</li> <li><b>slow</b>—Transmit packets every 30 seconds.</li> </ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li> <li><a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                          |

## pic

---

**Syntax**    `pic pic-number {  
              fibre-channel {  
                  port-range {  
                      port-range-low port-range-high;  
                  }  
              }  
              xe {  
                  (port port-number | port-range port-range-low port-range-high);  
              }  
              xle {  
                  (port port-number | port-range port-range-low port-range-high);  
              }`

**Hierarchy Level**    [edit **chassis fpc slot**]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Options **xe** and **xle** introduced in Junos OS 12.2X50-D20 for the QFX Series.

**Description**    (QFX3500 and QFX3600 standalone switches only) Enable the specified port on the Physical Interface Card (PIC) to perform in the specified operating mode.

**Options**    *pic-number*—Number of the PIC.

- On a QFX3500 standalone switch, specify **0** if the port type is **fibre-channel**, and **2** if the port type is **xle**.
- On a QFX3600 standalone switch, specify **0** if the port type is **xe**, and **1** if the port type is **xle**.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the QSFP+ Port Type on QFX3500 Standalone Switches on page 2543](#)
- [Configuring the Port Type on QFX3600 Standalone Switches on page 2545](#)



## port-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | port-mode (access   tagged-access   trunk);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family <a href="#">ethernet-switching</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure whether an interface on the switch operates in access, tagged access, or trunk mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | All switch interfaces are in access mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>access</b>—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>tagged-access</b>—Have the interface operate in access mode. In this mode, the interface can be in multiple VLANs. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>trunk</b>—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Reflective Relay on page 2044</a></li> <li>• <a href="#">Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## reflective-relay

---

|                                 |                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | reflective-relay;                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> ethernet-switching]                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                     |
| <b>Description</b>              | Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.                                                                                           |
| <b>Default</b>                  | Switch interfaces are not configured for reflective relay.                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Reflective Relay for Use with VEPA Technology on page 1958</a></li><li>• <a href="#">Configuring Reflective Relay on page 2044</a></li></ul> |

## routing-engine

---

|                                 |                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>routing-engine {<br/>  <a href="#">on-disk-failure</a> {<br/>    <a href="#">disk-failure-action</a> (halt   reboot);<br/>  }<br/>}</pre>                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit chassis]<br>[edit chassis <a href="#">interconnect-device</a> <i>name</i> ],<br>[edit chassis <a href="#">node-group</a> <i>name</i> ]                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors</a></li><li>• <a href="#">Junos OS High Availability Configuration Guide</a></li></ul>                                                                     |

---

## session-establishment-hold-time

---

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | session-establishment-hold-time <i>seconds</i> ;                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer</a> ],<br>[edit protocols <a href="#">iccp</a> ]                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                   |
| <b>Description</b>              | Specify the time during which an Interchassis Control Protocol (ICCP) connection must succeed between peers.        |
| <b>Options</b>                  | <b>seconds</b> —Specify the amount of time allotted for a successful ICCP connection.                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |

## speed

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>speed (auto-negotiation   <i>speed</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the speed of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | Autonegotiation is automatically enabled for Gigabit Ethernet interfaces. Autonegotiation is not an option for 10-Gigabit Ethernet interfaces. No explicit action is taken after the autonegotiation is complete or if the negotiation fails. If the <b>autonegotiation</b> statement at the [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ] hierarchy level is enabled, the <b>auto-negotiation</b> option is enabled by default.                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <ul style="list-style-type: none"><li>• <b>auto-negotiation</b>—Automatically negotiate the speed based on the speed of the other end of the link. Autonegotiation is automatically enabled for Gigabit Ethernet interfaces. Autonegotiation is not an option for 10-Gigabit Ethernet interfaces. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.</li><li>• <b>speed</b>—Specify the interface speed. This value sets the speed that is used on the link. If the <b>auto-negotiation</b> statement is enabled on a Gigabit Ethernet interface, configure the value to advertise to the interface at the other end of the link. If you disable autonegotiation on a Gigabit Ethernet interface, you must explicitly configure the speed to 1g.</li></ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">auto-negotiation</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <a href="#">Junos® OS Network Interfaces</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## status-control


---

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>status-control (active   standby);</code>                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <a href="#">aggregated-ether-options</a> <a href="#">mc-ae</a> ]                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                         |
| <b>Description</b>              | Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                   |

## targeted-broadcast


|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | targeted-broadcast;                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> family inet],<br>[edit <b>interfaces</b> <i>interface-range</i> <i>interface-range-name</i> <b>unit</b> <i>logical-unit-number</i> family inet]                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                             |
| <b>Description</b>              | Specify whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.                           |
| <b>Default</b>                  | When this statement is not included, broadcast packets are sent to the Routing Engine only.                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IP Directed Broadcast on an EX Series Switch</i></li> <li>• <i>Configuring IP Directed Broadcast (CLI Procedure)</i></li> <li>• <i>Understanding IP Directed Broadcast for EX Series Switches</i></li> </ul> |

## threshold (Detection Time)

|                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                               | threshold <i>milliseconds</i> ;                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                      | [edit protocols <b>iccp</b> <i>peer</i> <b>liveness-detection</b> <i>detection-time</i> ]                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                  | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                          | Specify the threshold for the adaptation of the detection time for a Bidirectional Forwarding Detection (BFD) session. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent. |
| <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval. values.</p> </div> </div> |                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                                                                                                                                                              | <i>milliseconds</i> — Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                             | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                              |

## traceoptions (Individual Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traceoptions {<br/>    flag <i>flag</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The <b>traceoptions</b> statement for interfaces does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system <b>syslog</b> file in the directory <b>/var/log</b>.</p>                                                                                                                                                     |
|                                 | <div> <b>NOTE:</b> The <b>traceoptions</b> statement is not supported on the QFX3000 QFabric system.</div>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                  | If you do not include this statement, no interface-specific tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All interface tracing operations</li><li>• <b>event</b>—Interface events</li><li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li><li>• <b>media</b>—Interface media changes</li><li>• <b>q921</b>—ISDN Q.921 frames</li><li>• <b>q931</b>—ISDN Q.931 frames</li></ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tracing Operations of an Individual Router or Switch Interface</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                       |

## traceoptions (ICCP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit <a href="#">protocols iccp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | Set Interchassis Control Protocol (ICCP) tracing options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>             | Tracing operations are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. By default, the log file is stored in <b>/var/log/</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file only</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The following are the tracing options:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All tracing operations</li> <li>• <b>config-internal</b>—Trace configuration internals.</li> <li>• <b>general</b>—Trace general events.</li> <li>• <b>normal</b>—All normal events.</li> </ul> <p><b>Default:</b> If you do not specify this option, only unusual or abnormal operations are traced.</p> <ul style="list-style-type: none"> <li>• <b>parse</b>—Trace configuration parsing.</li> <li>• <b>policy</b>—Trace policy operations and actions.</li> </ul> |

- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

---

## transmit-interval (Liveness Detection)

---

**Syntax**     `transmit-interval {  
                  minimum-interval milliseconds;  
                  threshold milliseconds;  
                  }`

**Hierarchy Level**     [edit protocols **iccp peer liveness-detection**]

**Release Information**     Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description**     Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |



---

## traps

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (traps   no-traps);                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> ],<br>[edit interfaces <i>interface-range</i> <i>interface-range-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                        |
| <b>Description</b>              | Enable or disable the sending of SNMP notifications when the state of the connection changes.                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling or Disabling SNMP Notifications on Physical Interfaces</i></li><li>• <i>Enabling or Disabling SNMP Notifications on Logical Interfaces</i></li></ul> |

## unit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> unit <i>logical-unit-number</i> {     family {         ethernet-switching {             filter input <i>filter-name</i>;             filter output <i>filter-name</i>;             native-vlan-id <i>vlan-id</i>;             port-mode <i>mode</i>;             vlan {                 members [ (all   <i>names</i>   <i>vlan-ids</i>) ];             }         }         fibre-channel {             port-mode (f-port   np-port);         }         inet {             address <i>address</i> {                 primary;             }             filter input <i>filter-name</i>;             filter output <i>filter-name</i>;             primary;             targeted-broadcast;         }     } </pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces <i>interface-name</i></a> ],<br>[edit <a href="#">interfaces <i>interface-range</i> <i>interface-range-name</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>logical-unit-number</i></b>—Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,384</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## uplink-failure-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>uplink-failure-detection {   group group-name {     link-to-monitor interface-name;     link-to-disable interface-name;   } }</pre>                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failure information to the downlink interfaces.</p> <p>The remaining statements are explained separately.</p>                                                                                                |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Uplink Failure Detection on page 2445</a></li> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2536</a></li> <li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2449</a></li> </ul> |

## version (Liveness Detection)

---

|                                 |                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | version (1   automatic);                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                              |
| <b>Description</b>              | Configure the Bidirectional Forwarding Detection (BFD) protocol version to detect.                                             |
| <b>Options</b>                  | <p>1—Use BFD protocol version 1.</p> <p><b>automatic</b>—Autodetect the BFD protocol version.</p>                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |

## vlan-id

---

|                            |                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>vlan-id <i>vlan-id-number</i>;</code>                                                                         |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                   |
| <b>Description</b>         | For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface. |



**NOTE:** The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.



|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>vlan-id-number</i> —Valid VLAN identifier.<br><b>Range:</b> 1 through 4094                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-tagging on page 2146</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2537</a></li><li>• <i>Junos® OS Network Interfaces</i></li></ul> |

## vlan-tagging

---

|                                 |                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan-tagging;</code>                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ]<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>              | Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.                                                                      |
| <b>Default</b>                  | VLAN tagging is disabled by default.                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-id on page 2624</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2537</a></li></ul> |

## xe (Port)

|                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                       | <pre>xe {   (port <i>port-number</i>   port-range <i>port-range-low</i> <i>port-range-high</i>); }</pre>                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                              | [edit <a href="#">chassis fpc slot pic pic-number</a> ]                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                                                                          | Statement introduced in Junos OS Release 12.2X50-D20 for the QFX Series.                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                  | (QFX3600 standalone switch only) Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet (xe) type ports.                                                                                                                                                                                                                        |
| <div>  <p><b>CAUTION:</b> The Packet Forwarding Engine on the switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.</p> </div>     |                                                                                                                                                                                                                                                                                                                                                               |
| <div>  <p><b>NOTE:</b> Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.</p> </div> |                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                                                                                      | <p><b><i>port-number</i></b>—Port number on which you want to configure the port type. Valid values are 0 through 15.</p> <p><b><i>port-range-low</i></b>—Lowest-numbered port in the range of ports. The lowest possible value is 0.</p> <p><b><i>port-range-high</i></b>—Highest-numbered port in the range of ports. The highest possible value is 15.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                     | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                        | <ul style="list-style-type: none"> <li><a href="#">Configuring the Port Type on QFX3600 Standalone Switches on page 2545</a></li> </ul>                                                                                                                                                                                                                       |

## xle (Port)

---

**Syntax**    `xle {  
                  (port port-number | port-range port-range-low port-range-high);  
                  }`

**Hierarchy Level**    [edit `chassis fpc slot pic pic-number`]

**Release Information**    Statement introduced in Junos OS Release 12.2X50-D20 for the QFX Series.

**Description**    (QFX3500 and QFX3600 standalone switches only) Configure a specific QSFP+ port or a range of QSFP+ ports to operate as 40-Gigabit Ethernet (*xle*) type ports.



**CAUTION:** The Packet Forwarding Engine on the switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

---

**Options**    *port-number*—Port number on which you want to configure the port type. On a QFX3500 standalone switch, specify a value from 0 through 3. On a QFX3600 standalone switch, specify a value from 0 through 15.

*port-range-low*—Lowest-numbered port in the range of ports. The lowest possible value is 0.

*port-range-high*—Highest-numbered port in the range of ports. The highest possible value is 3 on QFX3500 standalone switches, and 15 on QFX3600 standalone switches.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the QSFP+ Port Type on QFX3500 Standalone Switches on page 2543](#)
- [Configuring the Port Type on QFX3600 Standalone Switches on page 2545](#)

# Administration

- Routine Monitoring on page 2627
- Monitoring Commands on page 2633

## Routine Monitoring

- Monitoring System Process Information on page 2627
- Monitoring System Properties on page 2628
- Monitoring Interface Status and Traffic on page 2629
- Verifying That Layer 3 Logical Interfaces Are Working on page 2630
- Verifying the Status of a LAG Interface on page 2630
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631
- Verifying That Uplink Failure Detection Is Working Correctly on page 2632

## Monitoring System Process Information

- Purpose

View the processes running on the QFX Series.
- Action

To view the software processes running on the QFX Series:  
[edit system]  
  
user@switch> **show system processes**
- Meaning

Table 47 on page 315 summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 241: Summary of System Process Information Output Fields

| Field | Values                     |
|-------|----------------------------|
| PID   | Identifier of the process. |
| Name  | Owner of the process.      |

Table 241: Summary of System Process Information Output Fields (*continued*)

| Field              | Values                                                   |
|--------------------|----------------------------------------------------------|
| State              | Current state of the process.                            |
| CPU Load           | Percentage of the CPU that is being used by the process. |
| Memory Utilization | Amount of memory that is being used by the process.      |
| Start Time         | Time of day when the process started.                    |

- Related Documentation**
- [Monitoring System Properties on page 316](#)
  - [show system uptime on page 1067](#)

## Monitoring System Properties

**Purpose** View system properties such as the name and IP address of a QFX Series product and resource usage.

**Action** To monitor system properties in the CLI, enter the following commands:

- [show system uptime](#)
- [show system users](#)
- [show system storage](#)

**Meaning** [Table 48 on page 316](#) summarizes key output fields in the system properties display.

Table 242: Summary of Key System Properties Output Fields

| Field                      | Values                                                                                                  | Additional Information                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| <b>General Information</b> |                                                                                                         |                                                        |
| Serial Number              | Serial number for a QFX Series product.                                                                 |                                                        |
| Junos OS Version           | Version of Junos OS active on the switch, including whether the software is for domestic or export use. | Export software is for use outside the USA and Canada. |
| Hostname                   | Name of the QFX Series product.                                                                         |                                                        |
| IP Address                 | IP address of the QFX Series product.                                                                   |                                                        |
| Loopback Address           | Loopback address.                                                                                       |                                                        |
| Domain Name Server         | Address of the domain name server.                                                                      |                                                        |



Table 242: Summary of Key System Properties Output Fields (*continued*)

| Field                   | Values                                                                                                                                       | Additional Information                                                           |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Time Zone               | Time zone on the QFX Series product.                                                                                                         |                                                                                  |
| Time                    |                                                                                                                                              |                                                                                  |
| Current Time            | Current system time, in Coordinated Universal Time (UTC).                                                                                    |                                                                                  |
| System Booted Time      | Date and time when the QFX Series product was last booted and how long it has been running.                                                  |                                                                                  |
| Protocol Started Time   | Date and time when the protocols were last started and how long they have been running.                                                      |                                                                                  |
| Last Configured Time    | Date and time when a configuration was last committed. This field also shows the name of the user who issued the last <b>commit</b> command. |                                                                                  |
| Load Average            | CPU load average for 1, 5, and 15 minutes.                                                                                                   |                                                                                  |
| Storage Media           |                                                                                                                                              |                                                                                  |
| Internal Flash Memory   | Usage details of internal flash memory.                                                                                                      |                                                                                  |
| External Flash Memory   | Usage details of external USB flash memory.                                                                                                  |                                                                                  |
| Logged in Users Details |                                                                                                                                              |                                                                                  |
| User                    | Username of any user logged in to the switch.                                                                                                |                                                                                  |
| Terminal                | Terminal through which the user is logged in.                                                                                                |                                                                                  |
| From                    | System from which the user has logged in. A hyphen indicates that the user is logged in through the console.                                 |                                                                                  |
| Login Time              | Time when the user logged in.                                                                                                                | This is the <b>user@switch</b> field in <b>show system users</b> command output. |
| Idle Time               | How long the user has been idle.                                                                                                             |                                                                                  |

- Related Documentation**
- [Monitoring System Process Information on page 315](#)
  - [show system processes on page 983](#)

## Monitoring Interface Status and Traffic

- Purpose** View interface status to monitor interface bandwidth utilization and traffic statistics on the QFX Series product.

- Action**
- To view interface status for all the interfaces, enter [show interfaces xe](#).
  - To view status and statistics for a specific interface, enter [show interfaces xe interface-name](#).
  - To view status and traffic statistics for all interfaces, enter either [show interfaces xe detail](#) or [show interfaces xe extensive](#).

**Meaning** For details about output from the CLI commands, see [show interfaces xe](#).

## Verifying That Layer 3 Logical Interfaces Are Working

**Purpose** After configuring Layer 3 logical interfaces, verify that they are set up properly and transmitting data.

- Action**
- To determine if you have successfully created the logical interfaces and the links are up:

```
[edit interfaces]
user@switch> show interfaces interface-name terse
Interface Admin Link Proto Local Remote
ge-0/0/0 up up
ge-0/0/0.0 up up inet 1.1.1.1/24
ge-0/0/0.1 up up inet 2.1.1.1/24
ge-0/0/0.2 up up inet 3.1.1.1/24
ge-0/0/0.3 up up inet 4.1.1.1/24
ge-0/0/0.4 up up inet 5.1.1.1/24
ge-0/0/0.32767 up up
```

- Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the logical interface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

**Meaning** The output confirms that the logical interfaces have been created and the links are up.

**Related Documentation**

- [Configuring a Layer 3 Logical Interface on page 2537](#)

## Verifying the Status of a LAG Interface

**Purpose** Verify that a link aggregation group (LAG) (**ae0**) has been created on the switch.

- Action** To verify that the **ae0** LAG has been created:

```
[edit interfaces]
show interfaces ae0 terse
```

| Interface | Admin | Link | Proto | Local         | Remote |
|-----------|-------|------|-------|---------------|--------|
| ae0       | up    | up   |       |               |        |
| ae0.0     | up    | up   | inet  | 10.10.10.2/24 |        |

**Meaning** The output confirms that the **ae0** link is up and shows the family and IP address assigned to this link.

- Related Documentation**
- [Configuring Link Aggregation on page 2537](#)
  - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
  - [show lacp statistics interfaces \(View\) on page 2806](#)

## Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 2631](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 2632](#)

### Verifying the LACP Setup

**Purpose** Verify that the LACP has been set up correctly.

**Action** To verify that LACP has been enabled as active on one end:

```
user@switch>show lacp interfaces xe-0/0/0
Aggregated interface: ae0
LACP state:
xe-0/1/0 Actor No Yes No No No Yes Fast Active
xe-0/1/0 PartnerNo Yes No No No Yes Fast Passive
LACP protocol: Receive State Transmit State Mux State
xe-0/1/0 Defaulted Fast periodic Detached
```

**Meaning** This example shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, one side must be set as active in order for the bundled link to be up.

### Verifying That LACP Packets Are Being Exchanged

---

**Purpose** Verify that LACP packets are being exchanged between interfaces.

**Action** Use the **show lacp statistics interfaces *interface-name*** command to display LACP BPDU exchange information.

**show lacp statistics interfaces ae0**

Aggregated interface: ae0

| LACP Statistics: | LACP Rx | LACP Tx | Unknown Rx | Illegal Rx |
|------------------|---------|---------|------------|------------|
| xe-0/0/2         | 1352    | 2035    | 0          | 0          |
| xe-0/0/3         | 1352    | 2056    | 0          | 0          |

**Meaning** The output here shows that the link is up and that PDUs are being exchanged.

- Related Documentation**
- [Configuring Link Aggregation on page 2537](#)
  - [Verifying the Status of a LAG Interface on page 2630](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458](#)
  - [show lacp statistics interfaces \(View\) on page 2806](#)

### Verifying That Uplink Failure Detection Is Working Correctly

**Purpose** Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group : group1
Uplink : xe-0/0/0*
Downlink : xe-0/0/1*
Failure Action : Inactive
```



**NOTE:** The asterisk (\*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group : group1
Uplink : xe-0/0/0
Downlink : xe-0/0/1
Failure Action : Active
```

**Meaning** The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 2445](#)
  - [Configuring Interfaces for Uplink Failure Detection on page 2536](#)
  - [Example: Configuring Interfaces for Uplink Failure Detection on page 2449](#)

## Monitoring Commands

- [monitor interface](#)
- [show iccp](#)
- [show interfaces diagnostics optics](#)
- [show interfaces fabric](#)
- [show interfaces ge](#)
- [show interfaces mc-ae](#)
- [show interfaces statistics fabric](#)
- [show interfaces queue](#)
- [show interfaces queue fabric](#)

- [show interfaces xe](#)
- [show interfaces xle](#)
- [show lacp interfaces](#)
- [show lacp statistics interfaces \(View\)](#)
- [show uplink-failure-detection](#)

## monitor interface

**Syntax** monitor interface  
 <interface-name> | traffic <detail>>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.



**NOTE:** This command is not supported on the QFX3000 QFabric system.

**Options** none—Display real-time statistics for all interfaces.

**detail**—(Optional) With traffic option only, display detailed output.

**interface-name**—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

**traffic**—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

**Additional Information** The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 243 on page 2635](#). The keys are not case-sensitive.

**Table 243: Output Control Keys for the monitor interface Command**

| Key | Action                                                                                                                                                                                                                     |
|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| c   | Clears (returns to zero) the delta counters since <b>monitor interface</b> was started. This does not clear the accumulative counter. To clear the accumulative counter, use the <b>clear interfaces interval</b> command. |
| f   | Freezes the display, halting the display of updated statistics and delta counters.                                                                                                                                         |
| i   | Displays information about a different interface. The command prompts you for the name of a specific interface.                                                                                                            |

**Table 243: Output Control Keys for the monitor interface Command** (*continued*)

| Key      | Action                                                                                                                                                                                         |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n        | Displays information about the next interface. The <b>monitor interface</b> command displays the physical or logical interfaces in the same order as the <b>show interfaces terse</b> command. |
| q or Esc | Quits the command and returns to the command prompt.                                                                                                                                           |
| t        | Thaws the display, resuming the update of the statistics and delta counters.                                                                                                                   |

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 244 on page 2636](#). The keys are not case-sensitive.

**Table 244: Output Control Keys for the monitor interface traffic Command**

| Key      | Action                                                                                                               |
|----------|----------------------------------------------------------------------------------------------------------------------|
| b        | Displays the statistics in units of bytes and bits per second (bps).                                                 |
| c        | Clears (return to 0) the delta counters in the <b>Current Delta</b> column. The statistics counters are not cleared. |
| d        | Displays the <b>Current Delta</b> column (instead of the rate column) in Bps or packets per second (pps).            |
| p        | Displays the statistics in units of packets and packets per second (pps).                                            |
| q or Esc | Quits the command and returns to the command prompt.                                                                 |
| r        | Displays the rate column (instead of the <b>Current Delta</b> column) in Bps and pps.                                |

**Required Privilege Level** trace

**List of Sample Output** [monitor interface \(Physical\) on page 2638](#)  
[monitor interface \(OTN Interface\) on page 2639](#)  
[monitor interface \(Logical\) on page 2640](#)  
[monitor interface \(QFX3500 Switch\) on page 2640](#)  
[monitor interface traffic on page 2641](#)  
[monitor interface traffic \(QFX3500 Switch\) on page 2641](#)  
[monitor interface traffic detail \(QFX3500 Switch\) on page 2642](#)

**Output Fields** [Table 245 on page 2637](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.



Table 245: monitor interface Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>routerl</b>           | Hostname of the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Seconds</b>           | How long the monitor interface command has been running or how long since you last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Time</b>              | Current time (UTC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |
| <b>Delay x/y/z</b>       | Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> <li>• <b>x</b>—Time taken for the last polling (in milliseconds).</li> <li>• <b>y</b>—Minimum time taken across all pollings (in milliseconds).</li> <li>• <b>z</b>—Maximum time taken across all pollings (in milliseconds).</li> </ul>                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Interface</b>         | Short description of the interface, including its name, status, and encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Link</b>              | State of the link: <b>Up</b> , <b>Down</b> , or <b>Test</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels      |
| <b>Current delta</b>     | Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>Local Statistics</b>  | (Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| <b>Remote Statistics</b> | (Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                | All levels      |

Table 245: monitor interface Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| Description        | With the <b>traffic</b> option, displays the interface description configured at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |

## Sample Output

### monitor interface (Physical)

```

user@host> monitor interface so-0/0/0
router1 Seconds: 19 Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
 Input packets: 6045 (0 pps)
 Input bytes: 6290065 (0 bps)
 Output packets: 10376 (0 pps)
 Output bytes: 10365540 (0 bps)
Encapsulation statistics:
 Input keepalives: 1901
 Output keepalives: 1901
 NCP state: Opened
 LCP state: Opened
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 1
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
 LOS count 1
 LOF count 1
 SEF count 1
 ES-S 0
 SES-S 0
SONET statistics:
 BIP-B1 458871

```

```

BIP-B2 460072 [0]
REI-L 465610 [0]
BIP-B3 458978 [0]
REI-P 458773 [0]

```

Received SONET overhead:

```

F1 : 0x00 J0 : 0x00 K1 : 0x00
K2 : 0x00 S1 : 0x00 C2 : 0x00
C2(cmp) : 0x00 F2 : 0x00 Z3 : 0x00
Z4 : 0x00 S1(cmp) : 0x00

```

Transmitted SONET overhead:

```

F1 : 0x00 J0 : 0x01 K1 : 0x00
K2 : 0x00 S1 : 0x00 C2 : 0xcf
F2 : 0x00 Z3 : 0x00 Z4 : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
 Input bytes: 0 (0 bps)
 Output bytes: 0 (0 bps)
 Input packets: 0 (0 pps)
 Output packets: 0 (0 pps)
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 5
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Oversized frames 0
 Packet reject count 0
 DA rejects 0
 SA rejects 0
Output MAC/Filter Statistics:
 Unicast packets 0
 Broadcast packets 0
 Multicast packets 0
 Packet pad count 0
 Packet error count 0
OTN Link 0
 OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
 OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
 OTN OC - Seconds
 LOS 2

```

```

LOF 9
OTN OTU - FEC Statistics
 Corr err ratio N/A
 Corr bytes 0
 Uncorr words 0
OTN OTU - Counters
 BIP 0
 BBE 0
 ES 0
 SES 0
 UAS 422
OTN ODU - Counters
 BIP 0
 BBE 0
 ES 0
 SES 0
 UAS 422
OTN ODU - Received Overhead APSPCC 0-3: 0

```

### monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0
host name Seconds: 16 Time: 15:33:39
 Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics: Current delta
 Input bytes: 0 [0]
 Output bytes: 0 [0]
 Input packets: 0 [0]
 Output packets: 0 [0]
Remote statistics:
 Input bytes: 0 (0 bps) [0]
 Output bytes: 0 (0 bps) [0]
 Input packets: 0 (0 pps) [0]
 Output packets: 0 (0 pps) [0]
Traffic statistics:
 Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

### monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics: Current delta
 Input bytes: 0 (0 bps) [0]
 Output bytes: 0 (0 bps) [0]
 Input packets: 0 (0 pps) [0]
 Output packets: 0 (0 pps) [0]
Error statistics:
 Input errors: 0 [0]
 Input drops: 0 [0]
 Input framing errors: 0 [0]
 Policed discards: 0 [0]
 L3 incompletes: 0 [0]
 L2 channel errors: 0 [0]
 L2 mismatch timeouts: 0 [0]
 Carrier transitions: 0 [0]

```

```

Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
 Unicast packets 0 [0]
 Broadcast packets 0 Multicast packet [0]

Interface warnings:
 o Outstanding LINK alarm

```

### monitor interface traffic

```

user@host> monitor interface traffic
host name Seconds: 15 Time: 12:31:09

Interface Link Input packets (pps) Output packets (pps)
so-1/0/0 Down 0 (0) 0 (0)
so-1/1/0 Down 0 (0) 0 (0)
so-1/1/1 Down 0 (0) 0 (0)
so-1/1/2 Down 0 (0) 0 (0)
so-1/1/3 Down 0 (0) 0 (0)
t3-1/2/0 Down 0 (0) 0 (0)
t3-1/2/1 Down 0 (0) 0 (0)
t3-1/2/2 Down 0 (0) 0 (0)
t3-1/2/3 Down 0 (0) 0 (0)
so-2/0/0 Up 211035 (1) 36778 (0)
so-2/0/1 Up 192753 (1) 36782 (0)
so-2/0/2 Up 211020 (1) 36779 (0)
so-2/0/3 Up 211029 (1) 36776 (0)
so-2/1/0 Up 189378 (1) 36349 (0)
so-2/1/1 Down 0 (0) 18747 (0)
so-2/1/2 Down 0 (0) 16078 (0)
so-2/1/3 Up 0 (0) 80338 (0)
at-2/3/0 Up 0 (0) 0 (0)
at-2/3/1 Down 0 (0) 0 (0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

```

### monitor interface traffic (QFX3500 Switch)

```

user@switch> monitor interface traffic
switch Seconds: 7 Time: 16:04:37

Interface Link Input packets (pps) Output packets (pps)
ge-0/0/0 Down 0 (0) 0 (0)
ge-0/0/1 Up 392187 (0) 392170 (0)
ge-0/0/2 Down 0 (0) 0 (0)
ge-0/0/3 Down 0 (0) 0 (0)
ge-0/0/4 Down 0 (0) 0 (0)
ge-0/0/5 Down 0 (0) 0 (0)
ge-0/0/6 Down 0 (0) 0 (0)
ge-0/0/7 Down 0 (0) 0 (0)
ge-0/0/8 Down 0 (0) 0 (0)
ge-0/0/9 Up 392184 (0) 392171 (0)
ge-0/0/10 Down 0 (0) 0 (0)
ge-0/0/11 Down 0 (0) 0 (0)
ge-0/0/12 Down 0 (0) 0 (0)
ge-0/0/13 Down 0 (0) 0 (0)
ge-0/0/14 Down 0 (0) 0 (0)

```

|           |      |        |     |         |     |
|-----------|------|--------|-----|---------|-----|
| ge-0/0/15 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/16 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/17 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/18 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/19 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/20 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/21 | Down | 0      | (0) | 0       | (0) |
| ge-0/0/22 | Up   | 392172 | (0) | 392187  | (0) |
| ge-0/0/23 | Up   | 392185 | (0) | 392173  | (0) |
| vcp-0     | Down | 0      |     | 0       |     |
| vcp-1     | Down | 0      |     | 0       |     |
| ae0       | Down | 0      | (0) | 0       | (0) |
| bme0      | Up   | 0      |     | 1568706 |     |

## monitor interface traffic detail (QFX3500 Switch)

```
user@switch> monitor interface traffic detail
switch
```

Seconds: 74

Time: 16:03:02

| Interface   | Link | Input packets | (pps) | Output packets | (pps) |
|-------------|------|---------------|-------|----------------|-------|
| Description |      |               |       |                |       |
| ge-0/0/0    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1    | Up   | 392183        | (0)   | 392166         | (0)   |
| ge-0/0/2    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8    | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9    | Up   | 392181        | (0)   | 392168         | (0)   |
| ge-0/0/10   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21   | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22   | Up   | 392169        | (0)   | 392184         | (1)   |
| ge-0/0/23   | Up   | 392182        | (0)   | 392170         | (0)   |
| vcp-0       | Down | 0             |       | 0              |       |
| vcp-1       | Down | 0             |       | 0              |       |
| ae0         | Down | 0             | (0)   | 0              | (0)   |
| bme0        | Up   | 0             |       | 1568693        |       |

## show iccp

**Syntax** `show iccp <brief | detail>`

**Release Information** Command introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Display Interchassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.

**Options** none—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

brief—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

detail—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

**Required Privilege Level** view

**Related Documentation**

**List of Sample Output** [show iccp on page 2644](#)

**Output Fields** [Table 246 on page 2643](#) lists the output fields for the **show iccp** command. Output fields are listed in the approximate order in which they appear.

**Table 246: show iccp**

| Field Name                            | Field Description                                                                                       |
|---------------------------------------|---------------------------------------------------------------------------------------------------------|
| Redundancy Group Information for peer | Aggregated Ethernet interface name.                                                                     |
| TCP Connection                        | Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.                     |
| Liveness Detection                    | Specifies if liveness detection, also known as Bidirectional Forwarding Detection (BFD), is up or down. |
| Client Application                    | Specifies information regarding the state of the MCSNOOPD and ESWD client applications.                 |

## Sample Output

`show iccp`

```
user@switch> show iccp
Redundancy Group Information for peer 3.3.3.2
 TCP Connection : Established
 Liveliness Detection : Up

Client Application: MCSN00PD

Client Application: eswd
```



## show interfaces diagnostics optics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces diagnostics optics <i>interface-name</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, and QSFP+ optical transceivers installed in a QFX Series product. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li> <li>• <a href="#">Installing a Transceiver in a QFX Series Device</a></li> <li>• <a href="#">Removing a Transceiver from a QFX Series Device</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver) on page 2649</a><br><a href="#">show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver) on page 2650</a>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | lists the output fields for the <code>show interfaces diagnostics optics</code> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 247: show interfaces diagnostics optics Output Fields**

| Field Name                           | Field Description                                                                                                                                                  |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface                   | Displays the name of the physical interface.                                                                                                                       |
| Laser bias current                   | Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents. |
| Laser output power                   | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                                         |
| Module temperature                   | Displays the temperature, in Celsius and Fahrenheit.                                                                                                               |
| Module voltage                       | Displays the voltage, in volts.                                                                                                                                    |
| (Not available for XFP transceivers) |                                                                                                                                                                    |

Table 247: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                           | Field Description                                                                                             |
|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power</b><br>(Not available for SFP and SFP+ transceivers)               | Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).          |
| <b>Receiver signal average optical power</b><br>(Not available for XFP transceivers) | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm). |
| <b>Laser bias current high alarm</b>                                                 | Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .                         |
| <b>Laser bias current low alarm</b>                                                  | Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .                          |
| <b>Laser bias current high warning</b>                                               | Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .                       |
| <b>Laser bias current low warning</b>                                                | Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .                        |
| <b>Laser output power high alarm</b>                                                 | Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .                               |
| <b>Laser output power low alarm</b>                                                  | Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .                                |
| <b>Laser output power high warning</b>                                               | Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .                             |
| <b>Laser output power low warning</b>                                                | Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .                              |
| <b>Module temperature high alarm</b>                                                 | Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .                               |
| <b>Module temperature low alarm</b>                                                  | Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .                                |
| <b>Module temperature high warning</b>                                               | Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> .                             |
| <b>Module temperature low warning</b>                                                | Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .                              |
| <b>Module voltage high alarm</b><br>(Not available for XFP transceivers)             | Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .                                   |
| <b>Module voltage low alarm</b><br>(Not available for XFP transceivers)              | Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .                                    |
| <b>Module voltage high warning</b><br>(Not available for XFP transceivers)           | Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .                                 |
| <b>Module voltage low warning</b><br>(Not available for XFP transceivers)            | Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .                                  |
| <b>Laser rx power high alarm</b>                                                     | Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .                              |

Table 247: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                 | Field Description                                                                                                                                                                  |
|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Laser rx power low alarm                                                   | Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .                                                                                                    |
| Laser rx power high warning                                                | Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .                                                                                                 |
| Laser rx power low warning                                                 | Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .                                                                                                  |
| Laser bias current high alarm threshold                                    | Displays the vendor-specified threshold for the laser bias current high alarm.                                                                                                     |
| Module not ready alarm<br>(Not available for SFP and SFP+ transceivers)    | Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault.                                       |
| Module power down alarm<br>(Not available for SFP and SFP+ transceivers)   | Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module is in a limited power mode, low for normal operation.          |
| Tx data not ready alarm<br>(Not available for SFP and SFP+ transceivers)   | Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .                                              |
| Tx not ready alarm<br>(Not available for SFP and SFP+ transceivers)        | Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .                                                   |
| Tx laser fault alarm<br>(Not available for SFP and SFP+ transceivers)      | Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .                                                                                      |
| Tx CDR loss of lock alarm<br>(Not available for SFP and SFP+ transceivers) | Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .     |
| Rx not ready alarm<br>(Not available for SFP and SFP+ transceivers)        | Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .                                                    |
| Rx loss of signal alarm<br>(Not available for SFP and SFP+ transceivers)   | Receive loss of signal alarm. When <b>on</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> . |
| Rx CDR loss of lock alarm<br>(Not available for SFP and SFP+ transceivers) | Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                                 |
| Laser bias current low alarm threshold                                     | Displays the vendor-specified threshold for the laser bias current low alarm.                                                                                                      |
| Laser bias current high warning threshold                                  | Displays the vendor-specified threshold for the laser bias current high warning.                                                                                                   |

Table 247: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                    | Field Description                                                                |
|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Laser bias current low warning threshold                                      | Displays the vendor-specified threshold for the laser bias current low warning.  |
| Laser output power high alarm threshold                                       | Displays the vendor-specified threshold for the laser output power high alarm.   |
| Laser output power low alarm threshold                                        | Displays the vendor-specified threshold for the laser output power low alarm.    |
| Laser output power high warning threshold                                     | Displays the vendor-specified threshold for the laser output power high warning. |
| Laser output power low warning threshold                                      | Displays the vendor-specified threshold for the laser output power low warning.  |
| Module temperature high alarm threshold                                       | Displays the vendor-specified threshold for the module temperature high alarm.   |
| Module temperature low alarm threshold                                        | Displays the vendor-specified threshold for the module temperature low alarm.    |
| Module temperature high warning threshold                                     | Displays the vendor-specified threshold for the module temperature high warning. |
| Module temperature low warning threshold                                      | Displays the vendor-specified threshold for the module temperature low warning.  |
| Module voltage high alarm threshold<br>(Not available for XFP transceivers)   | Displays the vendor-specified threshold for the module voltage high alarm.       |
| Module voltage low alarm threshold<br>(Not available for XFP transceivers)    | Displays the vendor-specified threshold for the module voltage low alarm.        |
| Module voltage high warning threshold<br>(Not available for XFP transceivers) | Displays the vendor-specified threshold for the module voltage high warning.     |
| Module voltage low warning threshold<br>(Not available for XFP transceivers)  | Displays the vendor-specified threshold for the module voltage low warning.      |
| Laser rx power high alarm threshold                                           | Displays the vendor-specified threshold for the laser Rx power high alarm.       |
| Laser rx power low alarm threshold                                            | Displays the vendor-specified threshold for the laser Rx power low alarm.        |
| Laser rx power high warning threshold                                         | Displays the vendor-specified threshold for the laser Rx power high warning.     |

Table 247: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                           | Field Description                                                           |
|--------------------------------------|-----------------------------------------------------------------------------|
| Laser rx power low warning threshold | Displays the vendor-specified threshold for the laser Rx power low warning. |

## Sample Output

### show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver)

```

user@host> show interfaces diagnostics optics xe-0/0/1
Physical interface: xe-0/0/1
 Laser bias current : 4.968 mA
 Laser output power : 0.4940 mW / -3.06 dBm
 Module temperature : 27 degrees C / 81 degrees F
 Module voltage : 3.2310 V
 Receiver signal average optical power : 0.0000
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : On
 Laser rx power high warning : Off
 Laser rx power low warning : On
 Laser bias current high alarm threshold : 10.500 mA
 Laser bias current low alarm threshold : 2.000 mA
 Laser bias current high warning threshold : 9.000 mA
 Laser bias current low warning threshold : 2.500 mA
 Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
 Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
 Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
 Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : -5 degrees C / 23 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.630 V
 Module voltage low alarm threshold : 2.970 V
 Module voltage high warning threshold : 3.465 V
 Module voltage low warning threshold : 3.135 V
 Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
 Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
 Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
 Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

**show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver)**

```

user@host> show interfaces diagnostics optics node1:xe-0/0/1
Physical interface: node1:xe-0/0/1
 Laser bias current : 4.968 mA
 Laser output power : 0.4940 mW / -3.06 dBm
 Module temperature : 27 degrees C / 81 degrees F
 Module voltage : 3.2310 V
 Receiver signal average optical power : 0.0000
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : On
 Laser rx power high warning : Off
 Laser rx power low warning : On
 Laser bias current high alarm threshold : 10.500 mA
 Laser bias current low alarm threshold : 2.000 mA
 Laser bias current high warning threshold : 9.000 mA
 Laser bias current low warning threshold : 2.500 mA
 Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
 Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
 Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
 Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : -5 degrees C / 23 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.630 V
 Module voltage low alarm threshold : 2.970 V
 Module voltage high warning threshold : 3.465 V
 Module voltage low warning threshold : 3.135 V
 Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
 Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
 Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
 Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## show interfaces fabric

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces fabric &lt;interface-name&gt; &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   instance-name)&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display status information about the specified fabric interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>interface-name</b>—(QFabric systems only) Either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   instance-name)</b>—(Optional) Display all routing instances or the name of an individual routing instance.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 1160</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces fabric on page 2658</a></p> <p><a href="#">show interfaces fabric brief on page 2658</a></p> <p><a href="#">show interfaces fabric detail on page 2667</a></p> <p><a href="#">show interfaces fabric extensive on page 2668</a></p> <p><a href="#">show interfaces fabric terse on page 2670</a></p> <p><a href="#">show interfaces fabric device-name on page 2670</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Output Fields** Table 212 on page 2218 lists the output fields for the **show interfaces fabric** command. Output fields are listed in the approximate order in which they appear.

**Table 248: show interfaces fabric Output Fields**

| Field Name                | Field Description                                                                                                    | Level of Output       |
|---------------------------|----------------------------------------------------------------------------------------------------------------------|-----------------------|
| <b>Physical Interface</b> |                                                                                                                      |                       |
| Physical interface        | Name of the physical interface.                                                                                      | All levels            |
| Enabled                   | State of the interface.                                                                                              | All levels            |
| Type                      | Physical interface type; for example, Ethernet.                                                                      | All levels            |
| Interface index           | Index number of the physical interface, which reflects its initialization sequence.                                  | detail extensive none |
| SNMP ifIndex              | SNMP index number for the physical interface.                                                                        | detail extensive none |
| Link-level type           | Encapsulation being used on the physical interface.                                                                  | All levels            |
| MTU                       | Maximum transmission unit size on the physical interface.                                                            | All levels            |
| Clocking                  | Reference clock source.                                                                                              | detail                |
| Speed                     | Speed at which the interface is running.                                                                             | All levels            |
| Duplex                    | Duplex mode of the interface, either Full-Duplex or Half-Duplex.                                                     | All levels            |
| MAC-REWRITE Error         | Specifies if the encapsulation of the packet has been changed.                                                       | none                  |
| BPDU Error                | Specifies if a BPDU has been received on a blocked interface.                                                        | none                  |
| Loopback                  | Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.                     | All levels            |
| Source filtering          | Source filtering status: Enabled or Disabled.                                                                        | All levels            |
| Flow control              | Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured. | All levels            |
| Device flags              | Information about the physical device.                                                                               | All levels            |
| Interface flags           | Information about the interface.                                                                                     | All levels            |
| CoS queues                | Number of CoS queues configured.                                                                                     | detail extensive none |
| Hold-Times                | Current interface hold-time up and hold-time down, in milliseconds.                                                  | detail                |
| Current address           | Configured MAC address.                                                                                              | detail extensive none |



Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output       |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Hardware address        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail extensive none |
| Last flapped            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                              | detail extensive none |
| Statistics last cleared | Date, time, and how long ago the statistics for the interface were cleared. The format is <b>Statistics last cleared: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>2010-05-17 07:51:28 PDT (00:04:33 ago)</b> .                                                                                                                                                                                                                                                                      | detail extensive      |
| Traffic statistics      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                    | detail extensive      |
| IPv6 transit statistics | <p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail extensive      |

Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Input errors | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• Errors—Sum of the incoming frame aborts and FCS errors.</li> <li>• Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• Framing errors—Number of packets received with an invalid frame checksum (FCS).</li> <li>• Runt—Number of frames received that are smaller than the runt threshold.</li> <li>• Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <code>ignore-l3-incompletes</code> statement.</li> <li>• L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• Resource errors—Sum of transmit drops.</li> </ul> | extensive       |

Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output       |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Output errors                    | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>Errors—Sum of the outgoing frame aborts and FCS errors.</li> <li>Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces.</li> <li>MTU errors—Number of packets whose size exceeded the MTU of the interface.</li> <li>Resource errors—Sum of transmit drops.</li> </ul> | extensive             |
| Egress queues                    | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail extensive      |
| Queue counters                   | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>Queued packets—Number of queued packets.</li> <li>Transmitted packets—Number of transmitted packets.</li> <li>Dropped packets—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail extensive      |
| Input rate                       | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | None specified        |
| Output rate                      | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | None specified        |
| Active alarms and Active defects | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li><b>None</b>—There are no active defects or alarms.</li> <li><b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive none |

Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| MAC statistics                         | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets.</li> <li>CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>MAC control frames—Number of MAC control frames.</li> <li>MAC pause frames—Number of MAC control frames with pause operational code.</li> <li>Oversized frames—Number of packets that exceed the configured MTU.</li> <li>Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | extensive       |
| Packet Forwarding Engine Configuration | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>Destination slot—FPC slot number.</li> <li>CoS transmit queue—Queue number and its associated user-configured forwarding class name.</li> <li>Bandwidth %—Percentage of bandwidth allocated to the queue.</li> <li>Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>Priority—Queue priority: low or high.</li> <li>Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | extensive       |

---

**Logical Interface**


---

Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output       |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Item              | Type of QFabric system component being viewed. Possible values include Node group, Interconnect device, Fabric control, Fabric manager, Diagnostic routing engine, and Ungrouped Node device.                                                                                                                                                                                                      | none                  |
| Identifier        | Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.                                                                                                                                                                                                                                                                   | none                  |
| Connection        | Status of a QFabric system component: either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the listed component.                                                                                                                                                                                                                | none                  |
| Configuration     | Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured, Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown.                                                                                                                                                                     | none                  |
| Node group        | Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.                      | none                  |
| Fabric control    | Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine. | none                  |
| Logical interface | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                     | All levels            |
| Index             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                 | detail extensive none |
| SNMP ifIndex      | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                             | detail extensive none |
| Flags             | Information about the logical interface.<br><br>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | All levels            |
| Encapsulation     | Encapsulation method used on the logical interface.                                                                                                                                                                                                                                                                                                                                                | All levels            |

Table 248: show interfaces fabric Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output       |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Traffic statistics     | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail extensive      |
| Local statistics       | <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                 | detail extensive      |
| Transit statistics     | <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                 | detail extensive      |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is inet, the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                                                                                                       | brief                 |
| Generation             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                      | detail extensive      |
| Route table            | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                      | detail extensive none |

## Sample Output

### show interfaces fabric

```

user@switch> show interfaces fabric
Item Identifier Connection Configuration
Node group
 BBAK3775 Connected Configured
 NW-NG-0 Connected Configured
 P2659-C Connected Configured
 ptor-0 Connected Configured
Fabric control
 FC-0 Connected Configured
 FC-1 Connected Configured

```

### show interfaces fabric brief

```

user@switch> show interfaces fabric brief

```

Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up  
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
 Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface BBAK0372:fte-0/1/0.32768  
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
 eth-switch

Physical interface: BBAK0372:fte-0/1/2, Enabled, Physical link is Up  
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
 Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface BBAK0372:fte-0/1/2.32768  
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
 eth-switch

Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up  
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
 Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x4000

Logical interface BBAK0394:fte-0/1/0.32768  
 Flags: SNMP-Traps Encapsulation: ENET2  
 eth-switch

Physical interface: BBAK0394:fte-0/1/2, Enabled, Physical link is Up  
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
 Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x4000

Logical interface BBAK0394:fte-0/1/2.32768  
 Flags: SNMP-Traps Encapsulation: ENET2  
 eth-switch

Physical interface: BBAK3775:bme0, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,  
 Speed: Unspecified  
 Device flags : Present Running

Logical interface BBAK3775:bme0.0  
 Flags: LinkAddress 0-0 Encapsulation: ENET2  
 inet 128.0.0.1/2  
 128.0.0.16/2  
 128.0.32.0/2  
 tnp 0x10

Logical interface BBAK3775:bme0.1  
 Flags: LinkAddress 0-0 Encapsulation: ENET2  
 inet 128.0.0.13/2  
 128.0.130.0/2

Logical interface BBAK3775:bme0.2  
 Flags: Encapsulation: ENET2  
 inet 128.0.0.13/8  
 128.0.130.0/8

169.254.128.13/16  
169.254.193.0/16

Physical interface: BBAK3775:qfabric, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,  
Speed: Unspecified  
Device flags : Present Running  
Interface flags: SNMP-Traps

Logical interface BBAK3775:qfabric.0  
Flags: SNMP-Traps Encapsulation: ENET2  
inet  
mpls  
eth-switch

Physical interface: BBAK3775:vcp0, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:  
1000mbps  
Device flags : Present Running

Logical interface BBAK3775:vcp0.32769  
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:vcp1, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:  
1000mbps  
Device flags : Present Running

Logical interface BBAK3775:vcp1.32768  
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:vcp2, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:  
1000mbps  
Device flags : Present Running

Logical interface BBAK3775:vcp2.32768  
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:fte-0/1/0, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface BBAK3775:fte-0/1/0.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: EE3093:fte-0/1/0, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface EE3093:fte-0/1/0.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: EE3093:fte-0/1/2, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,



Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface EE3093:fte-0/1/2.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/0.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/4.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/6.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/13.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/15.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled

```
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/2.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/7.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/10.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
 Speed: Unspecified
 Device flags : Present Running

Logical interface IC-WS001:bme0.0
 Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.32.0 --> 0/0

Logical interface IC-WS001:bme0.1
 Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.0.7/2
 128.0.130.2/2

Logical interface IC-WS001:bme0.2
 Flags: Encapsulation: ENET2
 inet 128.0.0.7/8
 128.0.130.2/8
 169.254.128.7/16
 169.254.193.1/16

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

Logical interface IC-WS001:bme1.0
 Flags: Encapsulation: ENET2
 inet 128.0.0.1/2
 128.0.0.4/2
 128.0.0.16/2
 128.0.0.17/2
 128.0.0.24/2
```

```

128.0.0.25/2
128.0.0.26/2
128.0.0.28/2
128.0.0.29/2
128.0.0.31/2
tnp 0x4

```

```

Physical interface: IC-WS001:qfabric, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Logical interface IC-WS001:qfabric.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet
 mpls
 eth-switch

```

```

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000Mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000Mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000Mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000Mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000Mbps
 Device flags : Present Running

```

```

Logical interface IC-WS001:vcp0.32769
 Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000Mbps
 Device flags : Present Running

```

```

Logical interface IC-WS001:vcp1.32768
 Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:

```

```
1000mbps
Device flags : Present Running

Logical interface IC-WS001:vcp2.32768
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
Device flags : Present Running

Logical interface IC-WS001:vcp3.32768
Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
Device flags : Present Running

Logical interface IC-WS001:vcp4.32768
Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:bme0, Enabled, Physical link is Up
Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified
Device flags : Present Running

Logical interface NW-NG-0:bme0.0
Flags: LinkAddress 0-0 Encapsulation: ENET2
inet 128.0.0.1/2
 128.0.0.5/2
 128.0.32.0/2
tnp 0x5

Logical interface NW-NG-0:bme0.1
Flags: LinkAddress 0-0 Encapsulation: ENET2
inet 128.0.0.9/2
 128.0.128.4/2

Logical interface NW-NG-0:bme0.2
Flags: Encapsulation: ENET2
inet 128.0.0.9/8
 128.0.128.68/8
 169.254.128.9/16
 169.254.192.34/16

Physical interface: NW-NG-0:qfabric, Enabled, Physical link is Up
Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified
Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface NW-NG-0:qfabric.0
Flags: SNMP-Traps Encapsulation: ENET2
inet
mpls
eth-switch

Physical interface: NW-NG-0:vcp0, Enabled, Physical link is Up
Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000mbps
```

```

Device flags : Present Running

Logical interface NW-NG-0:vcp0.32769
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:vcp1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface NW-NG-0:vcp1.32768
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:bme0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified
 Device flags : Present Running

Logical interface P2659-C:bme0.0
 Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.0.1/2
 128.0.0.16/2
 128.0.32.0/2
 tnp 0x10

Logical interface P2659-C:bme0.1
 Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.0.8/2
 128.0.130.4/2

Logical interface P2659-C:bme0.2
 Flags: Encapsulation: ENET2
 inet 128.0.0.8/8
 128.0.130.4/8
 169.254.128.8/16
 169.254.193.2/16

Physical interface: P2659-C:qfabric, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified
 Device flags : Present Running
 Interface flags: SNMP-Traps

Logical interface P2659-C:qfabric.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet
 mpls
 eth-switch

Physical interface: P2659-C:vcp0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface P2659-C:vcp0.32769
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:vcp1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

```

Logical interface P2659-C:vcp1.32768  
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:vcp2, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps  
Device flags : Present Running

Logical interface P2659-C:vcp2.32768  
Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface P2659-C:fte-0/1/2.32768  
Flags: SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: ptor-0:bme0, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified  
Device flags : Present Running

Logical interface ptor-0:bme0.0  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
inet 128.0.0.1/2  
128.0.0.17/2  
128.0.32.0/2  
tnp 0x11

Logical interface ptor-0:bme0.1  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
inet 128.0.0.16/2  
128.0.130.18/2

Logical interface ptor-0:bme0.2  
Flags: Encapsulation: ENET2  
inet 128.0.0.16/8  
128.0.130.18/8  
169.254.128.16/16  
169.254.193.9/16

Physical interface: ptor-0:qfabric, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified  
Device flags : Present Running  
Interface flags: SNMP-Traps

Logical interface ptor-0:qfabric.0  
Flags: SNMP-Traps Encapsulation: ENET2  
inet  
mpls  
eth-switch

Physical interface: ptor-0:vcp0, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps

```

Device flags : Present Running

Logical interface ptor-0:vcp0.32769
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcp1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface ptor-0:vcp1.32768
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcp2, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface ptor-0:vcp2.32768
 Flags: LinkAddress 0-0 Encapsulation: ENET2

```

#### show interfaces fabric detail

```

user@switch> show interfaces fabric detail
Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up
 Interface index: 49165, SNMP ifIndex: 1212678666, Generation: 140
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 84:18:88:d1:fa:1f, Hardware address: 84:18:88:d1:fa:1f
 Last flapped : 2012-11-09 21:36:41 UTC (4d 00:23 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 14256654 0 bps
 Output bytes : 9618986 0 bps
 Input packets : 90511 0 pps
 Output packets : 60101 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets : 0
Egress queues: 12 supported, 5 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 3 fcoe         | 0              | 0                   | 0               |
| 4 no-loss      | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |
| 8 mcast        | 0              | 0                   | 0               |

```

Active alarms : None
Active defects : None

```

```

Logical interface BBAK0372:fte-0/1/0.32768 (Index 71) (SNMP ifIndex 1212678667)
(Generation 136)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 12450372
 Output bytes : 11986557
 Input packets: 90510
 Output packets: 62750
 Local statistics:
 Input bytes : 12450372
 Output bytes : 11986557
 Input packets: 90510
 Output packets: 62750
 Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 Protocol eth-switch, MTU: 0, Generation: 163, Route table: 0

```

#### show interfaces fabric extensive

```
user@switch> show interfaces fabric extensive
```

```

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up
 Interface index: 49176, SNMP ifIndex: 1209008655, Generation: 155
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06
 Last flapped : 2012-11-13 23:53:30 UTC (00:53:20 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 91179 0 bps
 Output bytes : 361268221791 952985992 bps
 Input packets: 590 0 pps
 Output packets: 2580487185 850880 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
 Resource errors: 0
 Output errors:
 Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
 Egress queues: 12 supported, 5 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets

 0 fabric_fcset 0 0 0
 1 fabric_fcset 0 0 0

```



```

2 fabric_fcset 0 0 0
3 fabric_fcset 0 0 0
4 fabric_fcset 0 0 0
5 fabric_fcset 0 0 0
6 fabric_fcset 0 0 0
7 fabric_fcset 0 0 0
8 fabric_fcset 0 2582632925 0
9 fabric_fcset 0 0 0
10 fabric_fcset 0 0
0
11 fabric_fcset 0 0
0
Active alarms : None
Active defects : None
MAC statistics:
 Receive Transmit
 Total octets 91179 361268221791
 Total packets 590 2580487185
 Unicast packets 590 2580487185
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
MAC Priority Flow Control Statistics:
Priority : 0 0 0
Priority : 1 0 0
Priority : 2 0 0
Priority : 3 0 0
Priority : 4 0 0
Priority : 5 0 0
Priority : 6 0 0
Priority : 7 0 0
Packet Forwarding Engine configuration:
Destination slot: 0
Direction : Output
CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 5 2000000000 5 0 low
none
3 fcoe 35 14000000000 35 0 low
none
4 no-loss 35 14000000000 35 0 low
none
7 network-control 5 2000000000 5 0 low
none

```

```

 8 mcast 20 8000000000 20 0 low
none

```

Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1209008656)  
(Generation 150)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

```

Input bytes : 79496
Output bytes : 179860
Input packets: 590
Output packets: 948

```

Local statistics:

```

Input bytes : 79496
Output bytes : 179860
Input packets: 590
Output packets: 948

```

Transit statistics:

```

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

```

Protocol eth-switch, MTU: 0, Generation: 178, Route table: 0

#### show interfaces fabric terse

```
user@switch> show interfaces fabric terse
```

| Item           | Identifier | Connection | Configuration |
|----------------|------------|------------|---------------|
| Node group     |            |            |               |
| BBAK3775       |            | Connected  | Configured    |
| NW-NG-0        |            | Connected  | Configured    |
| P2659-C        |            | Connected  | Configured    |
| ptor-0         |            | Connected  | Configured    |
| Fabric control |            |            |               |
| FC-0           |            | Connected  | Configured    |
| FC-1           |            | Connected  | Configured    |

#### show interfaces fabric device-name

```
user@switch> show interfaces fabric IC-WS001:fte-0/0/13
```

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up

Interface index: 49177, SNMP ifIndex: 1209008767

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU  
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:  
Disabled,

Flow control: Disabled

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

CoS queues : 12 supported, 12 maximum usable queues

Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d

Last flapped : 2012-11-13 23:55:15 UTC (00:55:38 ago)

Input rate : 0 bps (0 pps)

Output rate : 0 bps (0 pps)

Active alarms : None

Active defects : None

Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1209008768)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Input packets : 748

Output packets: 954

Protocol eth-switch, MTU: 0



## show interfaces ge

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces <i>device-name:type-fpc/pic/port</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;descriptions&gt;</code><br><code>&lt;media&gt;</code><br><code>&lt;routing-instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;snmp-index <i>snmp-index</i>&gt;</code><br><code>&lt;statistics&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b><i>device-name:type-fpc/pic/port</i></b>—The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing-instance or display all routing-instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li><li>• <a href="#">Troubleshooting Network Interfaces on page 1160</a></li><li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li><li>• <a href="#">Junos® OS Network Interfaces</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show interfaces on page 2680</a><br><a href="#">show interfaces brief on page 2680</a><br><a href="#">show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled) on page 2680</a><br><a href="#">show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled) on page 2681</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

[show interfaces extensive \(Symmetric Flow Control and Autonegotiation Enabled\) on page 2682](#)

[show interfaces extensive \(Asymmetric Flow Control and Autonegotiation Enabled\) on page 2684](#)

[show interfaces terse on page 2686](#)

[show interfaces terse \(QFabric Systems\) on page 2686](#)

**Output Fields** Table 249 on page 2673 lists the output fields for the **show interfaces ge** command. Output fields are listed in the approximate order in which they appear.

**Table 249: show interfaces ge Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output               |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Physical Interface</b>      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                               |
| <b>Physical interface</b>      | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                    |
| <b>Enabled</b>                 | State of the interface: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                    |
| <b>Interface index</b>         | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b>  |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive none</b>  |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>       |
| <b>Description</b>             | Optional user-specified description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive</b> |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                    |
| <b>MTU</b>                     | Maximum transmission unit size on the physical interface. The default is 1514.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                    |
| <b>Speed</b>                   | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                    |
| <b>Loopback</b>                | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                    |
| <b>Source filtering</b>        | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels                    |
| <b>Flow control</b>            | Flow control status: <b>Enabled</b> or <b>Disabled</b> .<br><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is not configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>       |
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ):<br><br><ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/> <b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/> <b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is configured. | <b>detail extensive</b>       |

Table 249: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                            | All levels                   |
| <b>Device flags</b>            | Information about the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                   |
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                   |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b> |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled on the switch.</p> | <b>detail extensive</b>      |

Table 249: show interfaces ge Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output  |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 249: show interfaces ge Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output              |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Output errors</b>                    | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>             |
| <b>Egress queues</b>                    | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Queue counters (Egress )</b>         | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Queue Number</b>                     | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |



Table 249: show interfaces ge Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| MAC statistics    | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | extensive       |
| Filter Statistics | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | extensive       |

Table 249: show interfaces ge Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output       |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Autonegotiation information            | <p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> <li>• <b>Link partner speed</b>—Speed of the link partner.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive             |
| Packet Forwarding Engine configuration | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | extensive             |
| Logical Interface                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                       |
| Logical interface                      | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels            |
| Index                                  | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail extensive none |
| SNMP ifIndex                           | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail extensive none |
| Generation                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail extensive      |
| Flags                                  | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels            |

Table 249: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>Protocol</b>                | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b> | If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.                                                                                                                                                                                                                                                                                           | <b>extensive</b>             |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>             |
| <b>Transit statistics</b>      | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>             |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |
| <b>Route Table</b>             | Route table in which the logical interface address is located. For example, 0 refers to the routing table <b>inet.0</b> .                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Input Filters</b>           | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Output Filters</b>          | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about protocol family flags.<br><br>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled. | <b>detail extensive</b>      |
| <b><i>protocol-family</i></b>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                              | <b>brief</b>                 |
| <b>Flags</b>                   | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Destination</b>             | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Local</b>                   | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Broadcast</b>               | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```
user@switch> show interfaces ge-0/0/9
Physical interface: ge-0/0/9, Enabled, Physical link is Down
 Interface index: 129, SNMP ifIndex: 21
 Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
 Last flapped : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
 Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 22)
 Flags: SNMP-Traps
 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch
 Flags: None
```

### show interfaces brief

```
user@switch> show interfaces ge-0/0/9 brief
Physical interface: ge-0/0/9, Enabled, Physical link is Down
 Description: voice priority and tcp and icmp traffic rate-limiting filter at i
 ngress port
 Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None

Logical interface ge-0/0/9.0
 Flags: Device-Down SNMP-Traps Encapsulation: ENET2
 eth-switch
```

### show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled)

```
user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
 Interface index: 193, SNMP ifIndex: 206, Generation: 196
 Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
```

```

Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 assured-forw 0 0 0
 5 expedited-fo 0 0 0
 7 network-cont 0 0 0

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,

```

#### show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on ,
Auto-negotiation: Enabled,

```

```

Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 assured-forw 0 0 0
5 expedited-fo 0 0 0
7 network-cont 0 0 0

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

#### show interfaces extensive (Symmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
Interface index: 49164, SNMP ifIndex: 574, Generation: 142

```

```

Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
Last flapped : 2011-02-25 00:45:03 UTC (22:42:48 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 2 no-loss      | 0              | 0                   | 0               |
| 3 fcoe         | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |

```

Queue number: Mapped forwarding classes
0 best-effort
2 no-loss
3 fcoe
7 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |

```

VLAN tagged frames 0
Code violations 0
MAC Priority Flow Control Statistics:
 Priority : 0 0 0
 Priority : 1 0 0
 Priority : 2 0 0
 Priority : 3 0 0
 Priority : 4 0 0
 Priority : 5 0 0
 Priority : 6 0 0
 Priority : 7 0 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Incomplete
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 0 best-effort 75 750000000 75 0 low
none
 7 network-control 5 500000000 5 0 low
none
 8 mcast-be 15 1500000000 15 0 low
none
 11 mcast-nc 5 500000000 5 0 low
none

```

#### show interfaces extensive (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
 Interface index: 49164, SNMP ifIndex: 574, Generation: 142
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on
 Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
 Last flapped : 2011-02-25 00:45:03 UTC (22:42:48 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes: 0 0 bps
 Input packets: 0 0 pps

```



```

Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 7 network-cont 0 0 0

Queue number: Mapped forwarding classes
 0 best-effort
 2 no-loss
 3 fcoe
 7 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
 Total octets Receive Transmit
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
MAC Priority Flow Control Statistics:
 Priority : 0 0 0
 Priority : 1 0 0
 Priority : 2 0 0
 Priority : 3 0 0
 Priority : 4 0 0
 Priority : 5 0 0
 Priority : 6 0 0
 Priority : 7 0 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0

```

```

Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link Partner:
 Link mode: Full-duplex, Flow control: None, Remote fault: OK,
 Link partner Speed: 1000 Mbps
Local resolution:
 Flow control: enable PAUSE transmit and Disable PAUSE receive, Remote
fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 75 750000000 75 0 low
none
7 network-control 5 50000000 5 0 low
none
8 mcast-be 15 150000000 15 0 low
none
11 mcast-nc 5 50000000 5 0 low
none

```

#### show interfaces terse

```

user@switch> show interfaces ge-0/0/12 terse
Interface Admin Link Proto Local Remote
ge-0/0/12 up up

```

#### show interfaces terse (QFabric Systems)

```

user@switch> show interfaces node1:ge-0/0/0 terse
Physical interface: node1:ge-0/0/0, Enabled, Physical link is Down
 Interface index: 129, SNMP ifIndex: 2884086
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
Interface flags: Internal: 0x4000
CoS queues : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

```

## show interfaces mc-ae

|                                 |                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show interfaces mc-ae id <i>identifier</i> unit <i>number</i></b>                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                      |
| <b>Description</b>              | On peers with multi-chassis aggregated Ethernet ( <b>mc-aeX</b> ) interfaces, use this command to display information about the <b>mc-aeX</b> interfaces.                            |
| <b>Options</b>                  | <p><b>identifier</b>—(Optional) Name of the multichassis aggregated Ethernet interface.</p> <p><b>number</b>—(Optional) Specify the logical interface by unit number.</p>            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show interfaces mc-ae on page 2688</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 250 on page 2687</a> lists the output fields for the <b>show interfaces mc-ae</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 250: show interfaces mc-ae Output Fields**

| Output Field Name                    | Field Description                                                                                                                                                                                |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Current State Machine's State</b> | Specifies the state of the MC-LAG initialization state machine.                                                                                                                                  |
| <b>Member Link</b>                   | Specifies the identifiers of the configured multichassis link aggregated interface members.                                                                                                      |
| <b>Local Status</b>                  | Specifies the status of the local link: <b>active</b> or <b>standby</b> .                                                                                                                        |
| <b>Peer Status</b>                   | Specifies the status of the peer link: <b>active</b> or <b>standby</b> .                                                                                                                         |
| <b>Peer State</b>                    | Specifies the status of the local and peer links in an <b>active/active</b> MC-LAG configuration                                                                                                 |
| <b>Logical Interface</b>             | Specifies the identifier and unit of the AE interface.                                                                                                                                           |
| <b>Topology Type</b>                 | Specifies the bridge configured on the AE.                                                                                                                                                       |
| <b>Local State</b>                   | Specifies if the local device is up or down.                                                                                                                                                     |
| <b>Peer State</b>                    | Specifies if the peer device is up or down.                                                                                                                                                      |
| <b>Peer Ip/MCP/State</b>             | Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet (MC-AE) interfaces that are part of the peer. |

## Sample Output

### show interfaces mc-ae

```
user@host> show interfaces mc-ae ae1 512
Member Link : ae0
Current State Machine's State: mcae active state
Local Status : active
Local State : up
Peer Status : active
Peer State : up
 Logical Interface : ae0.0
 Topology Type : bridge
 Local State : up
 Peer State : up
 Peer Ip/MCP/State : 3.3.3.2 ae1.0 up
```

## show interfaces statistics fabric

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces statistics fabric &lt;brief   detail   terse&gt; &lt;descriptions&gt; &lt;interface-name&gt; &lt;media&gt; &lt;routing-instance (all   instance-name)&gt; &lt;snmp-index snmp-index&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display status information about the specified fabric interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>brief   detail   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>interface-name</b>—(QFabric systems only) The interface name is either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   instance-name)</b>—(Optional) Display all routing instances or the name of an individual routing instance.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 1160</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces statistics fabric on page 2694</a></p> <p><a href="#">show interfaces statistics fabric brief on page 2701</a></p> <p><a href="#">show interfaces statistics fabric detail on page 2704</a></p> <p><a href="#">show interfaces statistics fabric terse on page 2705</a></p> <p><a href="#">show interfaces statistics fabric device-name on page 2706</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | <a href="#">Table 251 on page 2690</a> lists the output fields for the <b>show interfaces statistics fabric</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 251: show interfaces statistics fabric Output Fields

| Field Name                | Field Description                                                                                                                                                                                                                                    | Level of Output |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                                                                                                                                                                                      |                 |
| Physical interface        | Name of the physical interface.                                                                                                                                                                                                                      | All levels      |
| Enabled                   | State of the interface.                                                                                                                                                                                                                              | All levels      |
| Interface index           | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                  | detail none     |
| SNMP ifIndex              | SNMP index number for the physical interface.                                                                                                                                                                                                        | detail none     |
| Link-level type           | Encapsulation being used on the physical interface.                                                                                                                                                                                                  | All levels      |
| MTU                       | Maximum transmission unit size on the physical interface.                                                                                                                                                                                            | All levels      |
| Clocking                  | Reference clock source.                                                                                                                                                                                                                              | detail          |
| Speed                     | Speed at which the interface is running.                                                                                                                                                                                                             | All levels      |
| Duplex                    | Duplex mode of the interface, either Full-Duplex or Half-Duplex.                                                                                                                                                                                     | All levels      |
| MAC-REWRITE Error         | Specifies if the encapsulation of the packet has been changed.                                                                                                                                                                                       | none            |
| BPDU Error                | Specifies if a BPDU has been received on a blocked interface.                                                                                                                                                                                        | none            |
| Loopback                  | Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.                                                                                                                                                     | All levels      |
| Source filtering          | Source filtering status: Enabled or Disabled.                                                                                                                                                                                                        | All levels      |
| Flow control              | Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.                                                                                                                                 | All levels      |
| Device flags              | Information about the physical device.                                                                                                                                                                                                               | All levels      |
| Interface flags           | Information about the interface.                                                                                                                                                                                                                     | All levels      |
| CoS queues                | Number of CoS queues configured.                                                                                                                                                                                                                     | detail none     |
| Hold-Times                | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                  | detail          |
| Current address           | Configured MAC address.                                                                                                                                                                                                                              | detail none     |
| Hardware address          | Hardware MAC address.                                                                                                                                                                                                                                | detail none     |
| Last flapped              | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> . | detail none     |

Table 251: show interfaces statistics fabric Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Statistics last cleared | Date, time, and how long ago the statistics for the interface were cleared. The format is <b>Statistics last cleared: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail          |
| Traffic statistics      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail          |
| IPv6 transit statistics | <p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail          |
| Input errors            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• Errors—Sum of the incoming frame aborts and FCS errors.</li> <li>• Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• Framing errors—Number of packets received with an invalid frame checksum (FCS).</li> <li>• Runt—Number of frames received that are smaller than the runt threshold.</li> <li>• Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement.</li> <li>• L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• Resource errors—Sum of transmit drops.</li> </ul> | detail none     |

Table 251: show interfaces statistics fabric Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Output errors            | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>Errors—Sum of the outgoing frame aborts and FCS errors.</li> <li>Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces.</li> <li>MTU errors—Number of packets whose size exceeded the MTU of the interface.</li> <li>Resource errors—Sum of transmit drops.</li> </ul> | detail none     |
| Egress queues            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail          |
| Queue counters           | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>Queued packets—Number of queued packets.</li> <li>Transmitted packets—Number of transmitted packets.</li> <li>Dropped packets—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail          |
| Input rate               | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | None specified  |
| Output rate              | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | None specified  |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                 |
| Logical interface        | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| Index                    | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail none     |
| SNMP ifIndex             | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail none     |



Table 251: show interfaces statistics fabric Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Flags                  | <p>Information about the logical interface.</p> <p>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.</p>                                                                              | All levels      |
| Input packets          | Number of packets received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail none     |
| Output packets         | Number of packets transmitted on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail none     |
| Input packets          | Number of packets received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail none     |
| Output packets         | Number of packets transmitted on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail none     |
| Encapsulation          | <p>Encapsulation method used on the logical interface.</p> <ul style="list-style-type: none"> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                        | All levels      |
| Traffic statistics     | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |
| Local statistics       | <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                 | detail          |
| Transit statistics     | <ul style="list-style-type: none"> <li>Input bytes—Number of bytes received on the interface.</li> <li>Output bytes—Number of bytes transmitted on the interface.</li> <li>Input packets—Number of packets received on the interface.</li> <li>Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                 | detail          |
| Addresses, Flags       | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail none     |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                                                                                               | brief           |

Table 251: show interfaces statistics fabric Output Fields (*continued*)

| Field Name  | Field Description                                                                                                 | Level of Output |
|-------------|-------------------------------------------------------------------------------------------------------------------|-----------------|
| MTU         | Maximum transmission unit size on the physical interface.                                                         | All levels      |
| Destination | IP address of the remote side of the connection.                                                                  | detail none     |
| Local       | IP address of the logical interface.                                                                              | detail none     |
| Broadcast   | Broadcast address of the logical interlace.                                                                       | detail none     |
| Generation  | Unique number for use by Juniper Networks technical support only.                                                 | detail          |
| Route table | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0. | detail none     |

## Sample Output

### show interfaces statistics fabric

```

user@switch> show interfaces statistic fabric
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
 Interface index: 49174, SNMP ifIndex: 1208484473
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:19 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474)

 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
 Interface index: 49175, SNMP ifIndex: 1208484363
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:04, Hardware address: 00:00:00:00:00:04
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

```

Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/4.32768 (Index 84) (SNMP ifIndex 1208484364)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Input packets : 0

Output packets: 0

Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down

Interface index: 49176, SNMP ifIndex: 1208484367

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

Disabled, Flow control: Disabled

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

CoS queues : 12 supported, 12 maximum usable queues

Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06

Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)

Statistics last cleared: Never

Input rate : 0 bps (0 pps)

Output rate : 0 bps (0 pps)

Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1208484368)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Input packets : 0

Output packets: 0

Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down

Interface index: 49177, SNMP ifIndex: 1208484479

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

Disabled, Flow control: Disabled

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

CoS queues : 12 supported, 12 maximum usable queues

Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d

Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)

Statistics last cleared: Never

Input rate : 0 bps (0 pps)

Output rate : 0 bps (0 pps)

Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Input packets : 0

Output packets: 0

Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down

Interface index: 49178, SNMP ifIndex: 1208484475

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

Disabled, Flow control: Disabled

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

CoS queues : 12 supported, 12 maximum usable queues

Current address: 00:00:00:00:00:0f, Hardware address: 00:00:00:00:00:0f  
Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)  
Statistics last cleared: Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/15.32768 (Index 87) (SNMP ifIndex 1208484476)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
Input packets : 0  
Output packets: 0  
Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down  
Interface index: 49211, SNMP ifIndex: 1208484377  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU  
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:  
Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0  
CoS queues : 12 supported, 12 maximum usable queues  
Current address: 00:00:00:00:00:02, Hardware address: 00:00:00:00:00:02  
Last flapped : 2012-11-27 20:30:47 UTC (01:55:03 ago)  
Statistics last cleared: Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/2.32768 (Index 120) (SNMP ifIndex 1208484378)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
Input packets : 0  
Output packets: 0  
Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down  
Interface index: 49212, SNMP ifIndex: 1208484365  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU  
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:  
Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0  
CoS queues : 12 supported, 12 maximum usable queues  
Current address: 00:00:00:00:00:07, Hardware address: 00:00:00:00:00:07  
Last flapped : 2012-11-27 20:30:47 UTC (01:55:04 ago)  
Statistics last cleared: Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/7.32768 (Index 121) (SNMP ifIndex 1208484366)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
Input packets : 0  
Output packets: 0  
Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down  
Interface index: 49213, SNMP ifIndex: 1208484625  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU

Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled

Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x0  
 CoS queues : 12 supported, 12 maximum usable queues  
 Current address: 00:00:00:00:00:0a, Hardware address: 00:00:00:00:00:0a  
 Last flapped : 2012-11-27 20:30:47 UTC (01:55:04 ago)  
 Statistics last cleared: Never  
 Input rate : 0 bps (0 pps)  
 Output rate : 0 bps (0 pps)  
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/10.32768 (Index 122) (SNMP ifIndex 1208484626)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
 Input packets : 0  
 Output packets: 0  
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up

Interface index: 64, SNMP ifIndex: 1208483877  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1500  
 Device flags : Present Running  
 Current address: 02:00:00:00:40:06, Hardware address: 02:00:00:00:40:01  
 Last flapped : Never  
 Statistics last cleared: Never  
 Input packets : 0  
 Output packets: 26730  
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:bme0.0 (Index 4) (SNMP ifIndex 1208484065)

Flags: LinkAddress 0-0 Encapsulation: ENET2  
 Input packets : 2715  
 Output packets: 18  
 Protocol inet, MTU: 1482  
 Local: 128.0.32.0

Logical interface IC-WS001:bme0.1 (Index 5) (SNMP ifIndex 1208484091)

Flags: LinkAddress 0-0 Encapsulation: ENET2  
 Input packets : 0  
 Output packets: 999  
 Protocol inet, MTU: 1482  
 Addresses  
 Destination: 128/2, Local: 128.0.0.6, Broadcast: 191.255.255.255  
 Destination: 128/2, Local: 128.0.130.2, Broadcast: 191.255.255.255

Logical interface IC-WS001:bme0.2 (Index 6) (SNMP ifIndex 1208484092)

Flags: Encapsulation: ENET2  
 Input packets : 180408  
 Output packets: 23051  
 Protocol inet, MTU: 1486  
 Destination: 128/8, Local: 128.0.0.6, Broadcast: 128.255.255.255  
 Destination: 128/8, Local: 128.0.130.2, Broadcast: 128.255.255.255  
 Destination: 169.254/16, Local: 169.254.128.6, Broadcast: 169.254.255.255  
 Destination: 169.254/16, Local: 169.254.193.1, Broadcast: 169.254.255.255

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up

Interface index: 49156, SNMP ifIndex: 1208483949  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps  
 Device flags : Present Running

Interface flags: SNMP-Traps  
Link type : Full-Duplex  
Current address: 00:0d:0c:0f:00:03, Hardware address: 00:0d:0c:0f:00:03  
Last flapped : 1970-01-01 00:00:01 UTC (2238w5d 22:25 ago)  
Statistics last cleared: Never  
Input packets : 168885  
Output packets: 184712  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:bme1.0 (Index 3) (SNMP ifIndex 1208483950)  
Flags: Encapsulation: ENET2  
Input packets : 168885  
Output packets: 184712  
Protocol inet, MTU: 1500  
Destination: 128/2, Local: 128.0.0.1, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.5, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.16, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.17, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.24, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.25, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.26, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.28, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.29, Broadcast: 191.255.255.255  
Destination: 128/2, Local: 128.0.0.31, Broadcast: 191.255.255.255  
Protocol tnp, MTU: 1500  
Local: 0x5

Physical interface: IC-WS001:dcfabric, Enabled, Physical link is Up  
Interface index: 27, SNMP ifIndex: 1208484093  
Type: Ethernet, Link-level type: Ethernet, MTU: 1572  
Device flags : Present Running  
Interface flags: SNMP-Traps  
Current address: 00:0b:ca:fe:00:01, Hardware address: 00:0b:ca:fe:00:01  
Last flapped : Never  
Statistics last cleared: Never  
Input packets : 0  
Output packets: 0  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:dcfabric.0 (Index 64) (SNMP ifIndex 1208484094)  
Flags: SNMP-Traps Encapsulation: ENET2  
Input packets : 0  
Output packets: 0  
Protocol inet, MTU: 1558  
Protocol mpls, MTU: 1546, Maximum labels: 3  
Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up  
Interface index: 66, SNMP ifIndex: 1208484104  
Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps  
Device flags : Present Running  
Interface flags: SNMP-Traps  
Link type : Full-Duplex  
Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6  
Last flapped : Never  
Statistics last cleared: Never  
Input packets : 1007238  
Output packets: 63383  
Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up

```
Interface index: 67, SNMP ifIndex: 1208484105
Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7
Last flapped : Never
Statistics last cleared: Never
 Input packets : 1007118
 Output packets: 55381
Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down
Interface index: 68, SNMP ifIndex: 1208484106
Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8
Last flapped : 2012-11-27 02:52:03 UTC (19:33:54 ago)
Statistics last cleared: Never
 Input packets : 0
 Output packets: 0
Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down
Interface index: 69, SNMP ifIndex: 1208484107
Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9
Last flapped : 2012-11-27 02:52:03 UTC (19:33:54 ago)
Statistics last cleared: Never
 Input packets : 0
 Output packets: 0
Input errors: 0, Output errors: 0

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up
Interface index: 74, SNMP ifIndex: 1208484372
Type: Ethernet, Link-level type: 70, MTU: 1514, Speed: 1000mbps
Device flags : Present Running
Link type : Full-Duplex
Current address: 00:23:9c:f1:a2:e3, Hardware address: 00:23:9c:f1:a2:e3
Last flapped : Never
Statistics last cleared: Never
 Input packets : 121842
 Output packets: 3548
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp0.32769 (Index 11) (SNMP ifIndex 1208484376)
Flags: LinkAddress 0-0 Encapsulation: ENET2
Input packets : 13044
Output packets: 3548

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up
Interface index: 70, SNMP ifIndex: 1208484108
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
Device flags : Present Running
Link type : Full-Duplex
Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6
```

Last flapped : Never  
Statistics last cleared: Never  
Input packets : 767413  
Output packets: 46503  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp1.32768 (Index 7) (SNMP ifIndex 1208484109)  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
Input packets : 735889  
Output packets: 46503

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up  
Interface index: 71, SNMP ifIndex: 1208484369  
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps  
Device flags : Present Running  
Link type : Full-Duplex  
Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7  
Last flapped : Never  
Statistics last cleared: Never  
Input packets : 831710  
Output packets: 44548  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp2.32768 (Index 8) (SNMP ifIndex 1208484373)  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
Input packets : 737844  
Output packets: 44548

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down  
Interface index: 72, SNMP ifIndex: 1208484370  
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps  
Device flags : Present Running  
Link type : Full-Duplex  
Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8  
Last flapped : 2012-11-27 20:31:36 UTC (01:54:21 ago)  
Statistics last cleared: Never  
Input packets : 0  
Output packets: 0  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp3.32768 (Index 9) (SNMP ifIndex 1208484374)  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
Input packets : 0  
Output packets: 0

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down  
Interface index: 73, SNMP ifIndex: 1208484371  
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps  
Device flags : Present Running  
Link type : Full-Duplex  
Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9  
Last flapped : 2012-11-27 20:31:36 UTC (01:54:21 ago)  
Statistics last cleared: Never  
Input packets : 0  
Output packets: 0  
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp4.32768 (Index 10) (SNMP ifIndex 1208484375)  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
Input packets : 0  
Output packets: 0



**show interfaces statistics fabric brief**

```

user@switch> show interfaces statistics fabric brief
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/0.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/4.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/6.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/13.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/15.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/2.32768
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

```

### eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/7.32768  
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down  
Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,  
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled  
Device flags : Present Running  
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/10.32768  
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2  
eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,  
Speed: Unspecified  
Device flags : Present Running

Logical interface IC-WS001:bme0.0  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
inet 128.0.32.0 --> 0/0

Logical interface IC-WS001:bme0.1  
Flags: LinkAddress 0-0 Encapsulation: ENET2  
inet 128.0.0.6/2  
128.0.130.2/2

Logical interface IC-WS001:bme0.2  
Flags: Encapsulation: ENET2  
inet 128.0.0.6/8  
128.0.130.2/8  
169.254.128.6/16  
169.254.193.1/16

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up  
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,  
Speed: 1000mbps  
Device flags : Present Running  
Interface flags: SNMP-Traps

Logical interface IC-WS001:bme1.0  
Flags: Encapsulation: ENET2  
inet 128.0.0.1/2  
128.0.0.5/2  
128.0.0.16/2  
128.0.0.17/2  
128.0.0.24/2  
128.0.0.25/2  
128.0.0.26/2  
128.0.0.28/2  
128.0.0.29/2  
128.0.0.31/2

tnp 0x5

Physical interface: IC-WS001:dcfabric, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,  
 Speed: Unspecified  
 Device flags : Present Running  
 Interface flags: SNMP-Traps

Logical interface IC-WS001:dcfabric.0  
 Flags: SNMP-Traps Encapsulation: ENET2  
 inet  
 mpls  
 eth-switch

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,  
 Speed: 1000mbps  
 Device flags : Present Running  
 Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,  
 Speed: 1000mbps  
 Device flags : Present Running  
 Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,  
 Speed: 1000mbps  
 Device flags : Present Running  
 Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down  
 Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,  
 Speed: 1000mbps  
 Device flags : Present Running  
 Interface flags: SNMP-Traps

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:  
 1000mbps  
 Device flags : Present Running

Logical interface IC-WS001:vcp0.32769  
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:  
 1000mbps  
 Device flags : Present Running

Logical interface IC-WS001:vcp1.32768  
 Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up  
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:  
 1000mbps  
 Device flags : Present Running

Logical interface IC-WS001:vcp2.32768  
 Flags: LinkAddress 0-0 Encapsulation: ENET2

```

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface IC-WS001:vcp3.32768
 Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
 Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
 Device flags : Present Running

Logical interface IC-WS001:vcp4.32768
 Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

```

### show interfaces statistics fabric detail

```

user@switch> show interfaces statistics fabric detail
show interfaces statistics fabric detail
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
 Interface index: 49174, SNMP ifIndex: 1208484473, Generation: 153
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
 Last flapped : 2012-11-27 20:30:30 UTC (02:04:59 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 fabric_fcset 0 0 0
 1 fabric_fcset 0 0 0
 2 fabric_fcset 0 0 0
 3 fabric_fcset 0 0 0

```

```

4 fabric_fcset 0 0 0
5 fabric_fcset 0 0 0
6 fabric_fcset 0 0 0
7 fabric_fcset 0 0 0
8 fabric_fcset 0 0 0
9 fabric_fcset 0 0 0
10 fabric_fcset 0 0
0
11 fabric_fcset 0 0
0

```

Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474)  
(Generation 148)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Local statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Transit statistics:

```

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

```

Protocol eth-switch, MTU: 0, Generation: 176, Route table: 0

### show interfaces statistics fabric terse

```

user@switch> show interfaces statistics fabric terse
Interface Admin Link Proto Local Remote
IC-WS001:fte-0/0/0 up down
IC-WS001:fte-0/0/0.32768 up down eth-switch
IC-WS001:fte-0/0/4 up down
IC-WS001:fte-0/0/4.32768 up down eth-switch
IC-WS001:fte-0/0/6 up down
IC-WS001:fte-0/0/6.32768 up down eth-switch
IC-WS001:fte-0/0/13 up down
IC-WS001:fte-0/0/13.32768 up down eth-switch
IC-WS001:fte-0/0/15 up down
IC-WS001:fte-0/0/15.32768 up down eth-switch
IC-WS001:fte-1/0/2 up down
IC-WS001:fte-1/0/2.32768 up down eth-switch
IC-WS001:fte-1/0/7 up down
IC-WS001:fte-1/0/7.32768 up down eth-switch
IC-WS001:fte-1/0/10 up down
IC-WS001:fte-1/0/10.32768 up down eth-switch
IC-WS001:bme0 up up
IC-WS001:bme0.0 up up inet 128.0.32.0 --> 0/0
IC-WS001:bme0.1 up up inet 128.0.0.6/2
 128.0.130.2/2

```

```

IC-WS001:bme0.2 up up inet 128.0.0.6/8
 128.0.130.2/8
 169.254.128.6/16
 169.254.193.1/16

IC-WS001:bme1 up up
IC-WS001:bme1.0 up up inet 128.0.0.1/2
 128.0.0.5/2
 128.0.0.16/2
 128.0.0.17/2
 128.0.0.24/2
 128.0.0.25/2
 128.0.0.26/2
 128.0.0.28/2
 128.0.0.29/2
 128.0.0.31/2
 tnp 0x5

IC-WS001:dcfabric up up
IC-WS001:dcfabric.0 up up inet
 mpls
 eth-switch

IC-WS001:pme0 up up
IC-WS001:pme1 up up
IC-WS001:pme2 up down
IC-WS001:pme3 up down
IC-WS001:vcp0 up up
IC-WS001:vcp0.32769 up up
IC-WS001:vcp1 up up
IC-WS001:vcp1.32768 up up
IC-WS001:vcp2 up up
IC-WS001:vcp2.32768 up up
IC-WS001:vcp3 up down
IC-WS001:vcp3.32768 up down
IC-WS001:vcp4 up down
IC-WS001:vcp4.32768 up down

```

#### show interfaces statistics fabric device-name

```

user@switch> show interfaces statistics fabric IC-WS001:fte-0/0/13
Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
 Interface index: 49177, SNMP ifIndex: 1208484479
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
 Last flapped : 2012-11-27 20:30:30 UTC (02:09:53 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)

 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

```

## show interfaces queue

**Syntax** show interfaces queue  
 <aggregate | remaining-traffic>  
 <both-ingress-egress>  
 <egress>  
 <forwarding-class *forwarding-class*>  
 <ingress>  
 <interface-name *interface-name*>  
 <l2-statistics>  
 <remaining-traffic>

**Release Information** Command introduced before Junos OS Release 7.4.  
**both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
**l2-statistics** option introduced in Junos OS Release 12.1.

**Description** Display class-of-service (CoS) queue information for physical interfaces.

**Options** **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(optional) Display layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 252 on page 2708](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the layer 3 level. In the case of link fragmentation and interleaving (LFI) for which not fragmentation is applied, corresponding layer 2 overheads are added, as shown in [Table 252 on page 2708](#).

Table 252: Layer 2 Overhead, Transmitted Packets/Bytes

| Protocol       | Fragmentation       |                            | LFI |
|----------------|---------------------|----------------------------|-----|
|                | First fragmentation | Second to n fragmentations |     |
|                | Bytes               | Bytes                      |     |
| MLPPP (Long)   | 13                  | 12                         | 8   |
| MLPPP (short)  | 11                  | 10                         | 8   |
| MLFR (FRF15)   | 12                  | 10                         | 8   |
| MFR (FRF16)    | 10                  | 8                          | -   |
| MCMLPPP(Long)  | 13                  | 12                         | -   |
| MCMLPPP(Short) | 11                  | 10                         | -   |

## Layer 2 Statistics - Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header : 1 byte
HDLC flag and FCS bytes : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
Inner proto : 2 bytes
HDLC flag and FCS : 4 bytes

```

## Fragments 2 ...n :

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
HDLC flag and FCS : 4 bytes

```

## MFR (FRF16) Overhead details:

=====



```

Fragment 1:
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 Inner proto : 2 bytes
 HDLC flag and FCS : 4 bytes

Fragments 2 ...n :
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 HDLC flag and FCS : 4 bytes

```

## Overhead with LFI

```

MLPPP(Long & short sequence):
=====
 Outer PPP header : 4 bytes
 HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
 Framereelay header : 2 bytes
 Control,NLPID : 2 bytes
 HDLC flag and FCS : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

**Additional Information** On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular

physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos® OS Network Interfaces*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level** view

**List of Sample Output**

- [show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 2714](#)
- [show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 2716](#)
- [show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 2717](#)
- [show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 2717](#)
- [show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 2721](#)
- [show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 2724](#)
- [show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 2726](#)
- [show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 2727](#)
- [show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 2728](#)
- [show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 2731](#)
- [show interfaces queue \(QFX Series\) on page 2741](#)
- [show interfaces queue l2-statistics \(lsq interface\) on page 2742](#)

**Output Fields** [Table 253 on page 2710](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 253: show interfaces queue Output Fields**

| Field Name                          | Field Description                                                                                                                    |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>           | Name of the physical interface.                                                                                                      |
| <b>Enabled</b>                      | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . |
| <b>Interface index</b>              | Physical interface's index number, which reflects its initialization sequence.                                                       |
| <b>SNMP ifIndex</b>                 | SNMP index number for the interface.                                                                                                 |
| <b>Forwarding classes supported</b> | Total number of forwarding classes supported on the specified interface.                                                             |

Table 253: show interfaces queue Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding classes in use</b> | Total number of forwarding classes in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Ingress queues supported</b>  | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                    |
| <b>Ingress queues in use</b>     | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                       |
| <b>Output queues supported</b>   | Total number of output queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output queues in use</b>      | Total number of output queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Egress queues supported</b>   | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Egress queues in use</b>      | Total number of egress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Queue</b>                     | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Queue counters (Ingress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                           |
| <b>Burst size</b>                | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.                                                                                                                                                                                                                                                            |
| <b>Forwarding classes</b>        | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Queued Packets</b>            | <p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p>                               |
| <b>Queued Bytes</b>              | <p>Number of bytes queued to this queue. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a>.</p>                                                                                                                                                                                                                                                                                           |
| <b>Transmitted Packets</b>       | <p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2707</a></p> |

Table 253: show interfaces queue Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Transmitted Bytes    | <p>Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 2707</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tail-dropped packets | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RED-dropped packets  | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> |
| RED-dropped bytes    | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a>.</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> <li>• (J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority bytes dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority bytes dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul> </li> </ul>                           |

Byte counts vary by PIC type. [Table 254 on page 2713](#) shows how the byte counts on the outbound interfaces vary depending on the PIC type. [Table 254 on page 2713](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 254: Byte Count by PIC Type

| PIC Type                         | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Comments                                                                                                                                                                                                     |
|----------------------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Gigabit Ethernet IQ and IQE PICs | Interface                   | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
|                                  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | —                                                                                                                                                                                                            |
| Non-IQ PIC                       | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Switches:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.</li> </ul> | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>                                                                                                                  |

Table 254: Byte Count by PIC Type (*continued*)

| PIC Type                                             | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Comments                                                                                                                           |
|------------------------------------------------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| IQ and IQE PICs with a SONET/SDH interface           | Interface                   | Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes<br><br>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes<br><br>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes                                                                                                                                                                                                               | The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.                                                   |
|                                                      | Packet forwarding component | Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet<br><br>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes                                                                                                                                                                                                                                                                                                                      | For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.                      |
| Non-IQ PIC with a SONET/SDH interface                | Interface                   | T Series, TX Series, T1600, and MX Series routers:<br><ul style="list-style-type: none"><li>Queued: 478 bytes of Layer 3 packet.</li><li>Transmitted: 478 bytes of Layer 3 packet.</li></ul> M Series routers:<br><ul style="list-style-type: none"><li>Queued: 478 bytes of Layer 3 packet.</li><li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li><li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li></ul> | For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP). |
| Interfaces configured with Frame Relay Encapsulation | Interface                   | The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                    |
| 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs        | Interface                   | Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.<br><br>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.                                                                                                                                                                                                                                                                                                                | The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.                                               |
| 4-port 1G IQ2 and IQ2-E PICs                         | Packet forwarding component | Queued: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                     | —                                                                                                                                  |
| 8-port 1G IQ2 and IQ2-E PICs                         |                             | Transmitted: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                    |

## Sample Output

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```

user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 42603765 595484 pps
 Bytes : 5453281920 609776496 bps
 Transmitted:
 Packets : 42603765 595484 pps
 Bytes : 5453281920 609776496 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Transmitted:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: af11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: ef11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Transmitted:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
Queued:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Transmitted:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
Queue: 1, Forwarding classes: af1
Queued:
Packets : 2480391 1650 pps
Bytes : 1304685666 6945704 bps
Transmitted:
Packets : 2478740 1650 pps
Bytes : 1303817240 6945704 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 1651 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 1651 0 pps

```



```

RED-dropped bytes : 868426 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 868426 0 pps

```

### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Transmitted:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Transmitted:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Transmitted:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use

```

```

Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 148450735 947295 pps
 Bytes : 8016344944 409228848 bps
Transmitted:
 Packets : 76397439 487512 pps
 Bytes : 4125461868 210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets : 72053285 459783 pps
 Low : 72053285 459783 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 3890877444 198626472 bps
 Low : 3890877444 198626472 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 410278257 473940 pps
 Bytes : 22156199518 204742296 bps
Transmitted:
 Packets : 4850003 4033 pps
 Bytes : 261900162 1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets : 405425693 469907 pps
 Low : 405425693 469907 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 21892988124 203000040 bps
 Low : 21892988124 203000040 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 76605230 485376 pps
 Bytes : 5209211400 264044560 bps
 Transmitted:
 Packets : 76444631 484336 pps
 Bytes : 5198235612 263478800 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 160475 1040 pps
 Low : 160475 1040 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 10912300 565760 bps
 Low : 10912300 565760 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 4836136 3912 pps
 Bytes : 333402032 2139056 bps
 Transmitted:
 Packets : 3600866 1459 pps
 Bytes : 244858888 793696 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 1225034 2450 pps
 Low : 1225034 2450 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps

```

```

 High : 0 0 pps
 RED-dropped bytes : 83302312 1333072 bps
 Low : 83302312 1333072 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

## Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
 Packets : 77059796 486384 pps
 Bytes : 3544750624 178989576 bps
Transmitted:
 Packets : 77059797 486381 pps
 Bytes : 3544750670 178988248 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps

```

```

 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
Transmitted:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
 Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
 use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
Transmitted:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Transmitted:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

## Queue: 3, Forwarding classes: network-control

## Queued:

|         |   |         |          |
|---------|---|---------|----------|
| Packets | : | 77009   | 11 pps   |
| Bytes   | : | 6894286 | 7888 bps |

## Transmitted:

|         |   |         |          |
|---------|---|---------|----------|
| Packets | : | 77009   | 11 pps   |
| Bytes   | : | 6894286 | 7888 bps |

Tail-dropped packets : Not Available

|                    |   |   |       |
|--------------------|---|---|-------|
| RL-dropped packets | : | 0 | 0 pps |
|--------------------|---|---|-------|

|                  |   |   |       |
|------------------|---|---|-------|
| RL-dropped bytes | : | 0 | 0 bps |
|------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

## Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

## Queue: 0, Forwarding classes: best-effort

## Queued:

|         |   |        |       |
|---------|---|--------|-------|
| Packets | : | 1031   | 0 pps |
| Bytes   | : | 147328 | 0 bps |

## Transmitted:

|         |   |        |       |
|---------|---|--------|-------|
| Packets | : | 1031   | 0 pps |
| Bytes   | : | 147328 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 pps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 pps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 pps |
|-----------|---|---|-------|

RED-dropped bytes : 0 0 bps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 bps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 bps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 bps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 bps |
|-----------|---|---|-------|

## Queue: 1, Forwarding classes: expedited-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 pps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 pps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 pps |
|-----------|---|---|-------|

RED-dropped bytes : 0 0 bps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 bps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 bps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 bps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 bps |
|-----------|---|---|-------|

## Queue: 2, Forwarding classes: assured-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

```

 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
Transmitted:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps

```

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 254 0 pps
 Bytes : 16274 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```



```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps

```

```

Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
Transmitted:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 288 0 pps
 Bytes : 18450 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

```

Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 9538 0 pps
 Bytes : 3819840 0 bps
 Transmitted:
 Packets : 9538 0 pps
 Bytes : 3819840 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
 Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
Transmitted:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 109355853 471736 pps
 Bytes : 7436199152 256627968 bps
 Transmitted:
 Packets : 109355852 471736 pps
 Bytes : 7436198640 256627968 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps

```

```

RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps

Transmitted:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps

```

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

## Queue: 1, Forwarding classes: REALTIME

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 2, Forwarding classes: PRIVATE

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|



|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|         |   |      |       |
|---------|---|------|-------|
| Packets | : | 60   | 0 pps |
| Bytes   | : | 4560 | 0 bps |

Transmitted:

|                      |   |      |       |
|----------------------|---|------|-------|
| Packets              | : | 60   | 0 pps |
| Bytes                | : | 4560 | 0 bps |
| Tail-dropped packets | : | 0    | 0 pps |
| RED-dropped packets  | : | 0    | 0 pps |
| Low                  | : | 0    | 0 pps |
| Medium-low           | : | 0    | 0 pps |
| Medium-high          | : | 0    | 0 pps |
| High                 | : | 0    | 0 pps |
| RED-dropped bytes    | : | 0    | 0 bps |
| Low                  | : | 0    | 0 bps |
| Medium-low           | : | 0    | 0 bps |
| Medium-high          | : | 0    | 0 bps |
| High                 | : | 0    | 0 bps |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

|         |   |          |             |
|---------|---|----------|-------------|
| Packets | : | 371365   | 23620 pps   |
| Bytes   | : | 15597330 | 7936368 bps |

Transmitted:

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |
| High                 | : | 0        | 0 pps       |
| RED-dropped bytes    | : | 0        | 0 bps       |
| Low                  | : | 0        | 0 bps       |
| Medium-low           | : | 0        | 0 bps       |
| Medium-high          | : | 0        | 0 bps       |

|                                        |   |   |       |
|----------------------------------------|---|---|-------|
| High                                   | : | 0 | 0 bps |
| Queue: 1, Forwarding classes: REALTIME |   |   |       |
| Queued:                                |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Transmitted:                           |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Tail-dropped packets                   | : | 0 | 0 pps |
| RED-dropped packets                    | : | 0 | 0 pps |
| Low                                    | : | 0 | 0 pps |
| Medium-low                             | : | 0 | 0 pps |
| Medium-high                            | : | 0 | 0 pps |
| High                                   | : | 0 | 0 pps |
| RED-dropped bytes                      | : | 0 | 0 bps |
| Low                                    | : | 0 | 0 bps |
| Medium-low                             | : | 0 | 0 bps |
| Medium-high                            | : | 0 | 0 bps |
| High                                   | : | 0 | 0 bps |
| Queue: 2, Forwarding classes: PRIVATE  |   |   |       |
| Queued:                                |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Transmitted:                           |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Tail-dropped packets                   | : | 0 | 0 pps |
| RED-dropped packets                    | : | 0 | 0 pps |
| Low                                    | : | 0 | 0 pps |
| Medium-low                             | : | 0 | 0 pps |
| Medium-high                            | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|         |   |         |        |
|---------|---|---------|--------|
| Packets | : | 32843   | 0 pps  |
| Bytes   | : | 2641754 | 56 bps |

Transmitted:

|                      |   |         |        |
|----------------------|---|---------|--------|
| Packets              | : | 32843   | 0 pps  |
| Bytes                | : | 2641754 | 56 bps |
| Tail-dropped packets | : | 0       | 0 pps  |
| RED-dropped packets  | : | 0       | 0 pps  |
| Low                  | : | 0       | 0 pps  |
| Medium-low           | : | 0       | 0 pps  |
| Medium-high          | : | 0       | 0 pps  |
| High                 | : | 0       | 0 pps  |
| RED-dropped bytes    | : | 0       | 0 bps  |
| Low                  | : | 0       | 0 bps  |
| Medium-low           | : | 0       | 0 bps  |
| Medium-high          | : | 0       | 0 bps  |
| High                 | : | 0       | 0 bps  |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |



|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
 Queued:

```

|                                      |   |   |       |
|--------------------------------------|---|---|-------|
| Packets                              | : | 0 | 0 pps |
| Bytes                                | : | 0 | 0 bps |
| Transmitted:                         |   |   |       |
| Packets                              | : | 0 | 0 pps |
| Bytes                                | : | 0 | 0 bps |
| Tail-dropped packets : Not Available |   |   |       |
| Total-dropped packets:               |   | 0 | 0 pps |
| Total-dropped bytes :                |   | 0 | 0 bps |

### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 1 0 pps
 Bytes : 1001 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 1062 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: ef
 Queued:
 Packets : 1 0 pps
 Bytes : 1500 0 bps
 Transmitted:
 Packets : 6 0 pps
 Bytes : 1573 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: af
 Queued:
 Packets : 1 0 pps
 Bytes : 512 0 bps
 Transmitted:
 Packets : 3 0 pps
 Bytes : 549 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
=====

```

## show interfaces queue fabric

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show interfaces queue fabric<br><egress><br><forwarding-class <i>forwarding-class</i> ><br><interface-name <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display class-of-service (CoS) queue information for the fabric interfaces that are configured between Node devices and Interconnect devices.                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Show detailed CoS queue statistics for all physical interfaces.</p> <p><b>egress</b>—(Optional) Display egress queue statistics.</p> <p><b>forwarding-class <i>forwarding-class</i></b>—(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p><b>interface-name <i>interface-name</i></b>—(Optional) Show detailed CoS queue statistics for the specified interface.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show interfaces fabric on page 2651</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show interfaces queue fabric on page 2744</a><br><a href="#">show interfaces queue fabric egress on page 2753</a><br><a href="#">show interfaces queue fabric interface-name egress on page 2763</a><br><a href="#">show interfaces queue fabric interface-name egress forwarding-class forwarding-class-name on page 2764</a>                                                                                                                                           |
| <b>Output Fields</b>            | Table 253 on page 2710 lists the output fields for the <b>show interfaces queue fabric</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                          |

Table 255: show interfaces queue fabric Output Fields

| Field Name         | Field Description               |
|--------------------|---------------------------------|
| Physical interface | Name of the physical interface. |

Table 255: show interfaces queue fabric Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enabled</b>            | State of the interface. Possible values are: <ul style="list-style-type: none"> <li>Administratively down, Physical link is Down—The interface is turned off, and the physical link is inoperable.</li> <li>Administratively down, Physical link is Up—The interface is turned off, but the physical link is operational and can pass packets when it is enabled.</li> <li>Enabled, Physical link is Down—The interface is turned on, but the physical link is inoperable and cannot pass packets.</li> <li>Enabled, Physical link is Up—The interface is turned on, and the physical link is operational and can pass packets.</li> </ul> |
| <b>Interface index</b>    | Physical interface's index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Forwarding classes</b> | Number of forwarding classes supported and in use for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Egress queues</b>      | Number of output queues supported and in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Queue</b>              | CoS queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Transmitted</b>        | Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include: <ul style="list-style-type: none"> <li>Packets—Number of packets transmitted.</li> <li>Bytes—Number of bytes transmitted.</li> <li>Tail-dropped packets—Number of arriving packets dropped because output queue buffers were full.</li> <li>Total-dropped pkts—Number of transmitted packets dropped.</li> <li>Total dropped bytes—Number of transmitted bytes dropped.</li> </ul>                                                                                                                                        |
| <b>Queued</b>             | Number of packets and bytes queued to this queue. <ul style="list-style-type: none"> <li>Packets—Number of packets queued.</li> <li>Bytes—Number of bytes queued.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### show interfaces queue fabric

```

user@switch> show interfaces queue fabric
Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
 Interface index: 49178, SNMP ifIndex: 1208484475
 Forwarding classes: 16 supported, 5 in use
 Egress queues: 12 supported, 5 in use
 Queue: 0, Forwarding classes: fabric_fcset_be
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 62665971 0 pps

```

```

Bytes : 7770580404 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
Interface index: 49211, SNMP ifIndex: 1208484377
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
Queued:
 Packets : 0 0 pps

```

```

 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
Interface index: 49212, SNMP ifIndex: 1208484365
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use

```



```

Queue: 0, Forwarding classes: fabric_fcset_be
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up

```

Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available

```

```

Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps

```

```

Total-dropped bytes : 0 0 bps

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
Interface index: 49161, SNMP ifIndex: 1209008630
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

### show interfaces queue fabric egress

```
user@switch> show interfaces queue fabric egress
```

```

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
 Interface index: 49178, SNMP ifIndex: 1208484475
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 62665971 0 pps
 Bytes : 7770580404 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
Interface index: 49211, SNMP ifIndex: 1208484377
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```



```

Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:

```

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up

Interface index: 49212, SNMP ifIndex: 1208484365

Forwarding classes: 16 supported, 5 in use

Egress queues: 12 supported, 5 in use

Queue: 0, Forwarding classes: fabric\_fcset\_be

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Queue: 1, Forwarding classes: fabric\_fcset\_noloss1

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Queue: 2, Forwarding classes: fabric\_fcset\_noloss2

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Queue: 3, Forwarding classes: fabric\_fcset\_noloss3

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Queue: 4, Forwarding classes: fabric\_fcset\_noloss4

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | Not Available |       |
| Total-dropped pkts   | : | 0             | 0 pps |
| Total-dropped bytes  | : | 0             | 0 bps |

Queue: 5, Forwarding classes: fabric\_fcset\_noloss5

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:

```

```

Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available

```

```

Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps

```

```

 Total-dropped bytes : 0 0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up

Interface index: 49161, SNMP ifIndex: 1209008630

Forwarding classes: 16 supported, 5 in use

Egress queues: 12 supported, 5 in use

Queue: 0, Forwarding classes: best-effort

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Queue: 3, Forwarding classes: fcoe

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Queue: 4, Forwarding classes: no-loss

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Queue: 7, Forwarding classes: network-control

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps

```

Queue: 8, Forwarding classes: mcast

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:

```

```

Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps

```

#### show interfaces queue fabric interface-name egress

```

user@switch> show interfaces queue fabric BBAK0394:fte-0/1/0 egress
Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported,
5 in use Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 74777763341 844587 pps
 Bytes : 9272442654284 837830728 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped pkts : 0 0 pps

```

Total-dropped bytes : 0 0 bps

**show interfaces queue fabric interface-name egress forwarding-class forwarding-class-name**

```
user@switch> show interfaces queue fabric BBAK0394:fte-0/1/0 egress forwarding-class
best-effort
Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported,
5 in use Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 74793424543 844612 pps
Bytes : 9274384643332 837855936 bps
Tail-dropped packets : Not Available
Total-dropped pkts : 0 0 pps
Total-dropped bytes : 0 0 bps
```



## show interfaces xe

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>device-name:type-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   <i>instance-name</i>)&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>device-name:type-fpc/pic/port</i></b>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 1160</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces on page 2773</a></p> <p><a href="#">show interfaces (Asymmetric Flow Control) on page 2774</a></p> <p><a href="#">show interfaces brief on page 2774</a></p> <p><a href="#">show interfaces detail on page 2774</a></p> <p><a href="#">show interfaces detail (Asymmetric Flow Control) on page 2776</a></p> <p><a href="#">show interfaces extensive on page 2777</a></p> <p><a href="#">show interfaces extensive (Asymmetric Flow Control) on page 2779</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

[show interfaces terse on page 2781](#)

[show interfaces \(QFabric System\) on page 2781](#)

**Output Fields** Table 212 on page 2218 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

**Table 256: show interfaces xe Output Fields**

| Field Name                                                                              | Field Description                                                                                                                                                                             | Level of Output              |
|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b>                                                               |                                                                                                                                                                                               |                              |
| <b>Physical interface</b>                                                               | Name of the physical interface.                                                                                                                                                               | All levels                   |
| <b>Enabled</b>                                                                          | State of the interface.                                                                                                                                                                       | All levels                   |
| <b>Interface index</b>                                                                  | Index number of the physical interface, which reflects its initialization sequence.                                                                                                           | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                                                                     | SNMP index number for the physical interface.                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Generation</b>                                                                       | Unique number for use by Juniper Networks technical support only.                                                                                                                             | <b>detail extensive</b>      |
| <b>Link-level type</b>                                                                  | Encapsulation being used on the physical interface.                                                                                                                                           | All levels                   |
| <b>MTU</b>                                                                              | Maximum transmission unit size on the physical interface.                                                                                                                                     | All levels                   |
| <b>Speed</b>                                                                            | Speed at which the interface is running.                                                                                                                                                      | All levels                   |
| <b>Duplex</b>                                                                           | Duplex mode of the interface, either <b>Full-Duplex</b> or <b>Half-Duplex</b> .                                                                                                               | All levels                   |
| <b>Loopback</b>                                                                         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                | All levels                   |
| <b>Source filtering</b>                                                                 | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                  | All levels                   |
| <b>LAN-PHY mode</b>                                                                     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                   |
| <b>Unidirectional</b>                                                                   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.             | All levels                   |
| <b>Flow control</b>                                                                     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                      | All levels                   |
| <b>NOTE:</b> This field is only displayed if asymmetric flow control is not configured. |                                                                                                                                                                                               |                              |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output              |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ): <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/><b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/><b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <p><b>NOTE:</b> This field is only displayed if asymmetric flow control is configured.</p> | All levels                   |
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels                   |
| <b>Device flags</b>            | Information about the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels                   |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Wavelength</b>              | Configured wavelength, in nanometers (nm).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                   |
| <b>Frequency</b>               | Frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>Schedulers</b>              | Number of CoS schedulers configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | None specified               |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Input errors</b>       | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Output errors</b>            | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |
| <b>Egress queues</b>            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue Number</b>             | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Ingress queues</b>           | Total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| <b>Queue counters (Ingress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>        |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output              |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>PCS statistics</b>                   | Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>MAC statistics</b>                   | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b>             |
| <b>Filter statistics</b>                | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>             |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive). For asymmetric <b>PAUSE</b>, shows if the <b>PAUSE</b> transmit and <b>PAUSE</b> receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive       |

Table 256: show interfaces xe Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output              |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Packet Forwarding Engine configuration</b> | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li><b>Destination slot</b>—FPC slot number.</li> <li><b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li><b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li><b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li><b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li><b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li><b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li><b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                              |
| <b>Logical interface</b>                      | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels                   |
| <b>Index</b>                                  | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                           | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Flags</b>                                  | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels                   |
| <b>Encapsulation</b>                          | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Protocol</b>                               | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Traffic statistics</b>                     | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b>                | If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>             |
| <b>Local statistics</b>                       | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Transit statistics</b>                     | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>             |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Route Table</b>                            | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |



Table 256: show interfaces xe Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Input Filters</b>    | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Output Filters</b>   | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Flags</b>            | Information about protocol family flags.<br><br>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | <b>detail extensive</b>      |
| <b>Addresses, Flags</b> | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <i>protocol-family</i>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                           | <b>brief</b>                 |
| <b>Flags</b>            | Information about the address flag.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>Local</b>            | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Broadcast</b>        | Broadcast address of the logical interface.                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b> |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0

```

```
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode
```

### show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects: None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0
 Flags: Trunk-Mode
```

### show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None

Logical interface xe-0/0/1.0
 Flags: SNMP-Traps Encapsulation: ENET2
 eth-switch
```

### show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
```

```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6
 7 network-control
 8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps

```

```

Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Egress queues: 12 supported, 9 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

 Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6

```

```

7 network-control
8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

|                |   |   |   |
|----------------|---|---|---|
| 0 best-effort  | 0 | 0 | 0 |
| 1 fc7          | 0 | 0 | 0 |
| 2 no-loss      | 0 | 0 | 0 |
| 3 fcoe         | 0 | 0 | 0 |
| 4 fc4          | 0 | 0 | 0 |
| 5 fc5          | 0 | 0 | 0 |
| 6 fc6          | 0 | 0 | 0 |
| 7 network-cont | 0 | 0 | 0 |
| 8 mcast        | 0 | 0 | 0 |

Queue number:            Mapped forwarding classes

|   |                 |
|---|-----------------|
| 0 | best-effort     |
| 1 | fc7             |
| 2 | no-loss         |
| 3 | fcoe            |
| 4 | fc4             |
| 5 | fc5             |
| 6 | fc6             |
| 7 | network-control |
| 8 | mcast           |

Active alarms : None

Active defects : None

MAC statistics:

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

MAC Priority Flow Control Statistics:

|              |   |   |
|--------------|---|---|
| Priority : 0 | 0 | 0 |
| Priority : 1 | 0 | 0 |
| Priority : 2 | 0 | 0 |
| Priority : 3 | 0 | 0 |
| Priority : 4 | 0 | 0 |
| Priority : 5 | 0 | 0 |
| Priority : 6 | 0 | 0 |
| Priority : 7 | 0 | 0 |

Filter statistics:

|                      |   |   |
|----------------------|---|---|
| Input packet count   | 0 |   |
| Input packet rejects | 0 |   |
| Input DA rejects     | 0 |   |
| Input SA rejects     | 0 |   |
| Output packet count  |   | 0 |

```

Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low
none
7 network-control 5 500000000 5 0 low
none
8 mcast 20 2000000000 20 0 low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 fc7 0 0 0
2 no-loss 0 0 0
3 fcoe 0 0 0
4 fc4 0 0 0
5 fc5 0 0 0
6 fc6 0 0 0
7 network-cont 0 0 0
8 mcast 0 0 0

Queue number: Mapped forwarding classes
0 best-effort
1 fc7
2 no-loss
3 fcoe
4 fc4
5 fc5
6 fc6
7 network-control
8 mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets Receive Transmit
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
MAC Priority Flow Control Statistics:
Priority : 0 0 0
Priority : 1 0 0

```



```

Priority : 2 0 0
Priority : 3 0 0
Priority : 4 0 0
Priority : 5 0 0
Priority : 6 0 0
Priority : 7 0 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low none
7 network-control 5 500000000 5 0 low none
8 mcast 20 2000000000 20 0 low none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface Admin Link Proto Local Remote

xe-0/0/1 up up
xe-0/0/1.0 up up eth-switch

```

### show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000  
CoS queues : 8 supported, 8 maximum usable queues  
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00  
Last flapped : Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)

## show interfaces xle

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>device-name:type-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   <i>instance-name</i>)&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>device-name:type-fpc/pic/port</i></b>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 317</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 1160</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1161</a></li> <li>• <a href="#">Junos® OS Network Interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces on page 2791</a></p> <p><a href="#">show interfaces (Asymmetric Flow Control) on page 2792</a></p> <p><a href="#">show interfaces brief on page 2792</a></p> <p><a href="#">show interfaces detail on page 2792</a></p> <p><a href="#">show interfaces detail (Asymmetric Flow Control) on page 2794</a></p> <p><a href="#">show interfaces extensive on page 2795</a></p> <p><a href="#">show interfaces extensive (Asymmetric Flow Control) on page 2797</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

[show interfaces terse on page 2799](#)

[show interfaces \(QFabric System\) on page 2799](#)

**Output Fields** Table 212 on page 2218 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

**Table 257: show interfaces xe Output Fields**

| Field Name                                                                              | Field Description                                                                                                                                                                             | Level of Output              |
|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b>                                                               |                                                                                                                                                                                               |                              |
| <b>Physical interface</b>                                                               | Name of the physical interface.                                                                                                                                                               | All levels                   |
| <b>Enabled</b>                                                                          | State of the interface.                                                                                                                                                                       | All levels                   |
| <b>Interface index</b>                                                                  | Index number of the physical interface, which reflects its initialization sequence.                                                                                                           | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                                                                     | SNMP index number for the physical interface.                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Generation</b>                                                                       | Unique number for use by Juniper Networks technical support only.                                                                                                                             | <b>detail extensive</b>      |
| <b>Link-level type</b>                                                                  | Encapsulation being used on the physical interface.                                                                                                                                           | All levels                   |
| <b>MTU</b>                                                                              | Maximum transmission unit size on the physical interface.                                                                                                                                     | All levels                   |
| <b>Speed</b>                                                                            | Speed at which the interface is running.                                                                                                                                                      | All levels                   |
| <b>Duplex</b>                                                                           | Duplex mode of the interface, either <b>Full-Duplex</b> or <b>Half-Duplex</b> .                                                                                                               | All levels                   |
| <b>Loopback</b>                                                                         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                | All levels                   |
| <b>Source filtering</b>                                                                 | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                  | All levels                   |
| <b>LAN-PHY mode</b>                                                                     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                   |
| <b>Unidirectional</b>                                                                   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.             | All levels                   |
| <b>Flow control</b>                                                                     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                      | All levels                   |
| <b>NOTE:</b> This field is only displayed if asymmetric flow control is not configured. |                                                                                                                                                                                               |                              |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output              |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ): <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/><b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/><b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <p><b>NOTE:</b> This field is only displayed if asymmetric flow control is configured.</p> | All levels                   |
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels                   |
| <b>Device flags</b>            | Information about the physical device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels                   |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                   |
| <b>Wavelength</b>              | Configured wavelength, in nanometers (nm).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                   |
| <b>Frequency</b>               | Frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>Schedulers</b>              | Number of CoS schedulers configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | None specified               |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Input errors</b>       | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Output errors</b>            | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |
| <b>Egress queues</b>            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Queue Number</b>             | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Ingress queues</b>           | Total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| <b>Queue counters (Ingress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>        |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>PCS statistics</b>                   | Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>      |
| <b>MAC statistics</b>                   | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b>             |
| <b>Filter statistics</b>                | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>             |



Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive). For asymmetric <b>PAUSE</b>, shows if the <b>PAUSE</b> transmit and <b>PAUSE</b> receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive       |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output              |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Packet Forwarding Engine configuration</b> | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li><b>Destination slot</b>—FPC slot number.</li> <li><b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li><b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li><b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li><b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li><b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li><b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li><b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                              |
| <b>Logical interface</b>                      | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels                   |
| <b>Index</b>                                  | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                           | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Flags</b>                                  | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | All levels                   |
| <b>Encapsulation</b>                          | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Protocol</b>                               | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Traffic statistics</b>                     | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b>                | If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>             |
| <b>Local statistics</b>                       | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>             |
| <b>Transit statistics</b>                     | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>             |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Route Table</b>                            | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |

Table 257: show interfaces xe Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Input Filters</b>    | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Output Filters</b>   | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Flags</b>            | Information about protocol family flags.<br><br>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | <b>detail extensive</b>      |
| <b>Addresses, Flags</b> | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <i>protocol-family</i>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                           | <b>brief</b>                 |
| <b>Flags</b>            | Information about the address flag.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>Local</b>            | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Broadcast</b>        | Broadcast address of the logical interface.                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b> |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0

```

```
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode
```

### show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects: None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0
 Flags: Trunk-Mode
```

### show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None

Logical interface xe-0/0/1.0
 Flags: SNMP-Traps Encapsulation: ENET2
 eth-switch
```

### show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
```

```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6
 7 network-control
 8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps

```

```

Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
 Interface index: 49195, SNMP ifIndex: 591, Generation: 169
 Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
 Configured-flow-control tx-buffers: off rx-buffers: on
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 12 supported, 12 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
 Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes: 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes: 0
 Input packets: 0
 Output packets: 0
 Egress queues: 12 supported, 9 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets

 0 best-effort 0 0 0
 1 fc7 0 0 0
 2 no-loss 0 0 0
 3 fcoe 0 0 0
 4 fc4 0 0 0
 5 fc5 0 0 0
 6 fc6 0 0 0
 7 network-cont 0 0 0
 8 mcast 0 0 0

 Queue number: Mapped forwarding classes
 0 best-effort
 1 fc7
 2 no-loss
 3 fcoe
 4 fc4
 5 fc5
 6 fc6

```

```

7 network-control
8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

#### show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

|                |   |   |   |
|----------------|---|---|---|
| 0 best-effort  | 0 | 0 | 0 |
| 1 fc7          | 0 | 0 | 0 |
| 2 no-loss      | 0 | 0 | 0 |
| 3 fcoe         | 0 | 0 | 0 |
| 4 fc4          | 0 | 0 | 0 |
| 5 fc5          | 0 | 0 | 0 |
| 6 fc6          | 0 | 0 | 0 |
| 7 network-cont | 0 | 0 | 0 |
| 8 mcast        | 0 | 0 | 0 |

Queue number:            Mapped forwarding classes

|   |                 |
|---|-----------------|
| 0 | best-effort     |
| 1 | fc7             |
| 2 | no-loss         |
| 3 | fcoe            |
| 4 | fc4             |
| 5 | fc5             |
| 6 | fc6             |
| 7 | network-control |
| 8 | mcast           |

Active alarms : None

Active defects : None

MAC statistics:

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

MAC Priority Flow Control Statistics:

|              |   |   |
|--------------|---|---|
| Priority : 0 | 0 | 0 |
| Priority : 1 | 0 | 0 |
| Priority : 2 | 0 | 0 |
| Priority : 3 | 0 | 0 |
| Priority : 4 | 0 | 0 |
| Priority : 5 | 0 | 0 |
| Priority : 6 | 0 | 0 |
| Priority : 7 | 0 | 0 |

Filter statistics:

|                      |   |   |
|----------------------|---|---|
| Input packet count   | 0 |   |
| Input packet rejects | 0 |   |
| Input DA rejects     | 0 |   |
| Input SA rejects     | 0 |   |
| Output packet count  |   | 0 |



```

Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low
none
7 network-control 5 500000000 5 0 low
none
8 mcast 20 2000000000 20 0 low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 fc7 0 0 0
2 no-loss 0 0 0
3 fcoe 0 0 0
4 fc4 0 0 0
5 fc5 0 0 0
6 fc6 0 0 0
7 network-cont 0 0 0
8 mcast 0 0 0

Queue number: Mapped forwarding classes
0 best-effort
1 fc7
2 no-loss
3 fcoe
4 fc4
5 fc5
6 fc6
7 network-control
8 mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets Receive Transmit
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
MAC Priority Flow Control Statistics:
Priority : 0 0 0
Priority : 1 0 0

```

```

Priority : 2 0 0
Priority : 3 0 0
Priority : 4 0 0
Priority : 5 0 0
Priority : 6 0 0
Priority : 7 0 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 75 7500000000 75 0 low none
7 network-control 5 500000000 5 0 low none
8 mcast 20 2000000000 20 0 low none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface Admin Link Proto Local Remote

xe-0/0/1 up up
xe-0/0/1.0 up up eth-switch

```

### show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000  
CoS queues : 8 supported, 8 maximum usable queues  
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00  
Last flapped : Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)

## show lacp interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show lacp interfaces<br><interface-name>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p>none—Display LACP information for all interfaces.</p> <p><i>interface-name</i>—(Optional) Display LACP information for the specified interface:</p> <ul style="list-style-type: none"> <li>• Aggregated Ethernet—<i>aex</i></li> <li>• Gigabit Ethernet—<i>ge-fpc/pic/port</i></li> <li>• 10-Gigabit Ethernet—<i>xe-fpc/pic/port</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454</a></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure)</a></li> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure)</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2534</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2417</a></li> <li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp interfaces (EX Series Switches) on page 2803</a><br><a href="#">show lacp interfaces (QFX Series) on page 2804</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 246 on page 2643</a> lists the output fields for the <b>show lacp interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 258: show lacp interfaces Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated interface | Aggregated Ethernet interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| LACP State           | <p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> <li>For a child interface configured with the <b>force-up</b> statement, LACP state displays <b>FUP</b> along with the interface name.</li> <li><b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li><b>Actor</b>—Local device participating in the LACP negotiation.</li> <li><b>Partner</b>—Remote device participating in the LACP negotiation.</li> </ul> </li> <li><b>Exp</b>—Expired state. <b>Yes</b> indicates that the actor or partner is in an expired state. <b>No</b> indicates that the actor or partner is not in an expired state.</li> <li><b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. <b>No</b> indicates that the operational partner information in use has been received in an LACP PDU.</li> <li><b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li><b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li><b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. The link is currently not in the right aggregation.</li> <li><b>Aggr</b>—Ability of the aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li><b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li><b>Activity</b>—Actor's or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul> |

Table 258: show lacp interfaces Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of the state machine. The transmit state is one of the following values: <ul style="list-style-type: none"> <li><b>Fast periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—The multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting</b>—<b>Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing</b>—<b>Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—The multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul> |

## Sample Output

### show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
 LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity
 xe-2/0/7 Actor No No Yes Yes Yes Yes Fast Active
 xe-2/0/7 Partner No No Yes Yes Yes Yes Fast Passive

```

|          |         |    |    |    |     |     |     |      |         |
|----------|---------|----|----|----|-----|-----|-----|------|---------|
| xe-4/0/7 | Actor   | No | No | No | No  | No  | Yes | Fast | Active  |
| xe-4/0/7 | Partner | No | No | No | Yes | Yes | Yes | Fast | Passive |

|                    |               |                |                         |
|--------------------|---------------|----------------|-------------------------|
| LACP protocol:     | Receive State | Transmit State | Mux State               |
| xe-2/0/7(Active)   | Current       | Fast periodic  | Collecting distributing |
| xe-34/0/7(Standby) | Current       | Fast periodic  | Waiting                 |

### show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

| LACP state:       | Role    | Exp | Def | Dist | CoI | Syn | Aggr | Timeout | Activity |
|-------------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| node1:xe-0/0/1FUP | Actor   | No  | Yes | No   | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node1xe-0/0/1FUP  | Partner | No  | Yes | No   | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Actor   | No  | Yes | No   | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Partner | No  | Yes | No   | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |



|              | LACP protocol:           | Receive State | Transmit State | Mux State  |
|--------------|--------------------------|---------------|----------------|------------|
|              | node1:xe-0/0/1FUP        | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2           | Current       | Fast periodic  | Collecting |
| distributing | node1:xe-0/0/1 (active)  | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2 (standby) | Current       | Fast periodic  | WAITING    |

## show lacp statistics interfaces (View)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show lacp statistics interfaces</b> <i>interface-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command modified in Release 10.2 of Junos OS.<br>Command introduced in Release 11.1 of Junos OS for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>interface-name</i> —(Optional) Name of an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2537</a></li> <li>• <a href="#">Verifying the Status of a LAG Interface on page 2630</a></li> <li>• <a href="#">Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2631</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454</a></li> <li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2458</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp statistics interfaces on page 2807</a><br><a href="#">show lacp statistics interfaces (QFX Series) on page 2807</a><br><a href="#">show lacp statistics interfaces (QFabric Systems) on page 2807</a>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 259 on page 2806</a> lists the output fields for the <b>show lacp statistics interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 259: show lacp statistics interfaces Output Fields**

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Aggregated interface</b> | Aggregated interface value.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>LACP Statistics</b>      | <p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> |

## Sample Output

### show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
ge-2/0/0 1352 2035 0 0
ge-2/0/1 1352 2056 0 0
ge-2/2/0 1352 2045 0 0
ge-2/2/1 1352 2043 0 0
```

### show lacp statistics interfaces (QFX Series)

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
xe-0/0/2 1352 2035 0 0
xe-0/0/3 1352 2056 0 0
```

### show lacp statistics interfaces (QFabric Systems)

```
user@host> show lacp statistics interfaces nodegroup1:ae0
Aggregated interface: nodegroup1:ae0
LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
node1:xe-0/0/2 1352 2035 0 0
node2:xe-0/0/3 1352 2056 0 0
```

## show uplink-failure-detection

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show uplink-failure-detection</code><br><code>&lt;group group-name&gt;</code>                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                       |
| <b>Description</b>              | Display information about the uplink-failure-detection group, the member interfaces, and their status.                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display information about all groups configured for uplink failure detection.<br><b>group group-name</b> —(Optional) Display information about the specified group only.        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                         |
| <b>Related Documentation</b>    |                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show uplink-failure-detection on page 2808</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 260 on page 2808</a> lists the output fields for the <b>show uplink-failure-detection</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 260: show uplink-failure-detection Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                         |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group          | Name of the group.                                                                                                                                                                                                        |
| Uplink         | The uplink interface or interfaces configured as link-to-monitor.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                         |
| Downlink       | The downlink interface or interfaces configured as link-to-disable.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                       |
| Failure Action | Status of uplink failure detection: <ul style="list-style-type: none"> <li>Active—The switch has detected an uplink failure and has brought the downlink down.</li> <li>Inactive—The uplink or uplinks are up.</li> </ul> |

## Sample Output

### show uplink-failure-detection

```

user@switch> show uplink-failure-detection
Group : group1
Uplink : ge-0/0/0*
Downlink : ge-0/0/1*

```

Failure Action : Inactive

Group : group2

Uplink : ge-0/0/3.0

Downlink : ge-0/0/4.0

Failure Action : Active



## CHAPTER 31

# Troubleshooting

- [Troubleshooting Procedures on page 2811](#)

## Troubleshooting Procedures

---

- [Troubleshooting an Aggregated Ethernet Interface on page 2811](#)
- [Troubleshooting Multichassis Link Aggregation on page 2811](#)
- [Troubleshooting Network Interfaces on page 2817](#)

## Troubleshooting an Aggregated Ethernet Interface

**Problem**    **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution**    Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related Documentation**

- [Verifying the Status of a LAG Interface on page 2630](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)

## Troubleshooting Multichassis Link Aggregation

Use the following information to troubleshoot multichassis link aggregation configuration.

- [MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table on page 2812](#)
- [MC-LAG Peer Does Not Go into Standby Mode on page 2813](#)

- [Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive on page 2813](#)
- [Redirect Filters Take Priority over User-Defined Filters on page 2813](#)
- [Operational Command Output Is Wrong on page 2813](#)
- [ICCP Connection Might Take Up to 60 Seconds to Become Active on page 2814](#)
- [MAC Address Age Learned on an MC-AE Interface Is Reset to Zero on page 2814](#)
- [MAC Address Is Not Learned Remotely in a Default VLAN on page 2814](#)
- [Snooping Entries Learned on MC-AE Interfaces Are Not Removed on page 2814](#)
- [ICCP Does Not Come Up After You Add or Delete an Authentication Key on page 2815](#)
- [Local Status Is Standby When It Should Be Active on page 2815](#)
- [Packets Loop on the Server When ICCP Fails on page 2815](#)
- [Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change on page 2815](#)
- [No Commit Checks Are Done for ICL-PL Interfaces on page 2816](#)
- [Double Failover Scenario on page 2816](#)
- [Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up on page 2816](#)
- [Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer on page 2816](#)
- [AE Interfaces Go Down on page 2816](#)
- [Flooding of Upstream Traffic on page 2817](#)

### MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table

**Problem Description:** When both of the multichassis aggregated Ethernet (MC-AE) interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the MC-AE interfaces are not removed from the MAC address table. For example, if you disable the MC-AE interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the MC-AE interfaces of both MC-LAG peers:

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
VLAN MAC address Type Age Interfaces
v10 * Flood - All-members
v10 00:10:94:00:00:01 Learn(L) 3:55 ae0.0 (MCAE)
v10 00:10:94:00:00:02 Learn(R) 0 xe-0/0/9.0
v20 * Flood - All-members
v30 * Flood - All-members
v30 84:18:88:de:b1:2e Static - Router
```

```
user@switchB> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
VLAN MAC address Type Age Interfaces
v10 * Flood - All-members
```



|     |                   |          |                |
|-----|-------------------|----------|----------------|
| v10 | 00:10:94:00:00:01 | Learn(R) | 0 ae0.0 (MCAE) |
| v10 | 00:10:94:00:00:02 | Learn    | 40 xe-0/0/10.0 |
| v20 | *                 | Flood    | - All-members  |
| v30 | *                 | Flood    | - All-members  |
| v30 | 84:18:88:df:83:0a | Static   | - Router       |

**Solution** This is expected behavior.

### MC-LAG Peer Does Not Go into Standby Mode

**Problem** **Description:** A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Interchassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

**Solution** To prevent failure to enter standby mode, make sure the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

### Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet (MC-AE) interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's MC-AE interfaces with status control set to standby become inactive instead of active.

**Solution** This is expected behavior.

### Redirect Filters Take Priority over User-Defined Filters

**Problem** **Description:** Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters. This is expected behavior.

**Solution** This is expected behavior.

### Operational Command Output Is Wrong

**Problem** **Description:** After you deactivate the Interchassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
Redundancy Group IDs Joined: None
```

```
Client Application: lacpd
Redundancy Group IDs Joined: 1
```

```
Client Application: eswd
Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

**Solution** This is expected behavior.

---

#### ICCP Connection Might Take Up to 60 Seconds to Become Active

---

**Problem** **Description:** When the Interchassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

**Solution** This is expected behavior.

---

#### MAC Address Age Learned on an MC-AE Interface Is Reset to Zero

---

**Problem** **Description:** When you activate and then deactivate an interchassis control link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet (MC-AE) interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine (PFE). The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
 VLAN MAC address Type Age Interfaces
 --- -
v100 * Flood - All-members
v100 00:10:00:00:00:01 Learn(L) 0 ae0.0 (MCAE)
v100 00:10:00:00:00:02 Learn(L) 0 ae0.0 (MCAE)
```

**Solution** This is expected behavior.

---

#### MAC Address Is Not Learned Remotely in a Default VLAN

---

**Problem** **Description:** If a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, the Interchassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

**Solution** This is expected behavior.

---

#### Snooping Entries Learned on MC-AE Interfaces Are Not Removed

---

**Problem** **Description:** When multichassis aggregated Ethernet (MC-AE) interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the MC-AE interfaces on the VLAN are not cleared when the MC-AE interfaces go down.

This is done to speed up convergence time when the interfaces come up, or come up and go down.

**Solution** This is expected behavior.

---

#### ICCP Does Not Come Up After You Add or Delete an Authentication Key

---

**Problem** **Description:** The Interchassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

**Solution** Delete the ICCP configuration , and then add the ICCP configuration.

---

#### Local Status Is Standby When It Should Be Active

---

**Problem** **Description:** If the multichassis aggregated Ethernet (MC-AE) interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the MC-AE interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

**Solution** This is expected behavior.

---

#### Packets Loop on the Server When ICCP Fails

---

**Problem** **Description:** When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

**Solution** This is expected behavior.

---

#### Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

---

**Problem** **Description:** After a reboot or after a new Interchassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis aggregated Ethernet (MC-AE) interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

**Solution** This is expected behavior.

### No Commit Checks Are Done for ICL-PL Interfaces

---

**Problem** **Description:** There are no commit checks on the interface being configured as an interchassis control link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

**Solution** This is expected behavior.

### Double Failover Scenario

---

**Problem** **Description:** If the following events happen in this exact order—the Interchassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet (MC-AE) interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the MC-AE interface on the MC-LAG in active mode were up and blocks the interchassis control protocol-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

**Solution** This is expected behavior.

### Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

---

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) goes down and up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine (PFE) flag Ip4McastFloodMode for the VLAN is changed to MCAST\_FLOOD\_ALL. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

**Solution** This is expected behavior.

### Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

---

**Problem** **Description:** When the Interchassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

**Solution** This is expected behavior.

### AE Interfaces Go Down

---

**Problem** **Description:** When a multichassis aggregated Ethernet (MC-AE) interface is converted to an aggregated Ethernet (AE) interface, it retains some MC-AE properties. For example, the AE interface might retain the administrative key of the MC-AE. When this happens, the AE interface goes down.

**Solution** Restart the Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the AE interface to bring up the AE interface. Restarting LACP removes the MC-AE properties of the AE interface.

### Flooding of Upstream Traffic

**Problem** **Description:** When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

**Solution** Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)

## Troubleshooting Network Interfaces

### The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

**Problem** **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

**Cause** By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution** Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.



## PART 11

# Routing Options

- [Overview on page 2821](#)
- [Configuration on page 2825](#)
- [Administration on page 2955](#)
- [Troubleshooting on page 3139](#)





## CHAPTER 32

# Overview

- [Routing Options Overview on page 2821](#)

## Routing Options Overview

---

- [Overview of Routing Options on page 2821](#)
- [Understanding Virtual Router Routing Instances on page 2822](#)
- [Understanding Distributed Periodic Packet Management on page 2822](#)

## Overview of Routing Options

In addition to dynamic routing protocols, you can configure static routing on QFX Series switches. You can also configure a variety of protocol-independent routing properties, such as

- Per-packet load balancing (equal cost multipath routing)
- Autonomous system numbers
- Autonomous system confederation members
- Router identifiers
- Routing table groups
- Multicast scoping

## Understanding Virtual Router Routing Instances

Virtual router routing instances allow administrators to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

Each virtual router routing instance consists of sets of the following:

- Routing tables
- Interfaces that belong to these routing tables
- Routing protocol configurations
- Routing option configurations

You can use virtual router routing instances to isolate customer traffic on a network and to bind customer-specific routing instances to customer-owned interfaces. Each interface can belong to only one routing instance. QFX switches and QFabric systems support as many as 256 virtual router routing instances.

### Related Documentation

- [Configuring Virtual Router Routing Instances on page 2829](#)

## Understanding Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes on the QFX Series can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocol, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

---

**Related Documentation** • [Configuring Distributed Periodic Packet Management on page 2828](#)



## CHAPTER 33

# Configuration

- [Configuration Tasks on page 2825](#)
- [Configuration Examples on page 2830](#)
- [Configuration Statements on page 2854](#)

### Configuration Tasks

---

- [Configuring Static Routing on page 2826](#)
- [Configuring Per-Packet Load Balancing on page 2826](#)
- [Configuring Distributed Periodic Packet Management on page 2828](#)
- [Configuring Virtual Router Routing Instances on page 2829](#)

## Configuring Static Routing

Static routes are routes that are manually configured and entered into the routing table.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:  

```
[edit]
user@switch# set routing-options static route 0.0.0.0/0 next-hop 10.0.1.1
```
- To configure a static route and specify the next address to be used when routing traffic to the static route:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```
- To always keep the static route in the forwarding table:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```
- To prevent the static route from being readvertised:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```
- To remove inactive routes from the forwarding table:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

### Related Documentation

- [Monitoring Routing Information on page 2955](#)

## Configuring Per-Packet Load Balancing

When there are multiple equal-cost paths to the same destination for the active route, Junos OS chooses one of the next-hop addresses to install into the forwarding table in a random fashion by default. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use this feature to spread traffic across multiple paths.

When per-packet load balancing is configured, traffic is divided into individual flows (up to a maximum of 16). Packets for an individual flow are sent out a single interface. To determine flows, the switch examines each of the following packet fields:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The switch recognizes packets in which all of these parameters are identical and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.



**NOTE:** Load balancing is not supported on management interfaces.

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
 from {
 match-conditions;
 route-filter destination-prefix match-type <actions>;
 prefix-list name;
 }
 then {
 load-balance per-packet;
 }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {
 export policy-name;
}
```

When you enable per-packet load balancing on a QFabric system, packets might be switched across the fabric even though there is a local port with a same-cost route to the destination. For example, if per-packet load balancing is enabled and a packet arrives at network Node device A, it might be switched to network Node device B and forwarded from there even if there is a same-cost route through a port on Node device A to the destination. In this case, traffic transits the fabric needlessly. You can configure a QFabric

system to choose a locally switched route if one is available. To enable this feature, include the **ecmp-do-local-lookup** statement at the **[edit forwarding-options]** hierarchy level.

**Related  
Documentation**

- [Examples: Configuring Per-Packet Load Balancing on page 2830](#)

## Configuring Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes on the QFX Series can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for Link Aggregation Control Protocol (LACP) packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- [Disabling or Enabling Distributed Periodic Packet Management Globally on page 2828](#)
- [Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 2828](#)

### Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
user@switch# delete ppm no-delegate-processing
```

### Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.



To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

#### Related Documentation

- [Understanding Distributed Periodic Packet Management on page 2822](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2417](#)

## Configuring Virtual Router Routing Instances

Use virtual router routing instances to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

To configure virtual router routing instances:

1. Create a routing instance:

```
[edit routing-instances]user@switch# set routing-instance-name instance-type virtual-router
```



**NOTE:** The default routing instance, master, refers to the main inet.0 routing table. The master routing instance is reserved and cannot be specified as a routing instance.

2. Bind each routing instance to the corresponding interfaces:

```
[edit routing-instances]user@switch# set routing-instance-name interface
device-name:type-fpc/pic/port.logical-unit-number
```



**NOTE:**

- You must bind routing instances to interfaces from the Node devices assigned to the network Node group only. If you try to bind routing instances to interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
- You can bind an interface to one routing instance only.

3. Create each of the logical interfaces bound to each routing instance:

```
[edit interfaces]user@switch# set device-name:type-fpc/pic/port unit logical-unit-number
family inet address ip-address
```



**NOTE:** Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shut down.

4. (Optional) Configure routing protocols for the routing instance at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level.
5. (Optional) Configure routing options for the routing instance at the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level.

**Related Documentation**

- [Understanding Virtual Router Routing Instances on page 2822](#)
- [Understanding Interfaces on the QFabric System on page 1173](#)
- [Understanding Node Groups on page 1190](#)
- [Verifying That Virtual Router Routing Instances Are Working on page 2956](#)

---

## Configuration Examples

- [Examples: Configuring Per-Packet Load Balancing on page 2830](#)
- [Examples: Configuring BFD for Static Routes on page 2831](#)
- [Example: Configuring BFD Authentication for Static Routes on page 2846](#)

### Examples: Configuring Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
 policy-statement load-balancing-policy {
 then {
 load-balance per-packet;
 }
 }
}
routing-options {
 forwarding-table {
 export load-balancing-policy;
 }
}
```

Perform per-packet load balancing for a limited set of routes:

```
[edit]
policy-options {
 policy-statement load-balancing-policy {
 from {
 route-filter 192.168.10/24 orlonger;
 route-filter 9.114/16 orlonger;
 }
 then {
 load-balance per-packet;
 }
 }
}
routing-options {
 forwarding-table {
```

```

 export load-balancing-policy;
 }
}

```

#### Related Documentation

- [Configuring Per-Packet Load Balancing on page 2826](#)

## Examples: Configuring BFD for Static Routes

- [Understanding BFD for Static Routes on page 2831](#)
- [Example: Configuring BFD for Static Routes on page 2835](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes on page 2840](#)

### Understanding BFD for Static Routes

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes. In Junos OS Release 8.2 and later, BFD also supports multihop static routes.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is not supported for any other protocol.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route destination-prefix]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



**NOTE:** If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

---

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
  - For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
  - For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- 



**NOTE:** SRX Series devices do not support distributed BFD.

---

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval** **minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.



**NOTE:** You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



**NOTE:** We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.



**NOTE:** If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Junos OS also supports BFD over multihop static routes. For example, you can configure BFD over a Layer 3 path to provide path integrity over that path. You can limit the number of hops by specifying the time to live (TTL).

To configure BFD over multihop static routes, include the following statements:

```
static route destination-prefix {
 bfd-liveness-detection {
 local-address ip-address;
 minimum-receive-ttl number;
 }
}
```

To specify the source address for the multihop static route and to enable multihop BFD support, include the **local-address** statement.

To specify the number of hops, include the **minimum-receive-ttl** statement. You must configure this statement for a multihop BFD session. You can configure a value in the range from 1 through 255. It is optional for a single-hop BFD session. If you configure the **minimum-receive-ttl** statement for a single-hop session, the value must be 255.

---

### Example: Configuring BFD for Static Routes

---

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 2835](#)
- [Overview on page 2835](#)
- [Configuration on page 2836](#)
- [Verification on page 2839](#)

#### **Requirements**

In this example, no special configuration beyond device initialization is required.

#### **Overview**


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

[Figure 69 on page 2836](#) shows the sample network.

Figure 69: Customer Routes Connected to a Service Provider

**Configuration**

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device B**

```

set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Device D**

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.



- ```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```
2. On Device B, create a static route and set the next-hop address.


```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```
 3. On Device B, configure BFD for the static route.


```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
1000
```
 4. On Device B, configure tracing operations for BFD.


```
[edit protocols]
user@B# set bfd traceoptions file bfd-trace
user@B# set bfd traceoptions flag all
```
 5. If you are done configuring Device B, commit the configuration.


```
[edit]
user@B# commit
```
 6. On Device D, configure the interfaces.


```
[edit interfaces]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```
 7. On Device D, create a static route and set the next-hop address.


```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```
 8. On Device D, configure BFD for the static route.


```
[edit routing-options]
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
```
 9. On Device D, configure tracing operations for BFD.


```
[edit protocols]
user@D# set bfd traceoptions file bfd-trace
user@D# set bfd traceoptions flag all
```
 10. If you are done configuring Device D, commit the configuration.


```
[edit]
user@D# commit
```

Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B  user@B# show interfaces
          ge-1/2/0 {
            unit 0 {
              description B->D;
              family inet {
                address 172.16.1.1/24;
              }
            }
          }
          lo0 {
            unit 57 {
              family inet {
                address 10.0.0.1/32;
                address 10.0.0.2/32;
              }
            }
          }

          user@D# show protocols
          bfd {
            traceoptions {
              file bfd-trace;
              flag all;
            }
          }

          user@B# show routing-options
          static {
            route 192.168.47.0/24 {
              next-hop 172.16.1.2;
              bfd-liveness-detection {
                minimum-interval 1000;
              }
            }
          }

Device D  user@D# show interfaces
          ge-1/2/0 {
            unit 1 {
              description D->B;
              family inet {
                address 172.16.1.2/24;
              }
            }
          }
          lo0 {
            unit 2 {
              family inet {
                address 192.168.47.5/32;
                address 192.168.47.6/32;
              }
            }
          }

          user@D# show routing-options
          static {
```

```

route 0.0.0.0/0 {
  next-hop 172.16.1.1;
  bfd-liveness-detection {
    minimum-interval 1000;
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 2839](#)
- [Viewing Detailed BFD Events on page 2840](#)

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	lt-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 172
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

```
user@D> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.1	Up	lt-1/2/0.1	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:35
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 170
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 1, remote discriminator 2
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: ppm_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74
```

Meaning BFD messages are being written to the trace file.

Example: Enabling BFD on Qualified Next Hops in Static Routes

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

- [Requirements on page 2840](#)
- [Overview on page 2840](#)
- [Configuration on page 2841](#)
- [Verification on page 2844](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview


In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See [Figure 70 on page 2841](#).

Figure 70: BFD Enabled on Qualified Next Hops



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```

set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
  bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
  bfd-liveness-detection minimum-interval 60

```

Device D

```

set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60

```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@B# set unit 2 description secondary-B->D
user@B# set unit 2 family inet address 192.168.2.1/24
```

```
[edit interfaces ge-1/2/0]
user@B# set unit 0 description B->D
user@B# set unit 0 family inet address 172.16.1.1/24
```

2. On Device B, configure the static route with two next hops, both with BFD enabled.

```
[edit routing-options static route 192.168.47.0/24]
user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
60
user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
60
```

3. On Device D, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@D# set unit 3 description secondary-D->B
user@D# set unit 3 family inet address 192.168.2.2/24
```

```
[edit interfaces ge-1/2/0]
user@D# set unit 1 description D->B
user@D# set unit 1 family inet address 172.16.1.2/24
```

4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

In this case, BFD is enabled on the route, not on the next hops.

```
[edit routing-options static route 0.0.0.0/0]
user@D# set qualified-next-hop 192.168.2.1
user@D# set qualified-next-hop 172.16.1.1
user@D# set bfd-liveness-detection minimum-interval 60
```

Results Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/1/0 {
  unit 2 {
    description secondary-B->D;
    family inet {
      address 192.168.2.1/24;
    }
  }
}
```

```

ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    qualified-next-hop 192.168.2.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
    qualified-next-hop 172.16.1.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
  }
}

user@D# show interfaces
fe-0/1/0 {
  unit 3 {
    description secondary-D->B;
    family inet {
      address 192.168.2.2/24;
    }
  }
}
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}

user@D# show routing-options
static {
  route 0.0.0.0/0 {
    qualified-next-hop 192.168.2.1;
    qualified-next-hop 172.16.1.1;
    bfd-liveness-detection {
      minimum-interval 60;
    }
  }
}

```

If you are done configuring the devices, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 2844](#)
- [Verifying the BFD Sessions on page 2844](#)
- [Removing BFD from Device D on page 2844](#)
- [Removing BFD from One Next Hop on page 2845](#)

Checking the Routing Tables

Purpose Make sure that the static route appears in the routing table on Device B with two possible next hops.

Action `user@B> show route 192.168.47.0 extensive`
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}
 *Static Preference: 5
 Next hop type: Router
 Address: 0x9334010
 Next-hop reference count: 1
 Next hop: 172.16.1.2 via ge-1/2/0.0
 Next hop: 192.168.2.2 via fe-0/1/0.2, selected
 State: <Active Int Ext>
 Age: 9
 Task: RT
 Announcement bits (1): 3-KRT
 AS path: I

Meaning Both next hops are listed. The next hop 192.168.2.2 is the selected route.

Verifying the BFD Sessions

Purpose Make sure that the BFD sessions are up.

Action `user@B> show bfd session`

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	0.720	0.240	3
192.168.2.2	Up	fe-0/1/0.2	0.720	0.240	3

2 sessions, 2 clients
Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps

Meaning The output shows that the BFD sessions are up.

Removing BFD from Device D

Purpose Demonstrate what happens when the BFD session is down for both next hops.

- Action** 1. Deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Down	ge-1/2/0.0	3.000	1.000	3
192.168.2.2	Down	fe-0/1/0.2	3.000	1.000	3

```
2 sessions, 2 clients
```

```
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps
```

3. Rerun the **show route 192.168.47.0** command on Device B.

```
user@B> show route 192.168.47.0
```

Meaning As expected, when the BFD sessions are down, the static route is removed from the routing table.

Removing BFD from One Next Hop

Purpose Demonstrate what happens when only one next hop has BFD enabled.

- Action** 1. If it is not already deactivated, deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Deactivate BFD on one of the next hops on Device B.

```
[edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
user@B# deactivate bfd-liveness-detection
user@B# commit
```

3. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.2.2	Down	fe-0/1/0.2	3.000	1.000	3

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

```
user@B> show route 192.168.47.0 extensive
```

```
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```
192.168.47.0/24 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
```

```
*Static Preference: 5
```

```
Next hop type: Router, Next hop index: 624
```

```
Address: 0x92f0178
```

```
Next-hop reference count: 3
```

```
Next hop: 172.16.1.2 via ge-1/2/0.0, selected
State: <Active Int Ext>
Age: 2:36
Task: RT
Announcement bits (1): 3-KRT
AS path: I
```

Meaning As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

- Related Documentation**
- [Example: Configuring BFD Authentication for Static Routes on page 2846](#)
 - [Example: Configuring BFD for OSPF on page 3881](#)
 - [Example: Configuring BFD for BGP on page 3348](#)
 - [Example: Configuring BFD for IS-IS](#)
 - [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol](#)

Example: Configuring BFD Authentication for Static Routes

- [Understanding BFD Authentication for Static Routes on page 2846](#)
- [Example: Configuring BFD Authentication for Static Routes on page 2848](#)

Understanding BFD Authentication for Static Routes

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.



NOTE: We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 2847](#)
- [Security Authentication Keychains on page 2847](#)
- [Strict Versus Loose Authentication on page 2848](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and

associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for Static Routes

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

- [Requirements on page 2848](#)
- [Overview on page 2848](#)
- [Configuration on page 2849](#)
- [Verification on page 2852](#)

Requirements

Junos OS Release 9.6 or later (Canada and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Overview

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.
3. Configure the related security authentication keychain. This must be configured on the main router.




TIP: We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

[edit]

```
user@host> set routing-options static route ipv4 bfd-liveness-detection
authentication loose-check
```

Figure 71 on page 2849 shows the sample network.

Figure 71: Customer Routes Connected to a Service Provider



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
"$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
DkP5Ft0IQFcleV7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
"2011-1-12:00:00 -0800"
```

Device D

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain
  bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm
  keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
  "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
  DkP5FtOIQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
  "2011-1-1.12:00:00 -0800"

```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```

[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24

```

```

user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32

```

2. On Device B, create a static route and set the next-hop address.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2

```

3. On Device B, configure BFD for the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000

```

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1

```



NOTE: Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

- On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
```

- On Device B, specify the unique security authentication information for BFD sessions:
 - The matching keychain name as specified in Step 5.
 - At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
 - The secret data used to allow access to the session.
 - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-kc4]
user@B# set key 5 secret
"$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
DkP5Ft0IQFclev7N"
user@B# set key 5 start-time "2011-1-1.12:00:00 -0800"
```

- If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

- Repeat the configuration on Device D.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

Results

Confirm your configuration by issuing the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
```

```

        address 10.0.0.2/32;
    }
}

user@B# show routing-options
static {
    route 192.168.47.0/24 {
        next-hop 172.16.1.2;
        bfd-liveness-detection {
            minimum-interval 1000;
            authentication {
                key-chain bfd-kc4;
                algorithm keyed-sha-1;
            }
        }
    }
}

user@B# show security
authentication-key-chains {
    key-chain bfd-kc4 {
        key 5 {
            secret
            "$9$JhZHmn6Ap0ln/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
            DkP5Ft0IQFclev7N"; ## SECRET-DATA
            start-time "2011-1-1.12:00:00 -0800";
        }
    }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 2852](#)
- [Viewing Details About the BFD Session on page 2853](#)
- [Viewing Extensive BFD Session Information on page 2853](#)

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up.

Action From operational mode, enter the **show bfd session** command.

```
user@B> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning The command output shows that the BFD session is up.

Viewing Details About the BFD Session

Purpose View details about the BFD sessions and make sure that authentication is configured.

Action From operational mode, enter the **show bfd session detail** command.

```
user@B> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**
 Session up time 00:53:58
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22

1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive BFD Session Information

Purpose View more detailed information about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Time	Interval	Multiplier
172.16.1.2	Up	ge-1/2/0.0	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**
 keychain bfd-kc4, algo keyed-sha-1, mode strict
 Session up time 01:39:45
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 3, remote discriminator 4
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

Related Documentation

- [Examples: Configuring BFD for Static Routes on page 2831](#)

Configuration Statements

- [active](#) on page 2857
- [aggregate \(Routing\)](#) on page 2858
- [as-path \(Routing Options\)](#) on page 2860
- [autonomous-system](#) on page 2862
- [backup-pe-group](#) on page 2864
- [backups](#) on page 2865
- [bandwidth \(Multicast Flow Map\)](#) on page 2866
- [bfd-liveness-detection \(Routing Options Static Route\)](#) on page 2867
- [bgp-orf-cisco-mode](#) on page 2871
- [bmp](#) on page 2872
- [brief](#) on page 2873
- [centralized](#) on page 2874
- [community \(Routing Options\)](#) on page 2875
- [confederation](#) on page 2877
- [description \(Routing Instances\)](#) on page 2878
- [discard](#) on page 2879
- [export \(Routing Options\)](#) on page 2880
- [export-rib](#) on page 2881
- [fate-sharing](#) on page 2882
- [flow](#) on page 2883
- [flow-map](#) on page 2884
- [forwarding-cache \(Flow Maps\)](#) on page 2885
- [forwarding-cache \(Multicast\)](#) on page 2886
- [forwarding-table](#) on page 2887
- [generate](#) on page 2888
- [import \(Routing Options\)](#) on page 2889
- [import-policy](#) on page 2890
- [import-rib](#) on page 2891
- [indirect-next-hop](#) on page 2892
- [install \(Routing Options\)](#) on page 2893
- [instance-type](#) on page 2894
- [interface \(Multicast Static Routes\)](#) on page 2895
- [interface \(Routing Instances\)](#) on page 2896
- [interface \(Routing Options\)](#) on page 2897
- [interface-routes](#) on page 2898

- [local-address \(Routing Options\) on page 2899](#)
- [martians on page 2900](#)
- [maximum-bandwidth \(Routing Options\) on page 2901](#)
- [maximum-paths on page 2902](#)
- [maximum-prefixes on page 2904](#)
- [med-igp-update-interval on page 2905](#)
- [metric \(Aggregate, Generated, or Static Route\) on page 2906](#)
- [multicast \(Routing Options\) on page 2907](#)
- [no-qos-adjust on page 2908](#)
- [options \(Routing Options\) on page 2909](#)
- [pim-to-igmp-proxy on page 2910](#)
- [policy \(Aggregate and Generated Routes\) on page 2911](#)
- [policy \(Flow Maps\) on page 2912](#)
- [policy-options on page 2913](#)
- [policy-statement on page 2914](#)
- [ppm on page 2916](#)
- [ppm \(Ethernet Switching\) on page 2917](#)
- [preference \(Routing Options\) on page 2918](#)
- [prefix on page 2919](#)
- [protocols on page 2920](#)
- [qualified-next-hop \(Static Routes\) on page 2922](#)
- [readvertise on page 2924](#)
- [redundant-sources on page 2925](#)
- [resolution on page 2926](#)
- [resolution-ribs on page 2927](#)
- [resolve on page 2928](#)
- [retain on page 2929](#)
- [reverse-oif-mapping on page 2930](#)
- [rpf-check-policy \(Routing Options RPF\) on page 2931](#)
- [rib \(General\) on page 2932](#)
- [rib \(Route Resolution\) on page 2934](#)
- [rib-group \(Routing Options\) on page 2935](#)
- [rib-groups on page 2936](#)
- [route-record on page 2937](#)
- [router-id on page 2938](#)
- [routing-instances on page 2939](#)
- [routing-options on page 2939](#)

- [scope](#) on page 2940
- [scope-policy](#) on page 2941
- [source-routing](#) on page 2942
- [static \(Routes\)](#) on page 2943
- [subscriber-leave-timer](#) on page 2945
- [tag \(Routing Options\)](#) on page 2946
- [threshold \(Multicast Forwarding Cache\)](#) on page 2947
- [timeout \(Flow Maps\)](#) on page 2948
- [timeout \(Multicast\)](#) on page 2949
- [traceoptions \(Routing Options\)](#) on page 2950
- [upstream-interface](#) on page 2953

active

Syntax	(active passive);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with reject next hops when they are inactive.</p> <ul style="list-style-type: none"> • active—Remove a route from the routing and forwarding tables when it becomes inactive. • passive—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive. <p>Include the active statement when configuring an individual route in the route portion of the static statement to override a passive option specified in the defaults portion of the statement.</p>
Default	active
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i>

aggregate (Routing)

Syntax	<pre> aggregate { defaults { ... aggregate-options ... } route destination-prefix { policy policy-name; ... aggregate-options ... } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure aggregate routes.
Options	<p>aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in aggregate-options. Each option is explained separately.</p> <ul style="list-style-type: none"> • (active passive); • as-path <as-path> <origin (egp igp incomplete)> <atomic-aggregate> <aggregator as-number ip-address>; • (brief full); • community [<i>community-ids</i>]; • discard; • (metric metric2 metric3 metric4) <i>value</i> <type type>; • (preference preference2 color color2) <i>preference</i> <type type>; • tag string; <p>defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults</p>

and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Summarizing Routes Through Route Aggregation</i>

as-path (Routing Options)

Syntax	<code>as-path <as-path> <aggregator as-number ip-address> <atomic-aggregate> <origin (egp igp incomplete)>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.</p> <p>In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.</p>
Default	No AS path information is associated with static routes.
Options	<p>aggregator—(Optional) Attach the BGP aggregator path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for as-number, followed by the IP address of the BGP system that formed the aggregate route for ip-address.</p>

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for **as-path**. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • <i>Understanding a 4-Byte Capable Router AS Path Through a 2-Byte Capable Domain in the Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>
------------------------------	--

autonomous-system

Syntax	<code>autonomous-system <i>autonomous-system</i> <asdot-notation> <loops <i>number</i>> { independent-domain <no-attrset>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. asdot-notation option introduced in Junos OS Release 9.3. asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches. no-attrset option introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the routing device's AS number.

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (<http://www.isi.edu>).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options ***autonomous-system***—AS number. Use a number assigned to you by the NIC.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1



NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring External BGP Peering on page 3149• Examples: Configuring Internal BGP Peering on page 3172• <i>4-Byte Autonomous System Numbers Overview</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• <i>Juniper Networks Implementation of 4-Byte Autonomous System Numbers</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• <i>Configuring 4-Byte Autonomous System Numbers</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>

backup-pe-group

Syntax	<pre>backup-pe-group group-name { backups [addresses]; local-address address; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<p>backups <i>addresses</i>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>local-address <i>address</i>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>pe-group-name—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Ingress PE Redundancy</i>• <i>Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs</i>

backups

Syntax	<code>backups [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-options multicast backup-pe-group <i>group-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<i>addresses</i> —Addresses of other PEs in the backup group.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Ingress PE Redundancy</i>

bandwidth (Multicast Flow Map)

Syntax	<code>bandwidth (<i>bps</i> <i>adaptive</i>);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map], [edit routing-options multicast flow-map]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the bandwidth property for multicast flow maps.
Options	adaptive —Specify that the bandwidth is measured for the flows that are matched by the flow map. bps —Bandwidth, in bits per second, for the flow map. Range: 0 through any amount of bandwidth Default: 2 Mbps
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring a Multicast Flow Map</i>

bfd-liveness-detection (Routing Options Static Route)

Syntax

```

bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    holddown-interval milliseconds;
    local-address ip-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl number;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix qualified-next-hop (interface-name |
 address)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-options rib routing-table-name static route destination-prefix],
[edit routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit routing-options static route destination-prefix],

```

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

- Release Information** Statement introduced before Junos OS Release 7.4.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2.
local-address statement introduced in Junos OS Release 8.2.
minimum-receive-ttl statement introduced in Junos OS Release 8.2.
Support for logical routers introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5.
no-adaptation statement introduced in Junos OS Release 9.0.
Support for IPv6 static routes introduced in Junos OS Release 9.1.
authentication algorithm, **authentication key-chain**, and **authentication loose-check** statements introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
- Description** Configure bidirectional failure detection timers and authentication criteria for static routes.

- Options** **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.
- authentication key-chain** *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.
- authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.
- detection-time threshold** *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
- holddown-interval** *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.
Range: 0 through 255,000
Default: 0
- local-address** *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.
- minimum-interval** *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.
Range: 1 through 255,000
- minimum-receive-interval** *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.
Range: 1 through 255,000
- minimum-receive-ttl** *number*—Configure the time to live (TTL) for the multihop BFD session.
Range: 1 through 255
Default: 255
- multiplier** *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.
Range: 1 through 255
Default: 3

neighbor *address*—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000


version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• Example: Configuring BFD Authentication for Static Routes on page 2848
------------------------------	--

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<p> NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>
Default	Disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3313](#)


bmp

Syntax	<pre>bmp { memory limit <i>bytes</i>; station-address (<i>ip-address</i> <i>name</i>); station-port <i>port-number</i>; statistics-timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit routing-options]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.
Options	<p>memory-limit <i>bytes</i>—(Optional) Specify a threshold at which to stop collecting BMP data if the limit is exceeded.</p> <p>Default: 10 MB</p> <p>Range: 1,048,576 through 52,428,800</p> <p>station-address (<i>ip-address</i> <i>name</i>)—Specify the IP address or a valid URL for the monitoring where BMP data should be sent.</p> <p>station-port <i>port-number</i>—Specify the port number of the monitoring station to use when sending BMP data.</p> <p>statistics-timeout <i>seconds</i>—(Optional) Specify how often to send BMP data to the monitoring station.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the BGP Monitoring Protocol

brief

Syntax	(brief full);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.</p> <ul style="list-style-type: none"> • brief—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP ATOMIC_ATTRIBUTE path attribute is included with the aggregate route. • full—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the route portion of the generate statement to override a retain option specified in the defaults portion of the statement.
Default	full
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Summarizing Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • aggregate on page 2858 • generate on page 2888

centralized

Syntax	centralized;
Hierarchy Level	[edit protocols lacp ppm]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.</p> <p>This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing statement in the [edit routing-options ppm] hierarchy.</p>
	<div>BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.</div>
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i>• <i>Configuring Aggregated Ethernet LACP (CLI Procedure)</i>• Configuring Distributed Periodic Packet Management on page 2828• Configuring Link Aggregation on page 2537

community (Routing Options)

Syntax	<code>community ([<i>community-ids</i>] no-advertise no-export no-export-subconfed none);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Associate BGP community information with a static, aggregate, or generated route.
Default	No BGP community information is associated with static routes.
Options	<p><i>community-ids</i>—One or more community identifiers. The <i>community-ids</i> format varies according to the type of attribute that you use.</p> <p>The BGP community attribute format is <i>as-number:community-value</i>:</p> <ul style="list-style-type: none"> • <i>as-number</i>—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value. • <i>community-value</i>—Identifier of the community member. It can be a number from 0 through 65,535. <p>For more information about BGP community attributes, see the “Configuring the Extended Communities Attribute” section in the <i>Routing Policy Configuration Guide</i>.</p> <p>For specifying the BGP community attribute only, you also can specify <i>community-ids</i> as one of the following well-known community names defined in RFC 1997:</p> <ul style="list-style-type: none"> • no-advertise—Routes containing this community name are not advertised to other BGP peers. • no-export—Routes containing this community name are not advertised outside a BGP confederation boundary. • no-export-subconfed—Routes containing this community name are not advertised to external BGP peers, including peers in other members’ ASs inside a BGP confederation.



NOTE: Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the *Routing Policy Configuration Guide*.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• <i>Examples: Configuring Static Routes</i>• <i>Example: Summarizing Routes Through Route Aggregation</i>• <i>Example: Conditionally Generating Static Routes</i>• aggregate on page 2858• generate on page 2888• <i>static</i> |
|------------------------------|---|

confederation


Syntax	<code>confederation <i>confederation-autonomous-system</i> members [<i>autonomous-systems</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify the routing device's confederation AS number.</p> <p>If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more <i>confederations</i>. Each confederation is identified by its own AS number, which is called a <i>confederation AS number</i>. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.</p> <p>The BGP path attributes NEXT_HOP, LOCAL_PREF, and MULTI_EXIT_DISC, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.</p> <p>Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.</p> <p>Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.</p> <p>If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the confederation statement.</p> <p>Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.</p> <p>Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.</p>
Options	<p><i>autonomous-systems</i>—AS numbers of the confederation members. Range: 1 through 65,535</p> <p><i>confederation-autonomous-system</i>—Confederation AS number. Use one of the numbers assigned to you by the NIC. Range: 1 through 65,535</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Confederations on page 3422](#)

description (Routing Instances)

Syntax	description text;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the show route instance detail command and has no effect on the operation of the routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Instances on PE Routers in VPNs• show route instance on page 3057

discard

Syntax	discard;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</p> <p>To propagate static routes into the routing protocols, include the discard statement when you define the route, along with a routing policy.</p>
<div>  <p>NOTE: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as route <i>address/mask-length</i> next-hop 127.0.0.1) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.</p> </div>	
Default	When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Summarizing Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i>

- [aggregate on page 2858](#)
- [generate on page 2888](#)

export (Routing Options)

Syntax	<code>export [<i>policy-name</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Apply one or more policies to routes being exported from the routing table into the forwarding table.
Options	<i>policy-name</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Load Balancing BGP Traffic on page 3363• <i>Routing Policy Configuration Guide</i>

export-rib

Syntax	<code>export-rib <i>routing-table-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-options rib-groups <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the name of the routing table from which Junos OS should export routing information.
Options	<i>routing-table-name</i> —Routing table group name.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • import-rib on page 2891 • <i>passive</i>

fate-sharing

Syntax	<pre>fate-sharing { group <i>group-name</i> { cost <i>value</i>; from <i>address</i> <to <i>address</i>>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>
Options	<p>cost <i>value</i>—Cost assigned to the group.</p> <p>Range: 1 through 65,535</p> <p>Default: 1</p> <p>from <i>address</i>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate from addresses in the group.</p> <p>group <i>group-name</i>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p>to <i>address</i>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a from and a to address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Ingress Router for MPLS-Signaled LSPs</i>• <i>Junos OS MPLS Applications Configuration Guide</i>

flow

Syntax	<pre> flow { route <i>name</i> { match { <i>match-conditions</i>; } then { <i>actions</i>; } } firewall-install-disable; term-order (legacy standard); validation { traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } } </pre>
Hierarchy Level	<p>[edit routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. term-order statement introduced in Junos OS Release 10.0 Statement introduced in Junos OS Release 11.3 for the QFX Series. firewall-install-disable statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.</p>
Description	Configure a flow route.
Default	legacy
Options	<p>actions—An action to take if conditions match.</p> <p>firewall-install-disable—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).</p> <p>Default: If you omit the firewall-install-disable statement, the default behavior is firewall-install-disable mode.</p> <p>match-conditions—Match packets to these conditions.</p> <p>route <i>name</i>—Name of the flow route.</p> <p>standard—Specify to use version 7 or later of the flow-specification algorithm.</p> <p>term-order (legacy standard)—Specify the version of the flow-specification algorithm.</p>

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

then—Actions to take on matching packets.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Flow Routes*

flow-map

Syntax `flow-map flow-map-name {
 bandwidth (bps | adaptive);
 forwarding-cache {
 timeout (never | non-discard-entry-only | minutes);
 }
 policy [policy-names];
 redundant-sources [addresses];
}`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure multicast flow maps.

Options *flow-map-name*—Name of the flow-map.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Multicast Flow Map*

forwarding-cache (Flow Maps)

Syntax	<pre>forwarding-cache { timeout (minutes never non-discard-entry-only); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-options multicast flow-map <i>flow-map-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure multicast forwarding cache properties for the flow map.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

forwarding-cache (Multicast)

Syntax	<pre>forwarding-cache { family (inet inet6) { threshold { log-warning value; suppress value <reuse value>; } } threshold { log-warning value; suppress value <reuse value>; } timeout minutes; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits, the threshold at which a warning message is logged, and timeout values.</p> <p>Specify a value for the threshold at which to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value should be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.</p> <p>You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p>
Default	By default, there are no limits on the number of multicast forwarding cache entries.
Options	<p>family (inet inet6)—(Optional) Apply the configured thresholds to either IPv4 or IPv6 multicast forwarding cache entries.</p> <p>Default: By default, the configured thresholds are applied to both IPv4 and IPv6 multicast forwarding cache entries.</p>

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Multicast Forwarding Cache</i>

forwarding-table

Syntax	<pre>forwarding-table { export [policy--names]; (indirect-next-hop no-indirect-next-hop); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure information about the routing device's forwarding table. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Per-Packet Load Balancing on page 2826

generate

Syntax	<pre>generate { defaults { generate-options; } route destination-prefix { policy policy-name; generate-options; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>], [edit routing-options], [edit routing-options rib <i>routing-table-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure generated routes, which are used as routes of last resort.
Options	<p>defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the generate statement.</p> <p>generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in generate-options. Each option is explained separately.</p> <ul style="list-style-type: none">• (active passive);• as-path <i><as-path></i> <i><origin (egp igp incomplete)></i> <i><atomic-aggregate></i> <i><aggregator as-number in-address></i>;• (brief full);• community [<i>community-ids</i>];• discard;• (metric <i>metric2</i> <i>metric3</i> <i>metric4</i>) <i>value</i> <i><type type></i>;• (preference <i>preference2</i> color <i>color2</i>) <i>preference</i> <i><type type></i>;• tag <i>string</i>; <p>route destination-prefix—Configure a non-default generated route:</p> <ul style="list-style-type: none">• default—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0.

- *destination-prefix/prefix-length—/destination-prefix* is the network portion of the IP address, and *prefix-length* is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation • *Example: Conditionally Generating Static Routes*

import (Routing Options)

Syntax `import [policy-names];`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution **rib**],
[edit logical-systems *logical-system-name* routing-options resolution **rib**],
[edit routing-instances *routing-instance-name* routing-options resolution **rib**],
[edit routing-options resolution **rib**]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.


Description Specify one or more import policies to use for route resolution.

Options *policy-names*—Name of one or more import policies.


Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation • *Example: Configuring Route Resolution on PE Routers*


import-policy

Syntax	<code>import-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-options rib-groups <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to routes imported into the routing table group. The import-policy statement complements the import-rib statement and cannot be used unless you first specify the routing tables to which routes are being imported.
<hr/>	
<div> NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</div> <hr/>	
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i>• export-rib on page 2881• <i>passive</i>

import-rib

Syntax	<code>import-rib [<i>routing-table-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-options rib-groups <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be inet.0 or routing-instance-name.inet.0. For IPv6 import routing tables, the primary routing table must be inet6.0.</p> <p>In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.</p>
<div>  NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another. </div>	
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • export-rib on page 2881 • <i>passive</i>

indirect-next-hop

Syntax	(indirect-next-hop no-indirect-next-hop);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.
<div> NOTE:</div> <ul style="list-style-type: none">• When virtual private LAN service (VPLS) is configured on the routing device, the indirect-next-hop statement is configurable at the [edit routing-options forwarding-table] hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.• By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the no-indirect-next-hop statement.	
Default	Disabled.
Options	indirect-next-hop —Enable indirectly connected next hops. no-indirect-next-hop —Explicitly disable indirect next hops.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine</i>

install (Routing Options)

Syntax	(install no-install);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> static (defaults route)]</p> <p>[edit routing-options static (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.
Options	<p>install—Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the route portion of the static statement to override a no-install option specified in the defaults portion of the statement.</p> <p>no-install—Do not install the route into the forwarding table, even if it is the route with the lowest preference.</p> <p>Default: install</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>static</i>

instance-type

Syntax	instance-type virtual-router
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Specify the type of routing instance.
Options	virtual-router —Virtual router routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches</i>• <i>Configuring Virtual Routing Instances (CLI Procedure)</i>• Configuring Virtual Router Routing Instances on page 2829

interface (Multicast Static Routes)

Syntax	<pre> interface <i>interface-names</i> { disable; maximum-bandwidth <i>bps</i>; no-qos-adjust; reverse-oif-mapping { no-qos-adjust; } subscriber-leave-timer <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Enable multicast traffic on an interface.</p> <p>By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.</p> <p>You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.</p> <p>You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.</p> <p>Static routes must be configured before you can enable multicast on an interface. Configuring the interface statement alone does not install any routes into the routing table. This feature relies on the static route configuration.</p>
Options	<p><i>interface-names</i>—Name of one or more interfaces on which to enable multicast traffic.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Defining Interface Bandwidth Maximums</i> • <i>Example: Configuring Multicast with Subscriber VLANs</i>

interface (Routing Instances)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	For virtual router routing instances, configure an interface.

**NOTE:**

- You must configure only interfaces from the Node devices assigned to the network Node group. If you try to configure interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
 - You can configure an interface for one routing instance only.
-

Options	<i>interface-name</i> —Name of an interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Virtual Router Routing Instances on page 2829

interface (Routing Options)

Syntax	<pre>interface <i>interface-names</i> { maximum-bandwidth <i>bps</i>; no-qos-adjust; reverse-oif-mapping { no-qos-adjust; } subscriber-leave-timer <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable multicast traffic on an interface.



TIP: You cannot enable multicast traffic on an interface by using the `routing-options multicast interface` statement and configure PIM on the interface.

Options	<p><i>interface-name</i>—Names of the physical or logical interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Defining Interface Bandwidth Maximums</i> • <i>Example: Configuring Multicast with Subscriber VLANs</i>

interface-routes

Syntax

```
interface-routes {  
    family (inet | inet6) {  
        export {  
            lan;  
            point-to-point;  
        }  
    }  
    rib-group group-name;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Description Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0**, and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

Options **inet**—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Importing Direct and Static Routes Into a Routing Instance*
- *Example: Configuring Multiple Routing Instances of OSPF*
- *passive*

local-address (Routing Options)

Syntax local-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-options multicast **backup-pe-group** *group-name*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options **address**—Address of local PEs in the backup group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Ingress PE Redundancy*

martians

Syntax	<pre>martians { destination-prefix match-type <allow>; }</pre>
Hierarchy Level	<pre>[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options], [edit logical-systems logical-system-name routing-instances routing-instance-name routing-options rib routing-table-name], [edit logical-systems logical-system-name routing-options], [edit logical-systems logical-system-name routing-options rib routing-table-name], [edit routing-instances routing-instance-name routing-options], [edit routing-instances routing-instance-name routing-options rib routing-table-name], [edit routing-options], [edit routing-options rib routing-table-name]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure martian addresses.
Options	<p>allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The allow option is the only supported action.</p> <p>destination-prefix—Destination route you are configuring:</p> <ul style="list-style-type: none">• destination-prefix/prefix-length—destination-prefix is the network portion of the IP address, and prefix-length is the destination prefix length.• default—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address 0.0.0.0/0. <p>match-type—Criteria that the destination must match:</p> <ul style="list-style-type: none">• exact—Exactly match the route's mask length.• longer—The route's mask length is greater than the specified mask length.• orlonger—The route's mask length is equal to or greater than the specified mask length.• through destination-prefix—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.• upto prefix-length—The route's mask length falls between the two destination prefix lengths, inclusive.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related Documentation • *Example: Configuring Martian Addresses*

maximum-bandwidth (Routing Options)

Syntax	<code>maximum-bandwidth <i>bps</i>;</code>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances <i>instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options multicast interface <i>interface-name</i>]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <i>interface interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <i>interface interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <i>interface interface-name</i>],</p> <p>[edit routing-options multicast <i>interface interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>dynamic-profiles hierarchy level added in Junos OS Release 11.2.</p>
Description	Configure the multicast bandwidth for the interface.
Options	<p><i>bps</i>—Bandwidth rate, in bits per second, for the multicast interface.</p> <p>Range: 0 through any amount of bandwidth</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	• <i>Example: Defining Interface Bandwidth Maximums</i>

maximum-paths

Syntax	<code>maximum-paths <i>path-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>value</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure a limit for the number of routes installed in a routing table based upon the route path.



NOTE: The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

Options	log-interval <i>seconds</i> —(Optional) Minimum time interval (in seconds) between log messages. Range: 5 through 86,400
	log-only —(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.
	<i>path-limit</i> —Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected. Range: 1 through 4,294,967,295 ($2^{32} - 1$) Default: No default
	threshold <i>value</i> —(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the <i>path-limit</i> value that starts triggering the warnings. Range: 1 through 100




NOTE: When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

maximum-prefixes

Syntax	<code>maximum-prefixes <i>prefix-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>percentage</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Configure a limit for the number of routes installed in a routing table based upon the route prefix.</p> <p>Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.</p>
<div>  <p>NOTE: The <code>maximum-prefixes</code> statement is similar to the <code>maximum-paths</code> statement. The <code>maximum-prefixes</code> statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:</p> <pre> OSPF 10.10.10.0/24 ISIS 10.10.10.0/24 </pre> <p>These are two routes, but only one destination (prefix). The <code>maximum-paths</code> limit applies the total number of routes (two). The <code>maximum-prefixes</code> limit applies to the total number of unique prefixes (one).</p> </div>	
Options	<p>log-interval <i>seconds</i>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p>Range: 5 through 86,400</p> <p>log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><i>prefix-limit</i>—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.</p> <p>Range: 1 through 4,294,967,295</p> <p>Default: No default</p> <p>threshold <i>value</i>—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the <i>prefix-limit</i> value that starts triggering the warnings.</p>

Range: 1 through 100



NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs](#)

med-igp-update-interval

Syntax med-igp-update-interval *minutes*;

Hierarchy Level [edit routing-options]

Release Information Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp offset delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options *minutes*—Interval to delay MED updates.

Range: 10 through 600

Default: 10 minutes

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3236](#)
- [metric-out on page 3542](#)

metric (Aggregate, Generated, or Static Route)

Syntax	(metric metric2 metric3 metric4) <i>metric</i> <type type>;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)], [edit routing-options (aggregate generate static) (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2 , metric3 , and metric4 .
Options	metric —Metric value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) type type —(Optional) Type of route. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric. Range: 1 through 16
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Summarizing Static Routes Through Route Aggregation</i>• <i>Example: Conditionally Generating Static Routes</i>• aggregate on page 2858• generate on page 2888• <i>static</i>

multicast (Routing Options)

Syntax

```
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  interface interface-name {
    enable;
  }
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  ssm-groups {
    address;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure generic multicast properties.



NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347](#)
- *Examples: Configuring the Multicast Forwarding Cache*
- *Multicast Protocols Configuration Guide*
- ([indirect-next-hop on page 2892](#) | no-indirect-next-hop)

no-qos-adjust

Syntax	no-qos-adjust;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i> reverse-oif-mapping]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement added to [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], and [edit routing-options multicast interface <i>interface-name</i>] hierarchy levels in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multicast with Subscriber VLANs</i>

options (Routing Options)

Syntax	options { syslog (level <i>level</i> upto level <i>level</i>); }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.
Options	<p>level <i>level</i>—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:</p> <ul style="list-style-type: none"> • alert—Conditions that should be corrected immediately, such as a corrupted system database. • critical—Critical conditions, such as hard drive errors. • debug—Software debugging messages. • emergency—Panic or other conditions that cause the system to become unusable. • error—Standard error conditions. • info—Informational messages. • notice—Conditions that are not error conditions, but might warrant special handling. • warning—System warning messages. <p>upto level <i>level</i>—Log all messages up to a particular level.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • syslog in the <i>Junos OS System Basics Configuration Guide</i>

pim-to-igmp-proxy

Syntax	<pre>pim-to-igmp-proxy { upstream-interface [interface-names]; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-igmp-proxy statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring PIM-to-IGMP Message Translation

policy (Aggregate and Generated Routes)

Syntax	<code>policy <i>policy-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Associate a routing policy when configuring an aggregate or generated route's destination prefix in the routes part of the aggregate or generate statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.</p> <p>If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or <i>primary</i>, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.</p> <p>The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:</p> <ol style="list-style-type: none"> 1. Compare the protocol's preference of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table. 2. Compare the protocol's preference2 of the contributing routes. The lower preference2 value is better. If only one route has preference2, then this route is preferred. 3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values. <ol style="list-style-type: none"> a. The primary contributor is the numerically smallest prefix value. b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options	<i>policy-name</i> —Name of a routing policy.
Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Summarizing Routes Through Route Aggregation</i>• <i>Example: Conditionally Generating Static Routes</i>• aggregate on page 2858• generate on page 2888

policy (Flow Maps)

Syntax	<code>policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure a flow map policy.
Options	<i>policy-names</i> —Name of one or more policies for flow mapping.
Required Privilege	routing—To view this statement in the configuration.
Level	

policy-options

```
Syntax  policy-options
        application-maps application-map-name {
            application application-name {
                code-points [ aliases ] [ bit-patterns ];
            }
        }
        policy-statement policy-name {
            term term-name {
                from {
                    family family-name;
                    match-conditions;
                    policy subroutine-policy-name;
                    prefix-list prefix-list-name;
                    prefix-list-filter prefix-list-name match-type <actions>;
                    route-filter destination-prefix match-type <actions>;
                    source-address-filter source-prefix match-type <actions>;
                }
                to {
                    match-conditions;
                    policy subroutine-policy-name;
                }
                then actions;
            }
        }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for the EX Series.

Description Configure options such as application maps for DCBX application protocol exchange and policy statements.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

policy-statement

Syntax	<pre> policy-statement <i>policy-name</i> { term <i>term-name</i> { from { family <i>family-name</i>; match-conditions; policy <i>subroutine-policy-name</i>; prefix-list <i>prefix-list-name</i>; prefix-list-filter <i>prefix-list-name</i> match-type <actions>; route-filter <i>destination-prefix</i> match-type <actions>; source-address-filter <i>source-prefix</i> match-type <actions>; } to { match-conditions; policy <i>subroutine-policy-name</i>; } then <i>actions</i>; } } </pre>
Hierarchy Level	<p>[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches. inet-mdt option introduced in Junos OS Release 10.0R2. Statement introduced in Junos OS Release 11.3 for the QFX Series. route-target option introduced in Junos OS Release 12.2.</p>
Description	<p>Define a routing policy, including subroutine policies.</p> <p>To list the routing policies under the [edit policy-options] hierarchy level by policy-statement <i>policy-name</i> in alphabetical order, enter the show policy-options configuration command.</p>
Options	<p>actions—(Optional) One or more actions to take if the conditions match. The actions are described in <i>Configuring Flow Control Actions</i>.</p> <p>family <i>family-name</i>—(Optional) Specify an address family protocol. Specify inet for IPv4. Specify inet6 for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify iso. For IPv4 multicast VPN traffic, specify inet-mvpn. For IPv6 multicast VPN traffic, specify inet6-mvpn. For multicast-distribution-tree (MDT) IPv4 traffic, specify inet-mdt. For BGP route target VPN traffic, specify route-target.</p>



NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the **family** statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Configuring Match Conditions in Routing Policy Terms*.

policy subroutine-policy-name—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see *Configuring Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

prefix-list prefix-list-name —Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter prefix-list-name—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match (see *Configuring Prefix List Filters*), and **actions** is the action to take if the prefixes match.

route-filter destination-prefix match-type <actions>—(Optional) List of routes on which to perform an immediate match; **destination-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **destination-prefix** matches.

source-address-filter source-prefix match-type <actions>—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. **source-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **source-prefix** matches.

term term-name—Name that identifies the term.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Defining Routing Policies</i>• <i>Configuring Routing Policies and Policy Objects in the Dynamic Database</i>• <i>dynamic-db</i>

ppm

Syntax	<pre>ppm { no-delegate-processing; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>(M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).</p> <p>After you disable PPM, PPM processing continues to run on the Routing Engine.</p> <p>In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The no-delegate-processing statement disables the default behavior and restores the legacy behavior.</p>
Default	Distributed PPM processing is enabled for all protocols that use PPM.
Options	no-delegate-processing —Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i>• Configuring Distributed Periodic Packet Management on page 2828

ppm (Ethernet Switching)

Syntax	ppm { centralized; }
Hierarchy Level	[edit protocols lacp]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.1 for T Series devices.
Description	Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets. This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing configuration statement in the [edit routing-options ppm] statement hierarchy.
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i> • Configuring Distributed Periodic Packet Management on page 2828

preference (Routing Options)

Syntax	<code>(preference preference2 color color2) preference <type type>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value (preference2), as well as colors, which are even finer-grained preference values (color and color2).</p> <p>If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.</p>
Options	<p>preference—Preference value. A lower number indicates a more preferred route.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 5 (for static routes), 130 (for aggregate and generated routes)</p> <p>type type—(Optional) Type of route.</p> <p>Range: 1 through 16</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>Example: Summarizing Routes Through Route Aggregation</i> • <i>Example: Conditionally Generating Static Routes</i> • aggregate on page 2858 • generate on page 2888 • <i>static</i>

prefix

Syntax	<code>prefix destination-prefix;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit routing-options multicast scope <i>scope-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure the prefix for multicast scopes.
Options	destination-prefix —Address range for the multicast scope.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Administrative Scoping</i> • <i>Example: Creating a Named Scope for Multicast Scoping</i> • <i>multicast</i>

protocols

```

Syntax protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf-configuration ...
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf3-configuration ...
    }
    pim {
        ... pim-configuration ...
    }
    rip {
        ... rip-configuration ...
    }
    ripng {
        ... ripng-configuration ...
    }
    rstp {
        rstp-configuration;
    }
    vstp {
        vstp configuration;
    }
    vpls {
        vpls configuration;
    }
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Support for RIPv6 introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 11.1 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

Options **bgp**—Specify BGP as the protocol for a routing instance.

isis—Specify IS-IS as the protocol for a routing instance.

ldp—Specify LDP as the protocol for a routing instance.

l2vpn—Specify Layer 2 VPN as the protocol for a routing instance.

msdp—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.

mstp—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.

ospf—Specify OSPF as the protocol for a routing instance.

ospf3—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.



NOTE: OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

pim—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.

rip—Specify RIP as the protocol for a routing instance.

ripng—Specify RIP next generation (RIPv6) as the protocol for a routing instance.

rstp—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.

vstp—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.

vpls—Specify VPLS as the protocol for a routing instance.

Required Privilege Level **routing**—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multiple Routing Instances of OSPF*

qualified-next-hop (Static Routes)

Syntax `qualified-next-hop (address | interface-name) {
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
 simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 interface interface-name;
 metric metric;
 preference preference;
}`

Hierarchy Level `[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
 [edit logical-systems logical-system-name routing-options rib inet6.0 static route
 destination-prefix],
 [edit logical-systems logical-system-name routing-options static route destination-prefix],
 [edit routing-instances routing-instance-name routing-options static route destination-prefix],
 [edit routing-options rib inet6.0 static route destination-prefix],
 [edit routing-options static route destination-prefix]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options *address*—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop

interface for a static route, specify **qualified-next-hop *interface-name***, where *interface-name* is the name of the IPv4 or IPv6 unnumbered interface.



NOTE: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the `unnumbered-address <interface-name>` statement at the [edit interfaces <interface-name> unit <logical-unit-number> family <family-name>] hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Static Route Preferences and Qualified Next Hops • Example: Enabling BFD on Qualified Next Hops in Static Routes on page 2840


readvertise

Syntax	(readvertise no-readvertise);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure whether static routes are eligible to be readvertised by routing protocols:
Default	Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the no-readvertise statement.
Options	readvertise —Readvertise static routes. Include the readvertise statement when configuring an individual route in the route portion of the static statement to override a no-readvertise option specified in the defaults portion of the statement. no-readvertise —Mark a static route as being ineligible for readvertisement. Include the no-readvertise option when configuring the route.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Controlling Static Routes in Routing and Forwarding Tables</i>• <i>static</i>

redundant-sources

Syntax	<code>redundant-sources [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>],</p> <p>[edit routing-options multicast flow-map <i>flow-map-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a list of redundant sources for multicast flows defined by a flow map.
Options	<i>addresses</i> —List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring a Multicast Flow Map</i>

resolution

Syntax	<pre> resolution { rib routing-table-name { import [policy-names]; resolution-ribs [routing-table-names]; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.</p> <p>For example, you might want to direct inet.2 route resolution to use topology routing tables :red.inet.0 and :blue.inet.0 for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.</p> <p>You can specify up to two routing tables in the resolution-ribs statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3], inet.0 is checked first and then inet.3 is checked.</p>
	<p> NOTE: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the show system processes and the show task memory commands. For more information, see <i>Routing Protocol Process Memory FAQs</i>.</p>
	<p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Configuring Route Resolution on PE Routers*
 - *Example: Configuring Route Resolution on Route Reflectors*
 - *Example: Configuring Multitopology Routing Based on a Multicast Source*

resolution-ribs

Syntax	<code>resolution-ribs [<i>routing-table-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options resolution rib],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit routing-options resolution rib]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify one or more routing tables to use for route resolution.</p> <p>This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is inet.3, but you want to allow fallback resolution through inet.0. One example use case is overriding the bgp.rtarget.0 (family route-target) routing table resolution from using only inet.3 to using both inet.3 and inet.0.</p>
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Resolution on PE Routers</i> • <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>

resolve

Syntax	resolve;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the inet.0 and inet.3 routing tables.
Default	Static routes can point only to a directly connected next hop.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>static</i>

retain

Syntax	(no-retain retain);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:
Default	Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.
Options	<p>no-retain—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the no-retain statement. Include this statement when configuring an individual route in the route portion of the static statement to override a retain option specified in the defaults portion of the statement.</p> <p>retain—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Examples: Configuring Static Routes</i> • <i>static</i>

reverse-oif-mapping

Syntax	<code>reverse-oif-mapping { no-qos-adjust; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit routing-options multicast interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 9.2 for EX Series switches. The no-qos-adjust statement added in Junos OS Release 9.5. The no-qos-adjust statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multicast with Subscriber VLANs</i>

rpf-check-policy (Routing Options RPF)

Syntax	<code>rpf-check-policy [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.
Options	<i>policy-names</i> —Name of one or more multicast RPF check policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring RPF Policies</i>

rib (General)

```
Syntax  rib routing-table-name {  
        aggregate {  
            defaults {  
                ... aggregate-options ...  
            }  
            route destination-prefix {  
                policy policy-name;  
                ... aggregate-options ...  
            }  
        }  
        generate {  
            defaults {  
                generate-options;  
            }  
            route destination-prefix {  
                policy policy-name;  
                generate-options;  
            }  
        }  
        martians {  
            destination-prefix match-type <allow>;  
        }  
    }  
    static {  
        defaults {  
            static-options;  
        }  
        rib-group group-name;  
        route destination-prefix {  
            next-hop;  
            static-options;  
        }  
    }  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.



NOTE: The IPv4 multicast routing table (`inet.1`) and the IPv6 multicast routing table (`inet6.1`) are not supported for this statement.

Default If you do not specify a routing table name with the *routing-table-name* option, the software uses the default routing tables, which are `inet.0` for unicast routes and `inet.1` for the multicast cache.

Options *routing-table-name*—Name of the routing table, in the following format:
protocol [.identifier].

In a routing instance, the routing table name must include the routing instance name.

For example, if the routing instance name is `link0`, the routing table name might be `link0.inet6.0`.

- *protocol* is the protocol family. It can be `inet6` for the IPv6 family, `inet` for the IPv4 family, `iso` for the ISO protocol family, or *instance-name.iso.0* for an ISO routing instance.
- *identifier* is a positive integer that specifies the instance of the routing table.

Default: `inet.0`

The remaining statements are explained separately.

Required Privilege Level `routing`—To view this statement in the configuration.
`routing-control`—To add this statement to the configuration.

Related Documentation

- *Example: Creating Routing Tables*
- *passive*

rib (Route Resolution)

Syntax	<pre>rib <i>routing-table-name</i> { import [<i>policy-names</i>]; resolution-ribs [<i>routing-table-names</i>]; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution], [edit logical-systems <i>logical-system-name</i> routing-options resolution], [edit routing-instances <i>routing-instance-name</i> routing-options resolution], [edit routing-options resolution]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Specify a routing table name for route resolution. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Route Resolution on PE Routers</i>

rib-group (Routing Options)

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options interface-routes],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options interface-routes],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options interface-routes],</p> <p>[edit routing-options interface-routes],</p> <p>[edit routing-options rib <i>routing-table-name</i> static],</p> <p>[edit routing-options static]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure which routing table groups interface routes are imported into.
Options	<p><i>group-name</i>—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • interface-routes on page 2898 • rib-groups on page 2936

rib-groups

Syntax	<pre>rib-groups { group-name { export-rib group-name; import-policy [policy-names]; import-rib [group-names]; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.</p> <p>Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the import-rib statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the export-rib statement).</p> <p>The first routing table you specify is the <i>primary routing table</i>, and any additional routing tables are the <i>secondary routing tables</i>.</p> <p>The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify inet.0 as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify inet6.0 as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (inet6.x).</p> <p>In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the import-rib statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same import-rib statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the inet6.3 routing table with IPv6 addresses that are compatible with IPv4. Specify inet.0 as the primary routing table, and specify inet6.3 as a secondary routing table.</p>



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.



NOTE: If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- [rib-group on page 2935](#)

route-record

Syntax route-record;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Export the AS path and routing information to the traffic sampling process.


Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Flow Aggregation*
- *Junos Services Interfaces Configuration Release 12.3*

router-id

Syntax	<code>router-id address;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options],</code> <code>[edit routing-options]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify the routing device's IP address.</p> <p>The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.</p>
<div> NOTE: We strongly recommend that you configure the router identifier under the <code>[edit routing-options]</code> hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.</div>	
<p>For more information about the router identifier in OSPF, see “Example: Configuring an OSPF Router Identifier” on page 3810.</p>	
Options	address —IP address of the routing device. Default: Address of the first interface encountered by Junos OS
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring External BGP Peering on page 3149• Examples: Configuring Internal BGP Peering on page 3172

routing-instances

Syntax	<pre>routing-instances <i>routing-instance-name</i> { <i>description</i>; instance-type virtual-router; interface <i>interface-name</i>; protocols; routing-options }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(QFabric switches only) Configure a virtual router routing instance.
Options	<p><i>routing-instance-name</i>—Name of this routing instance.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Virtual Router Routing Instances on page 2829

routing-options

Syntax	routing-options { ... }
Hierarchy Level	[edit], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure protocol-independent routing properties.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

scope

Syntax	<pre>scope scope-name { interface [interface-names]; prefix destination-prefix; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure multicast scoping.
Options	<i>scope-name</i> —Name of the multicast scope. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Creating a Named Scope for Multicast Scoping</i>

scope-policy

Syntax `scope-policy [policy-names];`

Hierarchy Level [edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]



NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance at the [edit routing-instances *routing-instance-name* routing-options multicast] or [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast] hierarchy level.

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.


Options *policy-names*—Name of one or more multicast scope policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [scope on page 2940](#)
- *Example: Using a Scope Policy for Multicast Scoping*

source-routing

Syntax	source-routing { (ip ipv6) }
Hierarchy Level	[edit routing-options]
Release Information	Statement for IPv6 introduced in Junos OS Release 8.2. Statement for IPv4 introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Enable source routing.</p> <p>Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.</p> <div> NOTE: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.</div>
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Filter-Based Forwarding on the Source Address</i>

static (Routes)

Syntax	<pre> static { defaults { <i>static-options</i>; } rib-group <i>group-name</i>; route <i>destination-prefix</i> { next-hop <i>address</i>; <i>next-hop options</i>; qualified-next-hop <i>address</i> { metric <i>metric</i>; preference <i>preference</i>; } <i>static-options</i>; } } </pre>
Hierarchy Level	[edit routing-options], [edit routing-options rib <i>routing-table-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure static routes to be installed in the routing table. You can specify any number of routes within a single static statement, and you can specify any number of static options in the configuration.
Options	<p>defaults—Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the static statement. This part of the static statement is optional.</p> <p>route <i>destination-prefix</i>—Destination of the static route.</p> <ul style="list-style-type: none"> defaults—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0. <i>destination-prefix/prefix-length</i>—<i>destination-prefix</i> is the network portion of the IP address, and <i>prefix-length</i> is the destination prefix length. <i>next-hop address</i>—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET). <i>nsap-prefix</i>—<i>nsap-prefix</i> is the network service access point (NSAP) address for ISO. <p><i>next-hop options</i>—Additional information for how to manage forwarding of packets to the next hop.</p> <ul style="list-style-type: none"> discard—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

- **iso-net**—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table *routing-table-name***—Name of the next routing table to the destination.
- **receive**—Install a receive route for this destination into the routing table.
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in **static-options**. Each of the options is explained separately.

- **(active | passive);**
- **(install | no-install);**
- **(metric | metric2 | metric3 | metric4) *value* <type type>;**
- **(preference | preference2 | color | color2) *preference* <type type>;**
- **(resolve | no-resolve);**
- **(no-retain | retain);**

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static Routing on page 2826

subscriber-leave-timer

Syntax	<code>subscriber-leave-timer seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.
Options	<p>seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.</p> <p>Range: 0 through 30</p> <p>Default: 0 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Multicast with Subscriber VLANs</i>

tag (Routing Options)

Syntax	<code>tag string;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options aggregate generate static) (defaults route)],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</code> <code>[edit routing-options (aggregate generate static) (defaults route)],</code> <code>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Associate an OSPF tag with a static, aggregate, or generated route.
Default	No OSPF tag strings are associated with static routes.
Options	<i>string</i> —OSPF tag string.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Configuring Static Routes</i>• <i>Example: Summarizing Routes Through Route Aggregation</i>• <i>Example: Conditionally Generating Static Routes</i>• aggregate on page 2858• generate on page 2888• <i>static</i>

threshold (Multicast Forwarding Cache)

Syntax	<pre>threshold { log-warning <i>value</i>; suppress <i>value</i> <reuse <i>value</i>>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache (inet inet6)],</p> <p>[edit routing-options multicast forwarding-cache],</p> <p>[edit routing-options multicast forwarding-cache family (inet inet6)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>To confirm the configured threshold values, use the show multicast forwarding-cache statistics command.</p>
Options	<p>reuse <i>value</i>—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the suppress value.</p> <p>Range: 1 through 200,000</p> <p>suppress <i>value</i>—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the reuse value.</p> <p>Range: 1 through 200,000</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Examples: Configuring the Multicast Forwarding Cache</i>

timeout (Flow Maps)

Syntax	timeout (never non-discard-entry-only <i>minutes</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the timeout value for multicast forwarding cache entries associated with the flow map.
Options	minutes —Length of time that the forwarding cache entry remains active. Range: 1 through 720 never non-discard-entry-only —Specify that the forwarding cache entry always remain active. If you omit the non-discard-entry-only option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the non-discard-entry-only option, entries with forwarding states are kept forever, and entries with pruned states time out.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

timeout (Multicast)

Syntax	<code>timeout <i>minutes</i> <family (inet inet6)>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache], [edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache], [edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache], [edit routing-options multicast forwarding-cache]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the timeout value for multicast forwarding cache entries.
Options	<i>minutes</i> —Length of time that the forwarding cache limit remains active. Range: 1 through 720 <i>family (inet inet6)</i> —(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the timeout statement globally for the multicast forwarding cache or including the family statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive. Default: By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries.
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Multicast Forwarding Cache</i>

traceoptions (Routing Options)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Example: Tracing Global Routing Protocol Operations</i>• <i>Tracing Nonstop Active Routing Synchronization Events</i>
------------------------------	---

upstream-interface

Syntax	<code>upstream-interface [<i>interface-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</p> <p>[edit routing-options multicast pim-to-igmp-proxy],</p> <p>[edit routing-options multicast pim-to-mld-proxy]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the pim-to-igmp-proxy statement), or into corresponding MLD report or leave messages (if you include the pim-to-mld-proxy statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.</p>
Options	<p><i>interface-names</i>—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([]).</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PIM-to-IGMP Message Translation • Configuring PIM-to-MLD Message Translation

Administration

- [Routine Monitoring on page 2955](#)
- [Operational Commands on page 2957](#)

Routine Monitoring

- [Monitoring Routing Information on page 2955](#)
- [Verifying That Virtual Router Routing Instances Are Working on page 2956](#)

Monitoring Routing Information

- Purpose** Use the monitoring functionality to view the `inet.0` routing table on the routing device.
- Action** To view the routing table, enter the following commands in the CLI interface:
- `show route terse`
 - `show route detail`
- Meaning** [Table 261 on page 2955](#) describes the different filters, their functions, and the associated actions.
- [Table 262 on page 2956](#) summarizes key output fields in the routing information display.

Table 261: Filtering Route Messages

Field	Function	Your Action
Destination Address	Specifies the destination address of the route.	Enter the destination address.
Next hop address	Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	Enter the next hop address.
Best route	Specifies only the best route available.	Select the view details of the best route.
Inactive routes	Specifies the inactive routes.	Select the view details of inactive routes.
Exact route	Specifies the exact route.	Select the view details of the exact route.

Table 261: Filtering Route Messages (*continued*)

Field	Function	Your Action
Hidden routes	Specifies the hidden routes.	Select the view details of hidden routes.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click Search .

Table 262: Summary of Key Routing Information Output Fields

Field	Values	Additional Information
Static Route Addresses	The list of static route addresses.	
Protocol	Protocol from which the route was learned: Static , Direct , Local .	
Preference	The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.
Next-Hop	Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	<p>If a next hop is listed as Discard, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the discard attribute has been set.</p> <p>If a next hop is listed as Reject, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected.</p> <p>If a next hop is listed as Local, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).</p>
Age	How long the route has been active.	
State	Flags for this route.	There are many possible flags.

Related Documentation • [Configuring Static Routing on page 2826](#)

Verifying That Virtual Router Routing Instances Are Working

Purpose After creating a virtual router routing instance, verify that it has been set up properly.

Action 1. Use the **show route instance** command to list all the routing instances and their properties:

```
user@switch> show route instance
```


Instance	Type	Active/holdown/hidden
Primary RIB		
master	forwarding	
inet.0		4/0/1
__juniper_private1__	forwarding	
__juniper_private1__.inet.0		1/0/3
__juniper_private2__	forwarding	
__juniper_private2__.inet.0		0/0/1
__juniper_private3__	forwarding	
__juniper_private3__.inet.0		1/0/2
__juniper_private4__	forwarding	
__juniper_private4__.inet.0		4/0/2
__master.anon__	forwarding	
r1	virtual-router	
r2	virtual-router	

- Use the **show route forwarding-table** command to view the forwarding table information for each routing instance:

```
user@switch> show route forwarding-table
Routing table: r1---qfabric.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0          Type Index NhRef Netif
0.0.0.0/32       perm  0          dscd  1626   1
224.0.0.0/4      perm  0          mdsc  1627   1
224.0.0.1/32     perm  0 224.0.0.1    mcst  1623   1
255.255.255.255/32 perm  0          bcst  1624   1
```

Meaning The output displays the routing table information and confirms that the virtual router routing instances have been created and the links are up.

Related Documentation

- [Configuring Virtual Router Routing Instances on page 2829](#)

Operational Commands

- [clear ipv6 neighbors](#)
- [show as-path](#)
- [show as-path domain](#)
- [show as-path summary](#)
- [show ipv6 neighbors](#)
- [show ipv6 router-advertisement](#)
- [show route](#)
- [show route active-path](#)
- [show route all](#)

- [show route aspath-regex](#)
- [show route best](#)
- [show route brief](#)
- [show route community](#)
- [show route community-name](#)
- [show route damping](#)
- [show route detail](#)
- [show route exact](#)
- [show route export](#)
- [show route extensive](#)
- [show route flow validation](#)
- [show route forwarding-table](#)
- [show route inactive-path](#)
- [show route inactive-prefix](#)
- [show route instance](#)
- [show route label](#)
- [show route label-switched-path](#)
- [show route martians](#)
- [show route next-hop](#)
- [show route no-community](#)
- [show route protocol](#)
- [show route range](#)
- [show route receive-protocol](#)
- [show route resolution](#)
- [show route snooping](#)
- [show route source-gateway](#)
- [show route summary](#)
- [show route table](#)
- [show route terse](#)

clear ipv6 neighbors

Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Clear IPv6 neighbor cache information.
Options	<p>none—Clear all IPv6 neighbor cache information.</p> <p>all—(Optional) Clear all IPv6 neighbor cache information.</p> <p>host <i>hostname</i>—(Optional) Clear the information for the specified IPv6 neighbors.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ipv6 neighbors on page 2969
List of Sample Output	clear ipv6 neighbors on page 2959
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

show as-path

List of Syntax	Syntax on page 2960 Syntax (EX Series Switches) on page 2960
Syntax	<code>show as-path</code> <code><brief detail></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show as-path</code> <code><brief detail></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show as-path summary on page 2967
List of Sample Output	show as-path on page 2961 show as-path detail on page 2962
Output Fields	Table 263 on page 2961 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear.

Table 263: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	brief none
Bucket	Bucket number.	All levels
Count	Number of AS path entries in this bucket.	All levels
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregat—Routing device has summarized a range of prefixes. 	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	detail
neighbor as	AS peer address.	detail
length	Length of the AS path.	detail
segments	Length of the AS segment descriptor.	detail
references	Path reference count.	detail

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
  Bucket 0      Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I

```

```

14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I
14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

```

...

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I

```

```

    domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```

show as-path domain

List of Syntax	Syntax on page 2964 Syntax (EX Series Switches) on page 2964
Syntax	show as-path domain <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path domain
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path domain information.
Options	none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show as-path domain on page 2966
Output Fields	Table 264 on page 2964 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear

Table 264: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
Number Paths	Number of known AS paths.
Flags	Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance.
Local AS	AS number of the local routing device.

Table 264: show as-path domain Output Fields (*continued*)

Field Name	Field Description
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

show as-path summary

List of Syntax	Syntax on page 2967 Syntax (EX Series Switches) on page 2967
Syntax	<pre>show as-path summary <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show as-path summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display autonomous system (AS) path summary information.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—(Optional) Display AS path summary information for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show as-path on page 2960
List of Sample Output	show as-path summary on page 2968
Output Fields	<p>Table 265 on page 2967 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 265: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.

Table 265: show as-path summary Output Fields (*continued*)

Field Name	Field Description
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets  Max   Min   Avg   Std deviation
 30425    1024     95    12    29    6.481419
```

show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ipv6 neighbors on page 2959
List of Sample Output	show ipv6 neighbors on page 2969
Output Fields	Table 266 on page 2969 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 266: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up , down , incomplete , reachable , stale , or unreachable .
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no .
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no .
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no
fe-1/2/0.1

```

fe80::2a0:a514:0:24c fe-1/2/0.1	00:05:85:8f:c8:bd	stale	258	yes	no
fe80::2a0:a514:0:64c fe-1/2/1.5	00:05:85:8f:c8:bd	stale	111	yes	no
fe80::2a0:a514:0:a4c fe-1/2/2.9	00:05:85:8f:c8:bd	stale	327	yes	no

show ipv6 router-advertisement

Syntax	<pre>show ipv6 router-advertisement <conflicts> <interface <i>interface</i>> <logical-system (all <i>logical-system-name</i>)> <prefix <i>prefix/prefix length</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.
Options	<p>none—Display all IPv6 router advertisement information for all interfaces.</p> <p>conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.</p> <p>interface <i>interface</i>—(Optional) Display IPv6 router advertisement information for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>prefix/prefix length</i>—(Optional) Display IPv6 router advertisement information for the specified prefix.</p>
Additional Information	The display identifies conflicting information by enclosing the value the router is advertising in brackets.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ipv6 router-advertisement
List of Sample Output	<p>show ipv6 router-advertisement on page 2972</p> <p>show ipv6 router-advertisement conflicts on page 2973</p> <p>show ipv6 router-advertisement prefix on page 2973</p>
Output Fields	Table 267 on page 2971 describes the output fields for the show ipv6 router-advertisement command. Output fields are listed in the approximate order in which they appear.

Table 267: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and the elapsed time since they were sent.

Table 267: show ipv6 router-advertisement Output Fields (*continued*)

Field Name	Field Description
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
Managed	Managed address configuration flag: 0 (stateless) or 1 (stateful).
Other configuration	Other stateful configuration flag: 0 (stateless) or 1 (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: 0 (not onlink) or 1 (onlink).
Autonomous	Autonomous address configuration flag: 0 (not autonomous) or 1 (autonomous).

Sample Output

show ipv6 router-advertisement

```

user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
  Managed: 0
  Other configuration: 0 [1]
  Reachable time: 0 ms
  Default lifetime: 1800 sec

```



```
Retransmit timer: 0 ms
Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
  Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
  Managed: 0
  Other configuration: 0
  Reachable time: 0 ms
  Default lifetime: 180 sec [1800 sec]
  Retransmit timer: 0 ms
  Current hop limit: 64
  Prefix: 8040:1::/64
    Valid lifetime: 2592000 sec
    Preferred lifetime: 604800 sec
    On link: 1
    Autonomous: 1
```

show route

List of Syntax	Syntax on page 2974 Syntax (EX Series Switches) on page 2974
Syntax	<pre>show route <all> <destination-prefix> <logical-system (all logical-system-name)> <private></pre>
Syntax (EX Series Switches)	<pre>show route <all> <destination-prefix> <private></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option private introduced in Junos OS Release 9.5. Option private introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the active entries in the routing tables.
Options	<p>none—Display brief information about all active entries in the routing tables.</p> <p>all—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>private—(Optional) Display information only about all private, or internal, routing tables.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring RIP on page 4109• Example: Configuring RIPng• Example: Configuring IS-IS• Examples: Configuring Internal BGP Peering on page 3172• Examples: Configuring External BGP Peering on page 3149• Examples: Configuring OSPF Routing Policy on page 3921
List of Sample Output	show route on page 2977 show route destination-prefix on page 2978

[show route extensive on page 2978](#)

Output Fields [Table 268 on page 2975](#) describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 268: show route Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. <p>However, if you have configured advertisement of multiple routes (with the add-path or advertise-inactive statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> • hidden (routes that are not used because of a routing policy).
<i>destination-prefix</i>	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.

Table 268: show route Output Fields (*continued*)

Field Name	Field Description
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
<i>weeks:days</i> <i>hours:minutes:seconds</i>	How long the route been known (for example, 2w4d 13:11:14 , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 268: show route Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
to	<p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p>
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • lsp-path-name—Name of the LSP used to reach the next hop. • label-action—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).

Sample Output

show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    * [MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

show route destination-prefix

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                   > to 192.168.167.254 via fxp0.0
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
      Local AS: 65500 Peer AS: 65500
    Age: 3  Metric2: 1
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100
    Import Accepted
    Localpref: 100
    Router ID: 10.0.0.30
    Primary Routing Table bgp.mvpn.0
    Indirect next hops: 1
      Protocol next hop: 10.0.0.40 Metric: 1
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
      10.0.0.40/32 Originating RIB: inet.3
        Metric: 1                               Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route active-path

List of Syntax	Syntax on page 2979 Syntax (EX Series Switches) on page 2979
Syntax	<pre>show route active-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route active-path <brief detail extensive terse></pre>
Release Information	<p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.
Options	<p>none—Display all active routes.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route active-path on page 2979 show route active-path brief on page 2980 show route active-path detail on page 2980 show route active-path extensive on page 2981 show route active-path terse on page 2983
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32   *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
```

```
100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21   *[Direct/0] 21:33:52
                  > via fxp0.0
192.168.70.19/32  *[Local/0] 21:33:52
                  Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 2979](#).

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
```



```

Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

show route active-path extensive

```

user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397

```

```
Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
```

AS path: I

show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.255.70.19/32	D	0			>1o0.0	
*	10.255.71.50/32	I	15	10		>100.1.2.1	
*	100.1.2.0/24	D	0			>so-2/1/3.0	
*	100.1.2.2/32	L	0			Local	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.70.19/32	L	0			Local	

show route all

List of Syntax	Syntax on page 2984 Syntax (EX Series Switches) on page 2984
Syntax	show route all <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route all
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about all routes in all routing tables, including private, or internal, tables.
Options	none —Display information about all routes in all routing tables, including private, or internal, tables. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route all on page 2984
Output Fields	In Junos OS Release 9.5 and later, only the output fields for the show route all command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```

user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
```

show route aspath-regex

List of Syntax	Syntax on page 2986 Syntax (EX Series Switches) on page 2986
Syntax	<code>show route aspath-regex <i>regular-expression</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route aspath-regex <i>regular-expression</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.
Options	<i>regular-expression</i> —Regular expression that matches an entire AS path. <i>logical-system</i> (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Additional Information	<p>You can specify a regular expression as:</p> <ul style="list-style-type: none">• An individual AS number• A period wildcard used in place of an AS number• An AS path regular expression that is enclosed in parentheses <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none">• <i>{m,n}</i>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term.• <i>{m}</i>—Exactly <i>m</i> repetitions of the AS path term.• <i>{m,}</i>—<i>m</i> or more repetitions of the AS path term.• <i>*</i>—Zero or more repetitions of an AS path term.• <i>+</i>—One or more repetitions of an AS path term.• <i>?</i>—Zero or one repetition of an AS path term.• <i>aspath_term</i> <i>aspath_term</i>—Match one of the two AS path terms. <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 ."</pre>

Required Privilege Level	view
List of Sample Output	show route aspath-regex (Matching a Specific AS Number) on page 2987 show route aspath-regex (Matching Any Path with Two AS Numbers) on page 2987
Output Fields	For information about output fields, see the output field table for the show route command.

Sample Output

show route aspath-regex (Matching a Specific AS Number)

```

user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...

```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```

user@host> show route aspath-regex ?.* 234 3561.*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...

```

show route best

List of Syntax	Syntax on page 2988 Syntax (EX Series Switches) on page 2988
Syntax	<code>show route best <i>destination-prefix</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route best <i>destination-prefix</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route best on page 2988 show route best detail on page 2989 show route best extensive on page 2990 show route best terse on page 2990
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSPV/7] 1d 13:20:13, metric 2
```



```

> via so-0/3/0.0, label-switched-path green-r1-r3

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8      *[Direct/0] 2d 01:43:34
                 > via fxp2.0
                 [Direct/0] 2d 01:43:34
                 > via fxp1.0

```

show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  *OSPF   Preference: 10
          Next-hop reference count: 9
          Next hop: 10.31.1.6 via ge-3/1/0.0, selected
          Next hop: via so-0/3/0.0
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:06      Metric: 2
          Area: 0.0.0.0
          Task: OSPF
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 5
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100016
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:59      Metric: 2
          Task: RSVP
          Announcement bits (1): 1-Resolve tree 2
          AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp2.0, selected
          State: <Active Int>
          Age: 2d 1:44:20
          Task: IF
          AS path: I
  Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp1.0, selected
          State: <NotBest Int>
          Inactive reason: No difference
          Age: 2d 1:44:20

```

Task: IF
AS path: I

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 2989](#).

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  0 10      2          >10.31.1.6
                        so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  R  7      2          >so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.0.0.0/8        D  0          >fxp2.0
                    D  0          >fxp1.0
```

show route brief

List of Syntax	Syntax on page 2991 Syntax (EX Series Switches) on page 2991
Syntax	show route brief <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route brief <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display brief information about the active entries in the routing tables.
Options	none —Display all active entries in the routing table. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route brief on page 2991
Output Fields	For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22     *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18     *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22    *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0

```

```
192.168.164.51/32 *[Local/0] 2w4d 13:11:14
                  Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14
                  > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16    *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32   *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32     *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

show route community

List of Syntax	Syntax on page 2993 Syntax (EX Series Switches) on page 2993
Syntax	show route community <i>as-number:community-value</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route community <i>as-number:community-value</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.
Options	<p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show route detail on page 3002
List of Sample Output	show route community on page 2993
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community

```

user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

```

```
4.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 568 721 Incomplete
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16     *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1673 1675 1747 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

List of Syntax	Syntax on page 2995 Syntax (EX Series Switches) on page 2995
Syntax	show route community-name <i>community-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route community-name <i>community-name</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.
Options	<i>community-name</i> —Name of the community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route community-name on page 2995
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```

user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: 300 I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204

```

```
AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
*[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
AS path: 300 I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
*[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
*[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route damping

List of Syntax	Syntax on page 2997 Syntax (EX Series Switch and QFX Series) on page 2997
Syntax	show route damping (decayed history suppressed) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show route damping (decayed history suppressed) <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>decayed—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p>history—Display entries that have already been withdrawn, but have been logged.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp damping on page 3584 • show policy damping on page 3615
List of Sample Output	show route damping decayed detail on page 3000 show route damping history on page 3001 show route damping history detail on page 3001
Output Fields	Table 269 on page 2998 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.

Table 269: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0 .	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holdddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
destination-prefix (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
[protocol, preference]	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive

Table 269: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive

Table 269: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

  6-Resolve tree 2 7-Resolve tree 3
    AS path: 65490 65520 65525 65525 65525 65525 I ()
    Communities: 65501:390 65501:2000 65501:3000 65504:701
    Localpref: 100
    Router ID: 172.23.2.129
    Merit (last update/now): 1934/1790
    damping-parameters: damping-high

```

```

Last update:      00:03:28 First update:      00:06:40
Flaps: 2

```

show route damping history

```

user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0

```

show route damping history detail

```

user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:      00:01:05 First update:      00:01:05
        Flaps: 1

```

show route detail

List of Syntax	Syntax on page 3002 Syntax (EX Series Switches) on page 3002
Syntax	show route detail <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route detail <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display detailed information about the active entries in the routing tables.
Options	none —Display all active entries in the routing table on all systems. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route detail on page 3010 show route detail (with BGP Multipath) on page 3016
Output Fields	Table 270 on page 3002 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.

Table 270: show route detail Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 270: show route detail Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • --A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 271 on page 3006 .

Table 270: show route detail Output Fields (*continued*)

Field Name	Field Description
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 272 on page 3008 .
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metricn	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Table 270: show route detail Output Fields (*continued*)

Field Name	Field Description
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 273 on page 3010 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).

Table 270: show route detail Output Fields (*continued*)

Field Name	Field Description
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

[Table 271 on page 3006](#) describes all possible values for the **Next-hop Types** output field.

Table 271: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.

Table 271: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 272 on page 3008 describes all possible values for the **State** output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 272: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGp path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).

Table 272: State Output Field Values (*continued*)

Value	Description
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 273 on page 3010 describes the possible values for the **Communities** output field.

Table 273: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0 . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
<i>bandwidth: local AS number:link-bandwidth-number</i>	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
<i>domain-id</i>	Unique configurable number that identifies the OSPF domain.
<i>domain-id-vendor</i>	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535 .
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7 . Setting the least significant bit in the field indicates that the route carries a type 2 metric.
<i>origin</i>	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
<i>route-type-vendor</i>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is <i>area-number:ospf-route-type:options</i> .
<i>rte-type</i>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is <i>area-number:ospf-route-type:options</i> .
<i>target</i>	Defines which VPN the route participates in; target has the format <i>32-bit IP address:16-bit number</i> . For example, 10.19.0.0:100.
<i>unknown IANA</i>	Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized.
<i>unknown OSPF vendor community</i>	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail
```

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```

10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Local AS: 69
    Age: 1:30:17 Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

```

```
224.0.0.2/32 (1 entry, 1 announced)
  *PIM    Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:45
          Task: PIM Recv
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:43
          Task: IGMP
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
```



```

        State: <Active Int>
        Local AS: 69
        Age: 1:31:44
        Task: IF
        AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:31:45 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:29:30
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:29:30 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

```

```
abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:31:44
```

```

Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:25:49 Metric2: 1
    AIGP 210
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]

```

```
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: 12 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

show route detail (with BGP Multipath)

```
user@host> show route detail
```

```
10.1.1.8/30 (2 entries, 1 announced)
*BGP Preference: 170/-101
Next hop type: Router, Next hop index: 262142
Address: 0x901a010
Next-hop reference count: 2
Source: 10.1.1.2
Next hop: 10.1.1.2 via ge-0/3/0.1, selected
Next hop: 10.1.1.6 via ge-0/3/0.5
State: <Active Ext>
Local AS: 1 Peer AS: 2
Age: 5:04:43
Task: BGP_2.10.1.1.2+59955
Announcement bits (1): 0-KRT
AS path: 2 I
Accepted Multipath
Localpref: 100
Router ID: 1.1.1.2
BGP Preference: 170/-101
Next hop type: Router, Next hop index: 678
Address: 0x8f97520
Next-hop reference count: 9
Source: 10.1.1.6
Next hop: 10.1.1.6 via ge-0/3/0.5, selected
State: <NotBest Ext>
Inactive reason: Not Best in its group - Active preferred
Local AS: 1 Peer AS: 2
Age: 5:04:43
Task: BGP_2.10.1.1.6+58198
AS path: 2 I
Accepted MultipathContrib
Localpref: 100
Router ID: 1.1.1.3
```

show route exact

List of Syntax	Syntax on page 3017 Syntax (EX Series Switches) on page 3017
Syntax	show route exact <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route exact <i>destination-prefix</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display only the routes that exactly match the specified address or range of addresses.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . destination-prefix —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route exact on page 3017 show route exact detail on page 3017 show route exact extensive on page 3018 show route exact terse on page 3018
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```

user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0

```

show route exact detail

```

user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)

```

```
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2d 3:30:26
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 207.17.136.0/24  S  5                >192.168.71.254
```

show route export

List of Syntax	Syntax on page 3019 Syntax (EX Series Switches) on page 3019
Syntax	<pre>show route export <brief detail> <instance <instance-name> routing-table-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches)	<pre>show route export <brief detail> <instance <instance-name> routing-table-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.
Options	<p>none—(Same as brief.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route export inet command).</p>
Required Privilege Level	view
List of Sample Output	show route export on page 3020 show route export detail on page 3020 show route export instance detail on page 3020
Output Fields	Table 274 on page 3019 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear.

Table 274: show route export Output Fields

Field Name	Field Description	Level of Output
Table or table-name	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	brief none

Table 274: show route export Output Fields (*continued*)

Field Name	Field Description	Level of Output
Export	Whether the table is currently exporting routes to other tables: Y or N (Yes or No).	brief none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	detail
Flags	(instance keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> config auto-policy—The policy was deduced from the configured IGP export policies. cleanup—Configuration information for this instance is no longer valid. config—The instance was explicitly configured. 	detail
Options	(instance keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	detail
Import policy	(instance keyword only) Policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf .	detail
Instance	(instance keyword only) Name of the routing instance.	detail
Type	(instance keyword only) Type of routing instance: forwarding , non-forwarding , or vrf .	detail

Sample Output

show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

show route export detail

```

user@host> show route export detail
inet.0                                Routes:      0
black.inet.0                          Routes:      3
  Import: [ inet.0 ]
red.inet.0                            Routes:      4
  Import: [ inet.0 ]

```

show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [ (ospf-master-from-red || isis-master-from-black) ]

```


Instance: black
Instance: red

Type: non-forwarding
Type: non-forwarding

show route extensive

List of Syntax	Syntax on page 3022 Syntax (EX Series Switches) on page 3022
Syntax	<pre>show route extensive <destination-prefix> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches)	<pre>show route extensive <destination-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display extensive information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route extensive on page 3027 show route extensive (Access Route) on page 3033 show route extensive (Route Reflector) on page 3034 show route extensive (FRR and LFA) on page 3034 show route extensive (FRR and LFA) on page 3035
Output Fields	<p>Table 275 on page 3022 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.</p>

Table 275: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> active (routes that are active). holddown (routes that are in the pending state before being declared inactive). hidden (routes that are not used because of a routing policy).

Table 275: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
TSI	Protocol header information.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.

Table 275: show route extensive Output Fields (*continued*)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
label-operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.

Table 275: show route extensive Output Fields (*continued*)

Field Name	Field Description
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGP path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.

Table 275: show route extensive Output Fields (*continued*)

Field Name	Field Description
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. Recorded—The AS path is recorded by the sample process (sampled). ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.

Table 275: show route extensive Output Fields (*continued*)

Field Name	Field Description
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:34:06
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2

```

```

                                Next hop: via so-0/3/0.0, selected
                                State: <Active Int>
                                Local AS:    69
                                Age: 1:32:40
                                Task: IF
                                Announcement bits (1): 3-Resolve tree 2
                                AS path: I
OSPF Preference: 10
                                Next-hop reference count: 1
                                Next hop: via so-0/3/0.0, selected
                                State: <Int>
                                Inactive reason: Route Preference
                                Local AS:    69
                                Age: 1:32:40    Metric: 1
                                Area: 0.0.0.0
                                Task: OSPF
                                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
                                Next hop type: Local
                                Next-hop reference count: 7
                                Interface: so-0/3/0.0
                                State: <Active NoReadvrt Int>
                                Local AS:    69
                                Age: 1:32:43
                                Task: IF
                                Announcement bits (1): 3-Resolve tree 2
                                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
  *OSPF Preference: 10
                                Next-hop reference count: 9
                                Next hop: via so-0/3/0.0
                                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                                State: <Active Int>
                                Local AS:    69
                                Age: 1:32:19    Metric: 2
                                Area: 0.0.0.0
                                Task: OSPF
                                Announcement bits (2): 0-KRT 3-Resolve tree 2
                                AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
  *PIM Preference: 0
                                Next-hop reference count: 18
                                State: <Active NoReadvrt Int>
                                Local AS:    69
                                Age: 1:34:08
                                Task: PIM Recv
                                Announcement bits (2): 0-KRT 3-Resolve tree 2
                                AS path: I
```



```

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
           Next-hop reference count: 18
           State: <Active NoReadvrt Int>
           Local AS:    69
           Age: 1:34:06
           Task: IGMP
           Announcement bits (2): 0-KRT 3-Resolve tree 2
           AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP   Preference: 7
           Next-hop reference count: 6
           Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
           Label-switched-path green-r1-r3
           Label operation: Push 100096
           State: <Active Int>
           Local AS:    69
           Age: 1:28:12   Metric: 2
           Task: RSVP
           Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
           AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP   Preference: 7
           Next-hop reference count: 6
           Next hop: via so-0/3/0.0 weight 0x1, selected
           Label-switched-path green-r1-r2
           State: <Active Int>
           Local AS:    69
           Age: 1:28:12   Metric: 1
           Task: RSVP
           Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
           AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
    *Direct Preference: 0
           Next hop type: Interface
           Next-hop reference count: 1
           Next hop: via lo0.0, selected
           State: <Active Int>
           Local AS:    69
           Age: 1:34:07
           Task: IF
           AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

```

```

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:34:08 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:31:53
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:31:53 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Indirect next hops: 1
            Protocol next hop: 10.255.70.103 Metric: 2

```

```

Push 800012
Indirect next hop: 87272e4 1048574
Indirect path forwarding next hops: 1
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
10.255.70.103/32 Originating RIB: inet.3
    Metric: 2                      Node path count: 1
    Forwarding nexthops: 1
    Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
    *MLD Preference: 0

```

```

    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:34:07
    Task: IF
    AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
```

```

*L2VPN Preference: 170/-101
  Next-hop reference count: 5
  Protocol next hop: 10.255.71.52
  Indirect next hop: 0 -
  State: <Active Int Ext>
  Age: 1:34:03 Metric2: 1
  Task: green-l2vpn
  Announcement bits (1): 1-BGP.0.0.0+179
  AS path: I
  Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
  Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  55.0.0.0/24 -> {Push 300112}
    *BGP Preference: 170/-101
      Next hop type: Router
      Address: 0x925c208
      Next-hop reference count: 2
      Source: 10.0.0.9
      Next hop: 10.0.0.9 via ge-1/2/0.15, selected
      Label operation: Push 300112
      Label TTL action: prop-ttl
      State: <Active Ext>
      Local AS: 7019 Peer AS: 13979
      Age: 1w0d 23:06:56
      AIGP: 25
      Task: BGP_13979.10.0.0.9+56732
      Announcement bits (1): 0-KRT
      AS path: 13979 7018 I
      Accepted
      Route Label: 300112
      Localpref: 100
      Router ID: 10.9.9.1

```

show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
```

```

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
    *Access Preference: 13
        Next-hop reference count: 78472
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 12
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (2): 0-KRT 1-OSPFv2
        AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
    *BGP    Preference: 170/-101
        Source: 192.168.4.214
        Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
        State: <Active Int Ext>
        Local AS: 10458 Peer AS: 10458
        Age: 3:09      Metric: 0      Metric2: 0
        Task: BGP_10458.192.168.4.214+1033
        Announcement bits (2): 0-KRT 4-Resolve inet.0
        AS path: 3944 7777 I <Originator>
        Cluster list: 1.1.1.1
        Originator ID: 10.255.245.88
        Communities: 7777:7777
        Localpref: 100
        Router ID: 4.4.4.4
        Indirect next hops: 1
            Protocol next hop: 207.17.136.192 Metric: 0
            Indirect next hop: 84ac908 40
            Indirect path forwarding next hops: 0
            Next hop type: Discard

```

show route extensive (FRR and LFA)

```

user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP    Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792

```

```

Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

show route extensive (FRR and LFA)

```

user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
State: FlashAll
TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
*RSVP Preference: 7/1
Next hop type: Router, Next hop index: 1048574
Address: 0xbbbc010
Next-hop reference count: 5
Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
Label-switched-path europa-d-to-europa-e
Label operation: Push 299776
Label TTL action: prop-ttl
Session Id: 0x201
Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
Label-switched-path europa-d-to-europa-e
Label operation: Push 299792
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0

```

Task: OSPF
AS path: I

show route flow validation

List of Syntax [Syntax on page 3037](#)
[Syntax \(EX Series Switches\) on page 3037](#)

Syntax show route flow validation
 <brief | detail>
 <ip-prefix>
 <table table-name>
 <logical-system (all | logical-system-name)>

Syntax (EX Series Switches) show route flow validation
 <brief | detail>
 <ip-prefix>
 <table table-name>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display flow route information.

Options none—Display flow route information.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

ip-prefix—(Optional) IP address for the flow route.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table table-name—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route flow validation inet** command).

Required Privilege Level view

List of Sample Output [show route flow validation on page 3038](#)

Output Fields [Table 276 on page 3037](#) lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

Table 276: show route flow validation Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).	All levels
<i>prefix</i>	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels

Table 276: show route flow validation Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```

user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent

```

show route forwarding-table

List of Syntax	Syntax on page 3039 Syntax (MX Series Routers) on page 3039 Syntax (TX Matrix and TX Matrix Plus Routers) on page 3039
Syntax	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (MX Series Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p>

Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | bridge-domain-name)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc interface-name—(Optional) Display route entries for the specified circuit cross-connect interface.

destination destination-prefix—(Optional) Destination prefix.

family family—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name interface-name—(Optional) Display routing table entries for the specified interface.

label name—(Optional) Display route entries for the specified label.

lcc number—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (*default* | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

vlan (*all* | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output [show route forwarding-table on page 3044](#)
[show route forwarding-table detail on page 3045](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 3045](#)
[show route forwarding-table extensive on page 3046](#)
[show route forwarding-table extensive \(RPF\) on page 3047](#)
[show route forwarding-table family mpls on page 3048](#)
[show route forwarding-table family vpls on page 3048](#)
[show route forwarding-table family vpls extensive on page 3048](#)
[show route forwarding-table table default on page 3050](#)
[show route forwarding-table table](#)
[logical-system-name/routing-instance-name on page 3050](#)
[show route forwarding-table vpn on page 3051](#)

Output Fields [Table 277 on page 3042](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might

be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 277: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table <i>logical-system-name/routing-instance-name</i> option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive
Flags	Route type flags: <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface interface-number—Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • rt nh decoupled—Route has been decoupled from the next hop to the destination. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. 	extensive
Next hop	IP address of the next hop to the destination.	detail extensive

Table 277: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd) —Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   46    4
0.0.0.0/32       perm  0                dscd   44    1
1.1.1.0/24       ifdn  0                rslv   608   1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0         recv   606   1 ge-2/0/1.0
1.1.1.1/32       user  0                rjct   46    4
1.1.1.1/32       intf  0 1.1.1.1         locl   607   2
1.1.1.1/32       iddn  0 1.1.1.1         locl   607   2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff bcst   605   1 ge-2/0/1.0
10.0.0.0/24      intf  0                rslv   616   1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0         recv   614   1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1         locl   615   2
10.0.0.1/32      dest  0 10.0.0.1         locl   615   2
10.0.0.255/32    dest  0 10.0.0.255       bcst   613   1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                rslv   612   1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0         recv   610   1 ge-2/0/1.0
10.1.1.1/32      user  0                rjct   46    4
10.1.1.1/32      intf  0 10.1.1.1         locl   611   2
10.1.1.1/32      iddn  0 10.1.1.1         locl   611   2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff bcst   609   1 ge-2/0/1.0
10.206.0.0/16    user  0 10.209.63.254    ucst   419   20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0  ucst   419   20 fxp0.0
10.209.0.0/18    intf  0                rslv   418   1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0         recv   416   1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131     locl   417   2
10.209.2.131/32  dest  0 10.209.2.131     locl   417   2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2 ucst   435   1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca ucst   434   1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0  ucst   419   20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255    bcst   415   1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254    ucst   419   20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   27    1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                locl   28    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   6    1
ff00::/8         perm  0                mdsc   4    1
ff02::1/128      perm  0 ff02::1          mcst   3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  16    1
100004(top)fe-0/0/1.0

```


show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b ucst  132   4 fxp0.0
default          perm   0                               rjct  14    1
10.1.1.0/24      intf   0 ff.3.0.21         ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0          recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1          locl  321   1
10.1.1.255/32    dest   0 10.1.1.255        bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21         ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0        recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1        locl  325   1
10.21.21.255/32  dest   0 10.21.21.255      bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1         locl  320   1
172.17.28.19/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0
172.17.28.44/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0          recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4          locl  135   2
10.0.0.4/32      dest   0 10.0.0.4          locl  135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  22    1
ff00::/8         perm   0                               mdsc  21    1
ff02::1/128      perm   0 ff02::1          mcst  17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

Flags: sent to PFE		
Next-hop type: unicast	Index: 262143	Reference: 1
Nexthop: 4.4.4.4		
Next-hop type: unicast	Index: 335	Reference: 2
Next-hop interface: so-1/1/0.0	Weight: 22	Balance: 3
Nexthop: 145.12.1.2		
Next-hop type: unicast	Index: 337	Reference: 2
Next-hop interface: so-0/1/2.0	Weight: 33	Balance: 33

show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
```

```
Routing table: inet [Index 0]
```

```
Internet:
```

```
Destination: default
```

```
Route type: user
```

```
Route reference: 2
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Nexthop: 0:90:69:8e:b1:1b
```

```
Next-hop type: unicast
```

```
Index: 132      Reference: 4
```

```
Next-hop interface: fxp0.0
```

```
Destination: default
```

```
Route type: permanent
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: none
```

```
Next-hop type: reject
```

```
Index: 14      Reference: 1
```

```
Destination: 127.0.0.1/32
```

```
Route type: interface
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Nexthop: 127.0.0.1
```

```
Next-hop type: local
```

```
Index: 320      Reference: 1
```

```
...
```

```
Routing table: private1__inet [Index 1]
```

```
Internet:
```

```
Destination: default
```

```
Route type: permanent
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Next-hop type: reject
```

```
Index: 46      Reference: 1
```

```
Destination: 10.0.0.0/8
```

```
Route type: interface
```

```
Route reference: 0
```

```
Route interface-index: 3
```

```
Flags: sent to PFE
```

```
Next-hop type: resolve
```

```
Next-hop interface: fxp1.0
```

```
Index: 136      Reference: 1
```

```
...
```

```
Routing table: iso [Index 0]
```

```
ISO:
```

```
Destination: default
```

```
Route type: permanent
```

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6          swap 100001      fe-1/1/0.0
800002           user  0                  Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                  user  0                  indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                  dymn  0                  indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                  dymn  0                  ucst  354      2 fe-0/1/0.0

```

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood          Index: 289      Reference: 1
Next-hop type: unicast       Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast       Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Route interface-index: 72
Route interface-index: 0

```

```

Next-hop type: discard                                Index: 341      Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0                                    Route interface-index: 69
Flags: sent to PFE
Next-hop type: flood                                Index: 293      Reference: 1
Next-hop type: indirect                              Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect                              Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast                              Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0                                    Route interface-index: 70
Flags: sent to PFE
Next-hop type: flood                                Index: 292      Reference: 1
Next-hop type: indirect                              Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect                              Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast                              Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0                                    Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast                              Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:          6640   Byte count:      675786
Route used as source
  Packet count:          6894   Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0                                    Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast                              Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:          96     Byte count:      8079
Route used as source:
  Packet count:          296    Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0                                    Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect                              Index: 301      Reference: 5
Next hop: 10.31.3.2

```

Next-hop type: Push 800000
 Next-hop interface: fe-0/1/1.0

show route forwarding-table table default

```
user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13                ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0                          rslv  688  1 fe-0/1/3.0
10.0.60.12/32    dest  0 10.0.60.12                recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22          ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14                locl  687  2
10.0.60.14/32    dest  0 10.0.60.14                locl  687  2
10.0.60.15/32    dest  0 10.0.60.15                bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13                ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21                ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0                recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0                          rjct  36  2
10.0.80.2/32     intf  0 10.0.80.2                locl  675  1
10.0.80.3/32     dest  0 10.0.80.3                bcst  677  1 so-0/0/1.0
10.0.90.12/30    intf  0                          rslv  684  1 fe-0/1/0.0
10.0.90.12/32    dest  0 10.0.90.12                recv  682  1 fe-0/1/0.0
10.0.90.14/32    intf  0 10.0.90.14                locl  683  2
10.0.90.14/32    dest  0 10.0.90.14                locl  683  2
10.0.90.15/32    dest  0 10.0.90.15                bcst  681  1 fe-0/1/0.0
10.5.0.0/16      user  0 192.168.187.126          ucst  324  15 fxp0.0
10.10.0.0/16     user  0 192.168.187.126          ucst  324  15 fxp0.0
10.13.10.0/23    user  0 192.168.187.126          ucst  324  15 fxp0.0
10.84.0.0/16     user  0 192.168.187.126          ucst  324  15 fxp0.0
10.150.0.0/16    user  0 192.168.187.126          ucst  324  15 fxp0.0
10.157.64.0/19   user  0 192.168.187.126          ucst  324  15 fxp0.0
10.209.0.0/16    user  0 192.168.187.126          ucst  324  15 fxp0.0
...

Routing table: default.iso
ISO:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm  0                          rjct  60  1

Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm  0                          rjct  44  1
::/128           perm  0                          dscd  42  1
ff00::/8         perm  0                          mdsc  43  1
ff02::1/128      perm  0 ff02::1                mcst  39  1

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm  0                          dscd  50  1
```

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	loc1	770	2	
2.0.2.1/32	dest	0	2.0.2.1	loc1	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	

Logical system: R4

Routing table: vpn-red.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	708	1	
::/128	perm	0		dscd	706	1	
ff00::/8	perm	0		mdsc	707	1	
ff02::1/128	perm	0	ff02::1	mcst	704	1	

Logical system: R4

Routing table: vpn-red.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	638		

show route forwarding-table vpn

user@host> show route forwarding-table vpn VPN-A

Routing table:: VPN-A.inet

Internet:

Destination	Type	RtRef	Nexthop	Type	Index	NhRef	Netif
default	perm	0		rjct	4	4	
10.39.10.20/30	intf	0	ff.3.0.21	ucst	40	1	
so-0/0/0.0							
10.39.10.21/32	intf	0	10.39.10.21	loc1	36	1	
10.255.14.172/32	user	0		ucst	69	2	
so-0/0/0.0							
10.255.14.175/32	user	0		indr	81	3	
				Push	100004	Push	
100004(top) so-1/0/0.0							
224.0.0.0/4	perm	2		mdsc	5	3	
224.0.0.1/32	perm	0	224.0.0.1	mcst	1	8	
224.0.0.5/32	user	1	224.0.0.5	mcst	1	8	
255.255.255.255/32	perm	0		bcst	2	3	

show route inactive-path

List of Syntax	Syntax on page 3052 Syntax (EX Series Switches) on page 3052
Syntax	<code>show route inactive-path</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route inactive-path</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.
Options	none —Display all inactive routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route inactive-path on page 3052 show route inactive-path detail on page 3053 show route inactive-path extensive on page 3054 show route inactive-path terse on page 3054
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                    AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS: 1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)

```

```

BGP      Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1
        Next hop: 10.12.80.1 via ge-6/3/2.0, selected
        State: <Ext>
        Inactive reason: Route Preference
        Peer AS: 100
        Age: 4:39:13
        Task: BGP_100.10.12.80.1+179
        AS path: 100 I
        Localpref: 100
        Router ID: 10.0.0.0

```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 3053](#).

show route inactive-path terse

```

user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.100.12/30   0 10      1          >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.0.0.0/8        D  0          >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.80.0/30     B 170      100      >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-prefix

List of Syntax	Syntax on page 3055 Syntax (EX Series Switches) on page 3055
Syntax	<pre>show route inactive-prefix <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route inactive-prefix <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display inactive route destinations in each routing table.
Options	<p>none—Display all inactive route destination.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route inactive-prefix on page 3055 show route inactive-prefix detail on page 3055 show route inactive-prefix extensive on page 3056 show route inactive-prefix terse on page 3056
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
```

```
127.0.0.1/32 (1 entry, 0 announced)
  Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Hidden Martian Int>
    Age: 4:51
    Task: IF
    AS path: I00:04:54
      > via lo0.0
```

`show route inactive-prefix extensive`

The output for the `show route inactive-prefix extensive` command is identical to that of the `show route inactive-path detail` command. For sample output, see [show route inactive-prefix detail on page 3055](#).

`show route inactive-prefix terse`

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
127.0.0.1/32	D 0			>lo0.0	

show route instance

List of Syntax	Syntax on page 3057 Syntax (EX Series Switches and QFX Series) on page 3057
Syntax	<pre>show route instance <brief detail summary> <instance-name> <logical-system (all logical-system-name)> <operational></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show route instance <brief detail summary> <instance-name> <operational></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show route instance cust1 command).</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show route instance on page 3058 show route instance detail (Graceful Restart Complete) on page 3059 show route instance detail (Graceful Restart Incomplete) on page 3060 show route instance detail (VPLS Routing Instance) on page 3062 show route instance operational on page 3063 show route instance summary on page 3063
Output Fields	<p>Table 177 on page 1855 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.</p>

Table 278: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , virtual-router , or vrf .	All levels
State	State of the routing instance: active or inactive .	brief detail none
Interfaces	Name of interfaces belonging to this routing instance.	brief detail none
Restart State	Status of graceful restart for this instance: Pending or Complete .	detail
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300 .	detail
Tables	Tables (and number of routes) associated with this routing instance.	brief detail none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	detail
Vrf-import	VPN routing and forwarding instance import policy name.	detail
Vrf-export	VPN routing and forwarding instance export policy name.	detail
Vrf-import-target	VPN routing and forwarding instance import target community name.	detail
Vrf-export-target	VPN routing and forwarding instance export target community name.	detail
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: high , medium , or low . The default is low .	detail
Restart State	Restart state: <ul style="list-style-type: none"> Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. Complete—All protocols have restarted for this routing table. 	detail
Primary rib	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@host> show route instance
Instance           Type
      Primary RIB
master             forwarding
Active/holddown/hidden

```

```

inet.0                                16/0/1
iso.0                                  1/0/0
mpls.0                                0/0/0
inet6.0                               2/0/0
l2circuit.0                           0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0           12/0/0
__juniper_private1__.inet6.0          1/0/0

```

show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                 : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf             State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf             State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0         : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
      BGP-L.mpls.0         : 3 routes (3 active, 0 holddown, 0 hidden)
      Restart Complete
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn           State: Active
    Restart State: Complete Path selection timeout: 300

```

```
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0             : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
```



```

master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Pending Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf              State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf              State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0          : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn            State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf              State: Active
  Restart State: Pending Path selection timeout: 300

```

```
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0          : 5 routes (4 active, 1 holddown, 0 hidden)
Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0       : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0        : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0     : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls           State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Fast-reroute-priority: high
```

Tables:
 test-vpls.12vpn.0 : 3 routes (3 active, 0 holddown, 0 hidden)

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:
```

```
master
default
```

show route instance summary

```
user@host> show route instance summary
```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding		
		inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		13vpn.0	0/0/0
		inet6.0	2/0/0
		12vpn.0	0/0/0
		12circuit.0	0/0/0
BGP-INET	vrf		
		BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
BGP-L	vrf		
		BGP-L.inet.0	5/0/0
		BGP-L.iso.0	0/0/0
		BGP-L.mpls.0	4/0/0
		BGP-L.inet6.0	0/0/0
L2VPN	12vpn		
		L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.12vpn.0	2/0/0
LDP	vrf		
		LDP.inet.0	4/0/0
		LDP.iso.0	0/0/0
		LDP.mpls.0	0/0/0
		LDP.inet6.0	0/0/0
		LDP.12circuit.0	0/0/0
OSPF	vrf		
		OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf		
		RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf		
		STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0

show route label

List of Syntax	Syntax on page 3064 Syntax (EX Series Switches) on page 3064
Syntax	<code>show route label <i>label</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route label <i>label</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.
Options	<i>label</i> —Value of the MPLS label. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label on page 3064 show route label detail on page 3064 show route label extensive on page 3065 show route label terse on page 3065
Output Fields	For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                 > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail
```

```

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS:      1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2

```

show route label extensive

The output for the `show route label extensive` command is identical to that of the `show route label detail` command. For sample output, see [show route label detail on page 3064](#).

show route label terse

```
user@host> show route label 100016 terse
```

```

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	100016	V	170			>10.12.80.1	

show route label-switched-path

List of Syntax	Syntax on page 3066 Syntax (EX Series Switches) on page 3066
Syntax	show route label-switched-path <i>path-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route label-switched-path <i>path-name</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes used in an MPLS label-switched path (LSP).
Options	brief detail extensive terse —(Optional) Display the specified level of output. <i>path-name</i> —LSP tunnel name. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label-switched-path on page 3066
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```

user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

```

```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                   > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route martians

List of Syntax	Syntax on page 3068 Syntax (EX Series Switches) on page 3068
Syntax	<pre>show route martians <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i>></pre>
Syntax (EX Series Switches)	<pre>show route martians <table <i>routing-table-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display the martian (invalid and ignored) entries associated with each routing table.
Options	<p>none—Display standard information about route martians for all routing tables.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route martians table inet command).</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Martian Addresses
List of Sample Output	show route martians on page 3069
Output Fields	<p>Table 279 on page 3068 lists the output fields for the show route martians command. Output fields are listed in the approximate order in which they appear</p>

Table 279: show route martians Output Fields

Field Name	Field Description
<i>table-name</i>	Name of the route table in which the route martians reside.
<i>destination-prefix</i>	Route destination.
<i>match value</i>	Route match parameter.
<i>status</i>	Status of the route: allowed or disallowed .

Sample Output

show route martians

```

user@host> show route martians

inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...

```

show route next-hop

List of Syntax	Syntax on page 3070 Syntax (EX Series Switches) on page 3070
Syntax	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that are being sent to the specified next-hop address.
Options	<code>brief detail extensive terse</code> —(Optional) Display the specified level of output. <code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>next-hop</code> —Next-hop address.
Required Privilege Level	view
List of Sample Output	show route next-hop on page 3070 show route next-hop detail on page 3071 show route next-hop extensive on page 3073 show route next-hop terse on page 3074
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25
```

```

> to 192.168.71.254 via fxp0.0
192.168.102.0/23 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.0/24 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2

```

```
AS path: I

192.168.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16     S  5          5          >192.168.71.254
* 10.209.0.0/16    S  5          5          >192.168.71.254
* 172.16.0.0/12    S  5          5          >192.168.71.254

```

```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32 S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

List of Syntax	Syntax on page 3076 Syntax (EX Series Switches) on page 3076
Syntax	<code>show route no-community</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route no-community</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are not associated with any community.
Options	none —(Same as brief) Display the route entries in each routing table that are not associated with any community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route no-community on page 3076 show route no-community detail on page 3077 show route no-community extensive on page 3077 show route no-community terse on page 3078
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
> via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
> to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2
```



```

> to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 * [OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0
> via so-0/3/2.0
10.255.71.241/32 * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32 * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24      * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

show route no-community detail

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

show route no-community extensive

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

```

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

show route no-community terse

```
user@host> show route no-community terse
```

```

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.52/32	D	0			>100.0	
*	10.255.71.63/32	0	10	1		>35.1.1.2	
*	10.255.71.64/32	0	10	2		>35.1.1.2	
*	10.255.71.240/32	0	10	2		so-0/1/2.0	
						>so-0/3/2.0	
*	10.255.71.241/32	0	10	1		>so-0/1/2.0	
*	10.255.71.242/32	0	10	1		>so-0/3/2.0	
*	12.1.1.0/24	0	10	2		>so-0/3/2.0	
*	14.1.1.0/24	0	10	3		>35.1.1.2	
						so-0/1/2.0	
						so-0/3/2.0	
*	16.1.1.0/24	0	10	2		>so-0/1/2.0	
...							

show route protocol

List of Syntax	Syntax on page 3079 Syntax (EX Series Switches) on page 3079
Syntax	<pre>show route protocol <i>protocol</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route protocol <i>protocol</i> <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>flow option introduced in Junos OS Release 10.0.</p> <p>flow option introduced in Junos OS Release 10.0 for EX Series switches.</p>
Description	Display the route entries in the routing table that were learned from a particular protocol.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>protocol</i>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> • access—Access route for use by DHCP application • access-internal—Access-internal route for use by DHCP application • aggregate—Locally generated aggregate route • arp—Route learned through the Address Resolution Protocol • atmvpn—Asynchronous Transfer Mode virtual private network • bgp—Border Gateway Protocol • ccc—Circuit cross-connect • direct—Directly connected route • dvmrp—Distance Vector Multicast Routing Protocol • esis—End System-to-Intermediate System • flow—Locally defined flow-specification route • frr—Precomputed protection route or backup route used when a link goes down • isis—Intermediate System-to-Intermediate System • ldp—Label Distribution Protocol • l2circuit—Layer 2 circuit

- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
List of Sample Output	show route protocol access on page 3081 show route protocol access-internal extensive on page 3081 show route protocol arp on page 3081 show route protocol bgp on page 3082 show route protocol bgp detail on page 3082 show route protocol bgp extensive on page 3082 show route protocol bgp terse on page 3083 show route protocol direct on page 3083 show route protocol frr on page 3084 show route protocol l2circuit detail on page 3084 show route protocol l2vpn extensive on page 3085 show route protocol ldp on page 3086 show route protocol ldp extensive on page 3086 show route protocol ospf (Layer 3 VPN) on page 3087 show route protocol ospf detail on page 3088 show route protocol rip on page 3088 show route protocol rip detail on page 3088

[show route protocol ripng table inet6 on page 3089](#)

[show route protocol static detail on page 3089](#)

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
```

```

20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.11/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.12/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.13/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
...

```

show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0

```

show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24    (1 entry, 1 announced)
    *BGP          Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006436
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:42:42 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
    AS path: 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
    Accepted
    Localpref: 100
    Router ID: 207.17.136.192

```

show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
    *BGP          Preference: 170/-101
                   Next hop type: Indirect

```

```

Next-hop reference count: 1006502
Source: 192.168.69.71
Next hop type: Router, Next hop index: 324
Next hop: 192.168.167.254 via fxp0.0, selected
Protocol next hop: 192.168.69.71
Indirect next hop: 8e166c0 342
State: <Active Ext>
Local AS: 69 Peer AS: 10458
Age: 6d 10:44:45 Metric2: 0
Task: BGP_10458.192.168.69.71+179
Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
  192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
192.168.64.0/21	B 170	100		>100.1.3.2	10023 21 I

show route protocol direct

```
user@host> show route protocol direct
```

```
inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

8.8.8.0/24      *[Direct/0] 17w0d 10:31:49
                 > via fe-1/3/1.0
10.255.165.1/32 *[Direct/0] 25w4d 04:13:18
                 > via lo0.0
30.30.30.0/24   *[Direct/0] 17w0d 23:06:26
                 > via fe-1/3/2.0
192.168.164.0/22 *[Direct/0] 25w4d 04:13:20
                 > via fxp0.0

```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
   *[Direct/0] 25w4d 04:13:21
   > via lo0.0

```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop          Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
```



```

*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000, Push 100000(top)[0] Offset: -4
  Protocol next hop: 10.245.255.63
  Push 100000 Offset: -4
    Indirect next hop: 86af0c0 298
  State: <Active Int>
  Local AS: 99
  Age: 9:52
  Task: Common L2 VC
  Announcement bits (2): 0-KRT 1-Common L2 VC
  AS path: I

```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)

```

*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000[0]
  Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
  State: <Active Int>
  Local AS: 99
  Age: 10:21
  Task: l2 circuit
  Announcement bits (1): 0-LDP
  AS path: I
  VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive
```

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)

TSI:

KRT in-kernel 800001 /36 -> {so-0/0/0.0}

```

*L2VPN Preference: 7
  Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
  Label operation: Pop Offset: 4
  State: <Active Int>
  Local AS: 69
  Age: 7:48
  Task: Common L2 VC
  Announcement bits (1): 0-KRT
  AS path: I

```

so-0/0/0.0 (1 entry, 1 announced)

TSI:

KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}

```

*L2VPN Preference: 7
  Next hop: via so-0/0/1.0, selected
  Label operation: Push 800000 Offset: -4
  Protocol next hop: 10.255.14.220

```

```

Push 800000 Offset: -4
Indirect next hop: 85142a0 288
State: <Active Int>
Local AS: 69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP Preference: 9
    Next-hop reference count: 3
    Next hop: via t1-4/0/0.0, selected
    Label operation: Push 100000
    State: <Active Int>
    Local AS: 65500
    Age: 1d 23:03:58 Metric: 1
    Task: LDP
    Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
    AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP Preference: 9
    Next-hop reference count: 3
    Next hop: via t1-4/0/0.0, selected
    State: <Active Int>
    Local AS: 65500
    Age: 1d 23:03:58 Metric: 1
    Task: LDP

```

```

Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Pop
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Swap 100000
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.16.1/32

```

show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32 *[OSPF/10] 00:05:18, metric 4

```

```

> via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
> via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30    [OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
224.0.0.5/32    *[OSPF/10] 20:26:20, metric 1

```

show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
  OSPF   Preference: 10
         Nexthop: via so-0/2/2.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Age: 6:25      Metric: 1
         Area: 0.0.0.0
         Task: VPN-AB-OSPF
         AS path: I
         Communities: Route-Type:0.0.0.0:1:0

...

```

show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
> to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     *[RIP/100] 00:03:59, metric 1

```

show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
  *RIP   Preference: 100
         Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
         State: <Active Int>
         Age: 20:25:02  Metric: 2
         Task: VPN-AB-RIPv2
         Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
         AS path: I
         Route learned from 10.39.1.22 expires in 96 seconds

```

show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

show route protocol static detail

```

user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified

```

Task: RT
Announcement bits (1): 0-KRT
AS path: I

show route range

List of Syntax	Syntax on page 3091 Syntax (EX Series Switches) on page 3091
Syntax	<pre>show route range <brief detail extensive terse> <destination-prefix> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches)	<pre>show route range <brief detail extensive terse> <destination-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display routing table entries using a prefix range.
Options	<p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—(Optional) Destination and prefix mask for the range.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route range on page 3091 show route range destination-prefix on page 3092 show route range detail on page 3092 show route range extensive on page 3093 show route range terse on page 3094
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
```

```
10.255.71.14/32      > to 192.168.71.254 via fxp0.0
                    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:30:01
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16      *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:31:14
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
```



```

State: <Active Int>
Age: 30:05
Task: IF
AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

...

```

show route range extensive

```

user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kerne1 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22

```

```

Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Age: 30:17
Task: RT
Announcement bits (1): 0-KRT
AS path: I

```

```
...
```

show route range terse

```
user@host> show route range terse
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.14/32	D	0			>lo0.0	
*	172.16.0.0/12	S	5			>192.168.71.254	
*	192.168.0.0/16	S	5			>192.168.71.254	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.71.14/32	L	0			Local	
*	192.168.102.0/23	S	5			>192.168.71.254	
*	207.17.136.0/24	S	5			>192.168.71.254	
*	207.17.136.192/32	S	5			>192.168.71.254	

```
__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.0.0.0/8	D	0			>fxp2.0	
		D	0			>fxp1.0	
*	10.0.0.4/32	L	0			Local	

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152						
*		D	0			>lo0.0	

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	abcd::10:255:71:14/128						
*		D	0			>lo0.0	
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.0	

```
__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.16385	

show route receive-protocol

List of Syntax	Syntax on page 3095 Syntax (EX Series Switches) on page 3095
Syntax	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol neighbor-address</i> —Protocol transmitting the route (bgp , dvmrp , msdp , pim , rip , or ripng) and address of the neighboring router from which the route entry was received.
Additional Information	The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.
Required Privilege Level	view
List of Sample Output	show route receive-protocol bgp on page 3098 show route receive-protocol bgp extensive on page 3098 show route receive-protocol bgp table extensive on page 3098 show route receive-protocol bgp logical-system extensive on page 3099 show route receive-protocol bgp detail (Layer 2 VPN) on page 3100 show route receive-protocol bgp extensive (Layer 2 VPN) on page 3100 show route receive-protocol bgp (Layer 3 VPN) on page 3101 show route receive-protocol bgp detail (Layer 3 VPN) on page 3101 show route receive-protocol bgp extensive (Layer 3 VPN) on page 3102
Output Fields	Table 280 on page 3095 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear.

Table 280: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0.	All levels

Table 280: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.	All levels
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) 	All levels
Prefix	Destination prefix.	none brief
MED	Multiple exit discriminator value included in the route.	none brief
<i>destination-prefix (entry, announced)</i>	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	detail extensive
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels

Table 280: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
AS path	<p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	detail extensive
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	detail extensive
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	detail extensive
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	detail extensive
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the Attrset attribute at the originating routing device.	detail extensive
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	detail extensive
control flags	Control flags: none or Site Down .	detail extensive
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	detail extensive

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0        100      I
10.22.2.0/24     10.255.245.215    0        100      I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```

user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300144
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300160
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

```

show route receive-protocol bgp detail (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp extensive (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100

```



```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

show route resolution

List of Syntax	Syntax on page 3103 Syntax (EX Series Switches) on page 3103
Syntax	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Syntax (EX Series Switches)	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.
Options	<p>none—Display standard information about all entries in the next-hop resolution database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>index <i>index</i>—(Optional) Show the index of the resolution tree.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>network/destination-prefix</i>—(Optional) Display database entries for the specified address.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled.</p> <p>unresolved—(Optional) Display routes that could not be resolved.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Resolution on PE Routers</i>
List of Sample Output	show route resolution detail on page 3104 show route resolution summary on page 3105 show route resolution unresolved on page 3105

Output Fields Table 281 on page 3104 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 281: show route resolution Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3 , this field indicates which routing table, inet.0 or inet.3 , provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```

Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1

```

show route resolution summary

```

user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3

```

show route resolution unresolved

```

user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
    Protocol Nexthop: 10.255.71.238 Push 800000
    Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
    Protocol Nexthop: 10.255.70.103 Push 800008
    Indirect nexthop: 0 -
Tree Index 2

```

show route snooping

Syntax	<code>show route snooping</code> <code><brief detail extensive terse></code> <code><all></code> <code><best address/prefix></code> <code><exact address></code> <code><range prefix-range></code> <code><summary></code> <code><table table-name></code>
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	show route snooping detail on page 3106
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
  *IGMP    Preference: 0
           Next hop type: MultiRecv
           Next-hop reference count: 4
           State: <Active NoReadvrt Int>
```

```

Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113

```

```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```



```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```

```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

show route source-gateway

List of Syntax	Syntax on page 3114 Syntax (EX Series Switches) on page 3114
Syntax	<code>show route source-gateway address</code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route source-gateway address</code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from a particular address. The Source field in the <code>show route detail</code> command output lists the source for each route, if known.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . address —IP address of the system. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route source-gateway on page 3114 show route source-gateway detail on page 3115 show route source-gateway extensive on page 3117
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```

```

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

```

show route source-gateway detail

```

user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

```

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-green-12vpn
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS,
          control flags:, mtu: 0
          Label-base: 800008, range: 8
          Localpref: 100
          Router ID: 10.255.70.103
          Primary Routing Table bgp.12vpn.0
```

red.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-red-12vpn
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
          control flags:Site-Down, mtu: 0
          Label-base: 800016, range: 8
          Localpref: 0
          Router ID: 10.255.70.103
          Primary Routing Table bgp.12vpn.0
```

bgp.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS, control
          flags:, mtu: 0
```



```

Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

```

```
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-1
```

```
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

show route summary

List of Syntax	Syntax on page 3120 Syntax (EX Series Switches) on page 3120
Syntax	<pre>show route summary <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i>></pre>
Syntax (EX Series Switches)	show route summary
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the show route summary command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show route summary command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p>
Options	<p>none—Display summary statistics about the entries in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route summary table inet command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>
Required Privilege Level	view
List of Sample Output	show route summary on page 3121 show route summary table on page 3122 show route summary table (with Route Limits Configured for the Routing Table) on page 3122
Output Fields	<p>Table 282 on page 3120 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 282: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.

Table 282: show route summary Output Fields (*continued*)

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: <ul style="list-style-type: none"> active—Number of routes that are active. holddown—Number of routes that are in the hold-down state before being declared inactive. hidden—Number of routes that are not used because of routing policy.
Limit/Threshold	Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> destinations—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. routes—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
<i>protocol-name</i>	Name of the protocol from which the route was learned. For example, OSPF , RSVP , and Static .

Sample Output

show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
      Direct:      6 routes,      5 active
      Local:       4 routes,      4 active
      OSPF:        5 routes,      4 active
      Static:      7 routes,      7 active
      IGMP:        1 routes,      1 active
      PIM:         2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
      RSVP:        2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

```
Restart Complete
  Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
  MPLS:       3 routes,      3 active
  VPLS:       4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
  Direct:      2 routes,      2 active
  PIM:         2 routes,      2 active
  MLD:         1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         4 routes,      4 active
```

show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
  Direct:      6 routes,      5 active
  Local:       9 routes,      9 active
  OSPF:        3 routes,      1 active
  Static:     13 routes,     13 active
  ICMP:        1 routes,      1 active
  PIM:         2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
  Local:        1 routes,      1 active
  PIM:          2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
  Direct:      2 routes,      2 active
  Local:       1 routes,      1 active
```

OSPF:	4 routes,	3 active
BGP:	4 routes,	4 active
IGMP:	1 routes,	1 active
PIM:	2 routes,	2 active

show route table

List of Syntax	Syntax on page 3124 Syntax (EX Series Switches) on page 3124
Syntax	<code>show route table <i>routing-table-name</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route table <i>routing-table-name</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in a particular routing table.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>routing-table-name</i> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route table inet command).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route summary on page 3120
List of Sample Output	show route table bgp.l2.vpn on page 3125 show route table bgp.l3vpn.0 on page 3125 show route table bgp.l3vpn.0 detail on page 3125 show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 3127 show route table inet.0 on page 3127 show route table inet6.0 on page 3127 show route table inet6.3 on page 3128 show route table inetflow detail on page 3128 show route table l2circuit.0 on page 3128 show route table mpls on page 3129 show route table mpls extensive on page 3129 show route table mpls.0 on page 3129 show route table mpls.0 (RSVP Route—Transit LSP) on page 3130 show route table vpls_1 detail on page 3130 show route table vpn-a on page 3130 show route table vpn-a.mdt.0 on page 3131 show route table VPN-A detail on page 3131

[show route table VPN-AB.inet.0 on page 3131](#)
[show route table VPN_blue.mvpn-inet6.0 on page 3132](#)
[show route table VPN-A detail on page 3132](#)
[show route table inetflow detail on page 3133](#)

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route table bgp.l2.vpn

```

user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table bgp.l3vpn.0

```

user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)

```

show route table bgp.l3vpn.0 detail

```

user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
    AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

```

```
Communities: 2914:420 target:11111:1 origin:56:78
VPN Label: 182449
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
```

```

AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
    *[RTarget/5] 00:03:14
        Type Proxy
        for 10.255.165.103
        for 10.255.166.124
        Local

```

show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                   > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                   > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                   Local
12.12.12.21/32     *[Local/0] 00:51:57
                   Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                   > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                   Local
13.13.13.21/32     *[Local/0] 00:51:58
                   Local
13.13.13.22/32     *[Direct/0] 00:33:59
                   > via t3-5/2/0.0
127.0.0.1/32       [Direct/0] 00:51:58
                   > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                   > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                   Local

```

show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64   *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128  *[Local/0] 00:01:34
>Local

```

```
fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
    *[LDP/9] 00:00:22, metric 1
    > via so-1/0/0.0
::10.255.245.196/128
    *[LDP/9] 00:00:08, metric 1
    > via so-1/0/0.0, Push 100008
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Next-hop reference count: 2
        State: **Active Ext>
        Local AS: 65002 Peer AS: 65000
        Age: 4
        Task: BGP_65000.10.12.99.5+3792
        Announcement bits (1): 0-Flow
        AS path: 65000 I
        Communities: traffic-rate:0:0
        Validation state: Accept, Originator: 10.12.99.5
        Via: 10.12.44.0/24, Active
        Localpref: 100
        Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow Preference: 5
        Next-hop reference count: 2
        State: **Active>
        Local AS: 65002
        Age: 6:30
        Task: RT Flow
        Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
        AS path: I
        Communities: 1:1
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
```

```

                                via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                                *[LDP/9] 00:50:14
                                Discard

```

show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:13:55, metric 1
                  Receive
1                *[MPLS/0] 00:13:55, metric 1
                  Receive
2                *[MPLS/0] 00:13:55, metric 1
                  Receive
1024             *[VPN/0] 00:04:18
                  to table red.inet.0, Pop

```

show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
              Next hop: via so-1/0/0.0, selected
              Pop
              State: <Active Int>
              Age: 29:50      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 10.0.0.194/32

```

show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:45:09, metric 1
                  Receive
1                *[MPLS/0] 00:45:09, metric 1
                  Receive
2                *[MPLS/0] 00:45:09, metric 1
                  Receive
100000           *[L2VPN/7] 00:43:04
                  > via so-0/1/0.1, Pop
100001           *[L2VPN/7] 00:43:03
                  > via so-0/1/0.2, Pop      Offset: 4
100002           *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100002(S=0)      *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100003           *[LDP/9] 00:43:22, metric 1
                  > via so-0/1/2.0, Swap 100002
                  via so-0/1/3.0, Swap 100002
100004           *[LDP/9] 00:43:16, metric 1

```

```

                via so-0/1/2.0, Swap 100049
            > via so-0/1/3.0, Swap 100049
so-0/1/0.1      *[L2VPN/7] 00:43:04
                > via so-0/1/2.0, Push 100001, Push 100049(top)
                via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2      *[L2VPN/7] 00:43:03
                via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
            > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

show route table mpls.0 (RSVP Route—Transit LSP)

```

user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:37:31, metric 1
                Receive
1                *[MPLS/0] 00:37:31, metric 1
                Receive
2                *[MPLS/0] 00:37:31, metric 1
                Receive
13               *[MPLS/0] 00:37:31, metric 1
                Receive
300352            *[RSVP/7/1] 00:08:00, metric 1
                > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)       *[RSVP/7/1] 00:08:00, metric 1
                > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384            *[RSVP/7/2] 00:05:20, metric 1
                > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)       *[RSVP/7/2] 00:05:20, metric 1
                > to 8.64.1.106 via ge-1/0/0.0, Pop

```

show route table vpls_1 detail

```

user@host> show route table vpls_1 detail

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

show route table vpn-a

```

user@host> show route table vpn-a

vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1/96
                *[VPN/7] 05:48:27
                Discard
192.168.24.1:1:2:1/96
                *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1

```

```

AS path: I
> to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
*[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
AS path: I
> to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
*[MVPN/70] 01:23:05, metric2 1
Indirect
1:1:1:10.255.14.218:232.1.1.1/144
*[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
AS path: I
> via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
*[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.179.13:200
Next hop type: Indirect
Next-hop reference count: 5
Source: 10.255.179.13
Next hop type: Router, Next hop index: 732
Next hop: 10.39.1.14 via fe-0/3/0.0, selected
Label operation: Push 299824, Push 299824(top)
Protocol next hop: 10.255.179.13
Push 299824
Indirect next hop: 8f275a0 1048574
State: (Secondary Active Int Ext)
Local AS: 1 Peer AS: 1
Age: 3:41:06 Metric: 1 Metric2: 1
Task: BGP_1.10.255.179.13+64309
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:200 rte-type:0.0.0.0:1:0
Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.13vpn.0

```

show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30 *[OSPF/10] 00:07:24, metric 1
> via so-7/3/1.0

```

```

10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

show route table VPN_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

1:10.255.2.202:65535:10.255.2.202/432
   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
   AS path: I
   > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
   AS path: I
   > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
   *[MVPN/70] 00:57:23, metric2 1
   Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
   AS path: I
   > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
   *[PIM/105] 00:02:37
   Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
   *[MVPN/70] 00:02:37, metric2 1
   Indirect

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
   *BGP      Preference: 170/-101
              Route Distinguisher: 10.255.179.13:200
              Next hop type: Indirect
              Next-hop reference count: 5
              Source: 10.255.179.13
              Next hop type: Router, Next hop index: 732
              Next hop: 10.39.1.14 via fe-0/3/0.0, selected
              Label operation: Push 299824, Push 299824(top)

```



```

Protocol next hop: 10.255.179.13
Push 299824
Indirect next hop: 8f275a0 1048574
State: (Secondary Active Int Ext)
Local AS: 1 Peer AS: 1
Age: 3:41:06 Metric: 1 Metric2: 1
Task: BGP_1.10.255.179.13+64309
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:200 rte-type:0.0.0.0:1:0
Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.l3vpn.0

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Next-hop reference count: 2
        State: **Active Ext>
        Local AS: 65002 Peer AS: 65000
        Age: 4
        Task: BGP_65000.10.12.99.5+3792
        Announcement bits (1): 0-Flow
        AS path: 65000 I
        Communities: traffic-rate:0:0
        Validation state: Accept, Originator: 10.12.99.5
        Via: 10.12.44.0/24, Active
        Localpref: 100
        Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow Preference: 5
        Next-hop reference count: 2
        State: **Active>
        Local AS: 65002
        Age: 6:30
        Task: RT Flow
        Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
        AS path: I
        Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    * [VPLS/170] 1d 03:11:03, metric2 1
        Indirect
1.1.1.4:100:1.1.1.4/96 AD
    * [BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
        AS path: I, validation-state: unverified
        > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
    * [VPLS/170] 1d 03:11:03, metric2 1
        Indirect
1.1.1.4:100:1:0/96 MH

```

```

    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
      AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
      Discard
```

show route terse


List of Syntax	Syntax on page 3135 Syntax (EX Series Switches) on page 3135
Syntax	<pre>show route terse <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show route terse
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display a high-level summary of the routes in the routing table.
<div>  <p>NOTE: For BGP routes, the <code>show route terse</code> command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the show route extensive command.</p> </div>	
Options	<p>none—Display a high-level summary of the routes in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route terse on page 3137
Output Fields	Table 283 on page 3135 describes the output fields for the <code>show route terse</code> command. Output fields are listed in the approximate order in which they appear.

Table 283: show route terse Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 283: show route terse Output Fields (*continued*)

Field Name	Field Description
route key	<p>Key for the state of the route:</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	<p>Validation status of the route:</p> <ul style="list-style-type: none"> • ?—Not evaluated. Indicates that the route was not learned through BGP. • I—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • V—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	<p>Protocol through which the route was learned:</p> <ul style="list-style-type: none"> • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel
Prf	<p>Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Table 283: show route terse Output Fields (*continued*)

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. ?—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 1.0.1.1/32        0 10      1           >10.0.0.2      I
?                               B 170      100           I
  unverified                               >10.0.0.2
* ? 1.1.1.1/32        D 0           >10.0.0.2
* V 2.2.0.2/32        B 170      110          >10.0.0.2      200 I
  valid                               >10.0.0.2
* ? 10.0.0.0/30       D 0           >1t-1/2/0.1
?                               B 170      100           I
  unverified                               >10.0.0.2
* ? 10.0.0.1/32       L 0           Local
* ? 10.0.0.4/30       B 170      100           I
  unverified                               >10.0.0.2
* ? 10.0.0.8/30       B 170      100           I
  unverified                               >10.0.0.2
* I 172.16.1.1/32     B 170      90           >10.0.0.2      200 I
  invalid                               >10.0.0.2
* N 192.168.2.3/32    B 170      100           >10.0.0.2      200 I
  unknown                               >10.0.0.2
* ? 224.0.0.5/32      O 10      1           MultiRecv

```


Troubleshooting

- [Troubleshooting Procedures on page 3139](#)

Troubleshooting Procedures

- [Troubleshooting Virtual Routing Instances on page 3139](#)

Troubleshooting Virtual Routing Instances

- [Direct Routes Not Leaked Between Routing Instances on page 3139](#)

Direct Routes Not Leaked Between Routing Instances

Problem **Description:** Direct routes are not exported (leaked) between virtual routing instances. For example, consider the following scenario:

- QFX switch with two virtual routing instances:
 - Routing instance 1 connects to downstream device through interface xe-0/0/1.
 - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the QFX switch connects to the upstream device over a direct route and these routes are not leaked between instances.



NOTE: You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the QFX switch over indirect routes.

Solution This is expected behavior.

- Related Documentation**
- [Understanding Virtual Router Routing Instances on page 2822](#)
 - [Configuring Virtual Router Routing Instances on page 2829](#)
 - [rib-group on page 2935](#)

PART 12

Border Gateway Protocol

- [Overview on page 3143](#)
- [Configuration on page 3149](#)
- [Administration on page 3583](#)

CHAPTER 36

Overview

- [BGP Overview on page 3143](#)

BGP Overview

- [Understanding BGP on page 3144](#)
- [BGP Routes Overview on page 3146](#)
- [BGP Messages Overview on page 3147](#)

Understanding BGP

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes MP_REACH_NLRI and MP_UNREACH_NLRI, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 3144](#)
- [AS Paths and Attributes on page 3144](#)
- [External and Internal BGP on page 3145](#)
- [Multiple Instances of BGP on page 3145](#)

Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

AS Paths and Attributes

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate


routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*.

[Figure 72 on page 3145](#) illustrates ASs, IBGP, and EBGP.

Figure 72: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



NOTE: When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

- Related Documentation**
- [BGP Routes Overview on page 3146](#)
 - [BGP Messages Overview on page 3147](#)

BGP Routes Overview

A BGP route is a destination, described as an IP address prefix, and information that describes the path to the destination.

The following information describes the path:

- AS path, which is a list of numbers of the ASs that a route passes through to reach the local router. The first number in the path is that of the last AS in the path—the AS closest to the local router. The last number in the path is the AS farthest from the local router, which is generally the origin of the path.
- Path attributes, which contain additional information about the AS path that is used in routing policy.

BGP peers advertise routes to each other in update messages.

BGP stores its routes in the Junos OS routing table (**inet.0**). The routing table stores the following information about BGP routes:

- Routing information learned from update messages received from peers

- Local routing information that BGP applies to routes because of local policies
- Information that BGP advertises to BGP peers in update messages

For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

The BGP router that first advertises a route assigns it one of the following values to identify its origin. During route selection, the lowest origin value is preferred.

- **0**—The router originally learned the route through an IGP (OSPF, IS-IS, or a static route).
- **1**—The router originally learned the route through an EGP (most likely BGP).
- **2**—The route's origin is unknown.

**Related
Documentation**

- [Understanding BGP Path Selection on page 3332](#)
- [Example: Advertising Multiple Paths in BGP on page 3380](#)

BGP Messages Overview

All BGP messages have the same fixed-size header, which contains a marker field that is used for both synchronization and authentication, a length field that indicates the length of the packet, and a type field that indicates the message type (for example, open, update, notification, keepalive, and so on).

This section discusses the following topics:

- [Open Messages on page 3147](#)
- [Update Messages on page 3148](#)
- [Keepalive Messages on page 3148](#)
- [Notification Messages on page 3148](#)

Open Messages

After a TCP connection is established between two BGP systems, they exchange BGP open messages to create a BGP connection between them. Once the connection is established, the two systems can exchange BGP messages and data traffic.

Open messages consist of the BGP header plus the following fields:

- **Version**—The current BGP version number is 4.
- **Local AS number**—You configure this by including the **autonomous-system** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- **Hold time**—Proposed hold-time value. You configure the local hold time with the BGP **hold-time** statement.
- **BGP identifier**—IP address of the BGP system. This address is determined when the system starts and is the same for every local interface and every BGP peer. You can

configure the BGP identifier by including the **router-id** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. By default, BGP uses the IP address of the first interface it finds in the router.

- Parameter field length and the parameter itself—These are optional fields.

Update Messages

BGP systems send update messages to exchange network reachability information. BGP systems use this information to construct a graph that describes the relationships among all known ASs.

Update messages consist of the BGP header plus the following optional fields:

- Unfeasible routes length—Length of the withdrawn routes field
- Withdrawn routes—IP address prefixes for the routes being withdrawn from service because they are no longer deemed reachable
- Total path attribute length—Length of the path attributes field; it lists the path attributes for a feasible route to a destination
- Path attributes—Properties of the routes, including the path origin, the multiple exit discriminator (MED), the originating system's preference for the route, and information about aggregation, communities, confederations, and route reflection
- Network layer reachability information (NLRI)—IP address prefixes of feasible routes being advertised in the update message

Keepalive Messages

BGP systems exchange keepalive messages to determine whether a link or host has failed or is no longer available. Keepalive messages are exchanged often enough so that the hold timer does not expire. These messages consist only of the BGP header.

Notification Messages

BGP systems send notification messages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.

Related Documentation

- [Understanding BGP on page 3144](#)
- [BGP Routes Overview on page 3146](#)

CHAPTER 37

Configuration

- [Basic BGP Configuration on page 3149](#)
- [BGP Path Attribute Configuration on page 3194](#)
- [BGP Policy Configuration on page 3305](#)
- [BGP BFD Configuration on page 3348](#)
- [BGP Load Balancing Configuration on page 3362](#)
- [IBGP Scaling Configuration on page 3405](#)
- [BGP Security Configuration on page 3428](#)
- [BGP Flap Configuration on page 3449](#)
- [BGP Monitoring Configuration on page 3475](#)
- [Configuration Statements on page 3482](#)

Basic BGP Configuration

- [Examples: Configuring External BGP Peering on page 3149](#)
- [Examples: Configuring Internal BGP Peering on page 3172](#)

Examples: Configuring External BGP Peering


- [Understanding External BGP Peering Sessions on page 3149](#)
- [Example: Configuring External BGP Point-to-Point Peer Sessions on page 3150](#)
- [Example: Configuring External BGP on Logical Systems with IPv6 Interfaces on page 3157](#)

Understanding External BGP Peering Sessions

To establish point-to-point connections between peer autonomous systems (ASs), you configure a BGP session on each interface of a point-to-point link. Generally, such sessions are made at network exit points with neighboring hosts outside the AS.

[Figure 73 on page 3150](#) shows an example of a BGP peering session.

Figure 73: BGP Peering Session



In [Figure 73 on page 3150](#), Router A is a gateway router for AS 3, and Router B is a gateway router for AS 10. For traffic internal to either AS, an interior gateway protocol (IGP) is used (OSPF, for instance). To route traffic between peer ASs, a BGP session is used.

You arrange BGP routing devices into groups of peers. Different peer groups can have different group types, AS numbers, and route reflector cluster identifiers.

To define a BGP group that recognizes only the specified BGP systems as peers, statically configure all the system's peers by including one or more **neighbor** statements. The peer neighbor's address can be either an IPv6 or IPv4 address.

As the number of external BGP (EBGP) groups increases, the ability to support a large number of BGP sessions might become a scaling issue. The preferred way to configure a large number of BGP neighbors is to configure a few groups consisting of multiple neighbors per group. Supporting fewer EBGP groups generally scales better than supporting a large number of EBGP groups. This becomes more evident in the case of hundreds of EBGP groups when compared with a few EBGP groups with multiple peers in each group.

After the BGP peers are established, BGP routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP RIB and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

Example: Configuring External BGP Point-to-Point Peer Sessions

This example shows how to configure BGP point-to-point peer sessions.

- [Requirements on page 3150](#)
- [Overview on page 3151](#)
- [Configuration on page 3151](#)
- [Verification on page 3153](#)


Requirements

Before you begin, if the default BGP policy is not adequate for your network, configure routing policies to filter incoming BGP routes and to advertise BGP routes.

Overview

Figure 74 on page 3151 shows a network with BGP peer sessions. In the sample network, Device E in AS 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22 and have IP addresses 10.10.10.2, 10.10.10.6, and 10.10.10.10. Peer D resides in AS 79, at IP address 10.21.7.2. This example shows the configuration on Device E.

Figure 74: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/2/0 unit 0 description to-A
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/0/1 unit 5 description to-B
set interfaces ge-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces ge-0/1/0 unit 9 description to-C
set interfaces ge-0/1/0 unit 9 family inet address 10.10.10.9/30
set interfaces ge-1/2/1 unit 21 description to-D
set interfaces ge-1/2/1 unit 21 family inet address 10.21.7.1/30
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 22
set protocols bgp group external-peers neighbor 10.10.10.2
set protocols bgp group external-peers neighbor 10.10.10.6
set protocols bgp group external-peers neighbor 10.10.10.10
set protocols bgp group external-peers neighbor 10.21.7.2 peer-as 79
set routing-options autonomous-system 17
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the interfaces to Peers A, B, C, and D.

```
[edit interfaces]
user@E# set ge-1/2/0 unit 0 description to-A
user@E# set ge-1/2/0 unit 0 family inet address 10.10.10.1/30
user@E# set ge-0/0/1 unit 5 description to-B
user@E# set ge-0/0/1 unit 5 family inet address 10.10.10.5/30
user@E# set ge-0/1/0 unit 9 description to-C
user@E# set ge-0/1/0 unit 9 family inet address 10.10.10.9/30
user@E# set ge-1/2/1 unit 21 description to-D
user@E# set ge-1/2/1 unit 21 family inet address 10.21.7.1/30
```

2. Set the autonomous system (AS) number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

3. Create the BGP group, and add the external neighbor addresses.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10
```

4. Specify the autonomous system (AS) number of the external AS.

```
[edit protocols bgp group external-peers]
user@E# set peer-as 22
```

5. Add Peer D, and set the AS number at the individual neighbor level.

The neighbor configuration overrides the group configuration. So, while **peer-as 22** is set for all the other neighbors in the group, **peer-as 79** is set for neighbor 10.21.7.2.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.21.7.2 peer-as 79
```

6. Set the peer type to external BGP (EBGP).

```
[edit protocols bgp group external-peers]
user@E# set type external
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@E# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-A;
    family inet {
      address 10.10.10.1/30;
```

```

    }
  }
}
ge-0/0/1 {
  unit 5 {
    description to-B;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
ge-0/1/0 {
  unit 9 {
    description to-C;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
ge-1/2/1 {
  unit 21 {
    description to-D;
    family inet {
      address 10.21.7.1/30;
    }
  }
}

[edit]
user@E# show protocols
bgp {
  group external-peers {
    type external;
    peer-as 22;
    neighbor 10.10.10.2;
    neighbor 10.10.10.6;
    neighbor 10.10.10.10;
    neighbor 10.21.7.2 {
      peer-as 79;
    }
  }
}

[edit]
user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3154](#)
- [Verifying BGP Groups on page 3156](#)
- [Verifying BGP Summary Information on page 3156](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, run the **show bgp neighbor** command.

```

user@E> show bgp neighbor
Peer: 10.10.10.2+179 AS 22      Local: 10.10.10.1+65406 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.10.2      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 0
  BFD: disabled, down
  Local Interface: ge-1/2/0.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0
    Suppressed due to damping: 0
    Advertised prefixes:      0
  Last traffic (seconds): Received 10   Sent 6   Checked 1
  Input messages: Total 8522   Updates 1   Refreshes 0   Octets 161922
  Output messages: Total 8433   Updates 0   Refreshes 0   Octets 160290
  Output Queue[0]: 0

Peer: 10.10.10.6+54781 AS 22   Local: 10.10.10.5+179 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.10.6      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 1
  BFD: disabled, down
  Local Interface: ge-0/0/1.5
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)

```

```

Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 12   Sent 6   Checked 33
Input messages: Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8430   Updates 0   Refreshes 0   Octets 160233
Output Queue[0]: 0

Peer: 10.10.10.10+55012 AS 22 Local: 10.10.10.9+179 AS 17
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.10.10.10 Local ID: 10.10.10.1 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 2
BFD: disabled, down
Local Interface: fe-0/1/0.9
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 15   Sent 6   Checked 37
Input messages: Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8429   Updates 0   Refreshes 0   Octets 160214
Output Queue[0]: 0

Peer: 10.21.7.2+61867 AS 79 Local: 10.21.7.1+179 AS 17
Type: External State: Established Flags: <ImportEval Sync>

```

```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.21.7.2          Local ID: 10.10.10.1      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 3
BFD: disabled, down
Local Interface: ge-1/2/1.21
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 79)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 28   Sent 24   Checked 47
Input messages: Total 8521   Updates 1   Refreshes 0   Octets 161943
Output messages: Total 8427   Updates 0   Refreshes 0   Octets 160176
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, run the **show bgp group** command.

```

user@E> show bgp group
Group Type: External                      Local AS: 17
  Name: external-peers   Index: 0         Flags: <>
  Holdtime: 0
  Total peers: 4         Established: 4
  10.10.10.2+179
  10.10.10.6+54781
  10.10.10.10+55012
  10.21.7.2+61867
  inet.0: 0/0/0/0

Groups: 1   Peers: 4   External: 4   Internal: 0   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending
inet.0     0          0          0           0        0        0        0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, run the **show bgp summary** command.

```
user@E> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          0          0          0          0          0          0          0
Peer           AS           InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2      22          8559     8470      0        0 2d 16:12:56
0/0/0/0         0/0/0/0
10.10.10.6      22          8566     8468      0        0 2d 16:12:12
0/0/0/0         0/0/0/0
10.10.10.10     22          8565     8466      0        0 2d 16:11:31
0/0/0/0         0/0/0/0
10.21.7.2       79          8560     8465      0        0 2d 16:10:58
0/0/0/0         0/0/0/0
```

Example: Configuring External BGP on Logical Systems with IPv6 Interfaces

This example shows how to configure external BGP (EBGP) point-to-point peer sessions on logical systems with IPv6 interfaces.

- [Requirements on page 3157](#)
- [Overview on page 3157](#)
- [Configuration on page 3158](#)
- [Verification on page 3167](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Junos OS supports EBGP peer sessions by means of IPv6 addresses. An IPv6 peer session can be configured when an IPv6 address is specified in the **neighbor** statement. This example uses EUI-64 to generate IPv6 addresses that are automatically applied to the interfaces. An EUI-64 address is an IPv6 address that uses the IEEE EUI-64 format for the interface identifier portion of the address (the last 64 bits).



NOTE: Alternatively, you can configure EBGP sessions using manually assigned 128-bit IPv6 addresses.

If you use 128-bit link-local addresses for the interfaces, you must include the **local-interface** statement. This statement is valid only for 128-bit IPv6 link-local addresses and is mandatory for configuring an IPv6 EBGP link-local peer session.

Configuring EBGP peering using link-local addresses is only applicable for directly connected interfaces. There is no support for multihop peering.


After your interfaces are up, you can use the **show interfaces terse** command to view the EUI-64-generated IPv6 addresses on the interfaces. You must use these generated

addresses in the BGP **neighbor** statements. This example demonstrates the full end-to-end procedure.

In this example, Frame Relay interface encapsulation is applied to the logical tunnel (lt) interfaces. This is a requirement because only Frame Relay encapsulation is supported when IPv6 addresses are configured on the lt interfaces.

Figure 75 on page 3158 shows a network with BGP peer sessions. In the sample network, Router R1 has five logical systems configured. Device E in autonomous system (AS) 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22. This example shows the step-by-step configuration on Logical System A and Logical System E.

Figure 75: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-E
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation frame-relay
set logical-systems A interfaces lt-0/1/0 unit 1 dlci 1
```

```

set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 25
set logical-systems A interfaces lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64
  eui-64
set logical-systems A interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set logical-systems A protocols bgp group external-peers type external
set logical-systems A protocols bgp group external-peers peer-as 17
set logical-systems A protocols bgp group external-peers neighbor
  2001:db8:0:1:2a0:a502:0:19da
set logical-systems A routing-options router-id 1.1.1.1
set logical-systems A routing-options autonomous-system 22

```

Device B

```

set logical-systems B interfaces lt-0/1/0 unit 6 description to-E
set logical-systems B interfaces lt-0/1/0 unit 6 encapsulation frame-relay
set logical-systems B interfaces lt-0/1/0 unit 6 dlci 6
set logical-systems B interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems B interfaces lt-0/1/0 unit 6 family inet6 address 2001:db8:0:2::/64
  eui-64
set logical-systems B interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set logical-systems B protocols bgp group external-peers type external
set logical-systems B protocols bgp group external-peers peer-as 17
set logical-systems B protocols bgp group external-peers neighbor
  2001:db8:0:2:2a0:a502:0:5da
set logical-systems B routing-options router-id 2.2.2.2
set logical-systems B routing-options autonomous-system 22

```

Device C

```

set logical-systems C interfaces lt-0/1/0 unit 10 description to-E
set logical-systems C interfaces lt-0/1/0 unit 10 encapsulation frame-relay
set logical-systems C interfaces lt-0/1/0 unit 10 dlci 10
set logical-systems C interfaces lt-0/1/0 unit 10 peer-unit 9
set logical-systems C interfaces lt-0/1/0 unit 10 family inet6 address 2001:db8:0:3::/64
  eui-64
set logical-systems C interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set logical-systems C protocols bgp group external-peers type external
set logical-systems C protocols bgp group external-peers peer-as 17
set logical-systems C protocols bgp group external-peers neighbor
  2001:db8:0:3:2a0:a502:0:9da
set logical-systems C routing-options router-id 3.3.3.3
set logical-systems C routing-options autonomous-system 22

```

Device D

```

set logical-systems D interfaces lt-0/1/0 unit 7 description to-E
set logical-systems D interfaces lt-0/1/0 unit 7 encapsulation frame-relay
set logical-systems D interfaces lt-0/1/0 unit 7 dlci 7
set logical-systems D interfaces lt-0/1/0 unit 7 peer-unit 21
set logical-systems D interfaces lt-0/1/0 unit 7 family inet6 address 2001:db8:0:4::/64
  eui-64
set logical-systems D interfaces lo0 unit 4 family inet6 address 2001:db8::4/128
set logical-systems D protocols bgp group external-peers type external
set logical-systems D protocols bgp group external-peers peer-as 17
set logical-systems D protocols bgp group external-peers neighbor
  2001:db8:0:4:2a0:a502:0:15da
set logical-systems D routing-options router-id 4.4.4.4
set logical-systems D routing-options autonomous-system 79

```

Device E

```

set logical-systems E interfaces lt-0/1/0 unit 5 description to-B
set logical-systems E interfaces lt-0/1/0 unit 5 encapsulation frame-relay

```

```

set logical-systems E interfaces lt-0/1/0 unit 5 dlci 6
set logical-systems E interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems E interfaces lt-0/1/0 unit 5 family inet6 address 2001:db8:0:2::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 9 description to-C
set logical-systems E interfaces lt-0/1/0 unit 9 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 9 dlci 10
set logical-systems E interfaces lt-0/1/0 unit 9 peer-unit 10
set logical-systems E interfaces lt-0/1/0 unit 9 family inet6 address 2001:db8:0:3::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 21 description to-D
set logical-systems E interfaces lt-0/1/0 unit 21 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 21 dlci 7
set logical-systems E interfaces lt-0/1/0 unit 21 peer-unit 7
set logical-systems E interfaces lt-0/1/0 unit 21 family inet6 address 2001:db8:0:4::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 25 description to-A
set logical-systems E interfaces lt-0/1/0 unit 25 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 25 dlci 1
set logical-systems E interfaces lt-0/1/0 unit 25 peer-unit 1
set logical-systems E interfaces lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64
    eui-64
set logical-systems E interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
set logical-systems E protocols bgp group external-peers type external
set logical-systems E protocols bgp group external-peers peer-as 22
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:1:2a0:a502:0:1da
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:2:2a0:a502:0:6da
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:3:2a0:a502:0:ada
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:4:2a0:a502:0:7da peer-as 79
set logical-systems E routing-options router-id 5.5.5.5
set logical-systems E routing-options autonomous-system 17

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Run the **show interfaces terse** command to verify that the physical router has a logical tunnel (lt) interface.

```

user@R1> show interfaces terse
Interface           Admin Link Proto  Local          Remote
...
lt-0/1/0             up    up
...

```

2. On Logical System A, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System E.

```

user@R1> set cli logical-system A
Logical system: A
[edit]

```

```

user@R1:A> edit
Entering configuration mode
[edit]
user@R1:A# edit interfaces
[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 encapsulation frame-relay
user@R1:A# set lt-0/1/0 unit 1 dlci 1
user@R1:A# set lt-0/1/0 unit 1 peer-unit 25

```

3. On Logical System A, configure the network address for the link to Peer E, and configure a loopback interface.

```

[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 description to-E
user@R1:A# set lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:A# set lo0 unit 1 family inet6 address 2001:db8::1/128

```

4. On Logical System E, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System A.

```

user@R1> set cli logical-system E
Logical system: E
[edit]
user@R1:E> edit
Entering configuration mode
[edit]
user@R1:E# edit interfaces
[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 encapsulation frame-relay
user@R1:E# set lt-0/1/0 unit 25 dlci 1
user@R1:E# set lt-0/1/0 unit 25 peer-unit 1

```

5. On Logical System E, configure the network address for the link to Peer A, and configure a loopback interface.

```

[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 description to-A
user@R1:E# set lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:E# set lo0 unit 5 family inet6 address 2001:db8::5/128

```

6. Run the **show interfaces terse** command to see the IPv6 addresses that are generated by EUI-64.

The 2001 addresses are used in this example in the BGP **neighbor** statements.



NOTE: The fe80 addresses are link-local addresses and are not used in this example.

```

user@R1:A> show interfaces terse
Interface          Admin Link Proto  Local              Remote
Logical system: A

betsy@tp8:A> show interfaces terse
Interface          Admin Link Proto  Local              Remote
lt-0/1/0
lt-0/1/0.1          up    up    inet6    2001:db8:0:1:2a0:a502:0:1da/64

```

```

fe80::2a0:a502:0:1da/64
1o0
1o0.1                up    up    inet6  2001:db8::1
                                fe80::2a0:a50f:fc56:1da

user@R1:E> show interfaces terse
Interface           Admin Link Proto  Local              Remote
1t-0/1/0
1t-0/1/0.25         up    up    inet6  2001:db8:0:1:2a0:a502:0:19da/64

                                fe80::2a0:a502:0:19da/64
1o0
1o0.5                up    up    inet6  2001:db8::5
                                fe80::2a0:a50f:fc56:1da

```

7. Repeat the interface configuration on the other logical systems.

Configuring the External BGP Sessions

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. On Logical System A, create the BGP group, and add the external neighbor address.


```
[edit protocols bgp group external-peers]
user@R1:A# set neighbor 2001:db8:0:1:2a0:a502:0:19da
```
2. On Logical System E, create the BGP group, and add the external neighbor address.


```
[edit protocols bgp group external-peers]
user@R1:E# set neighbor 2001:db8:0:1:2a0:a502:0:1da
```
3. On Logical System A, specify the autonomous system (AS) number of the external AS.


```
[edit protocols bgp group external-peers]
user@R1:A# set peer-as 17
```
4. On Logical System E, specify the autonomous system (AS) number of the external AS.


```
[edit protocols bgp group external-peers]
user@R1:E# set peer-as 22
```
5. On Logical System A, set the peer type to EBGP.


```
[edit protocols bgp group external-peers]
user@R1:A# set type external
```
6. On Logical System E, set the peer type to EBGP.


```
[edit protocols bgp group external-peers]
user@R1:E# set type external
```
7. On Logical System A, set the autonomous system (AS) number and router ID.


```
[edit routing-options]
user@R1:A# set router-id 1.1.1.1
user@R1:A# set autonomous-system 22
```

8. On Logical System E, set the AS number and router ID.

```
[edit routing-options]
user@R1:E# set router-id 5.5.5.5
user@R1:E# set autonomous-system 17
```

9. Repeat these steps for Peers A, B, C, and D.

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show logical-systems
A {
  interfaces {
    lt-0/1/0 {
      unit 1 {
        description to-E;
        encapsulation frame-relay;
        dlci 1;
        peer-unit 25;
        family inet6 {
          address 2001:db8:0:1::/64 {
            eui-64;
          }
        }
      }
    }
  }
  lo0 {
    unit 1 {
      family inet6 {
        address 2001:db8::1/128;
      }
    }
  }
  protocols {
    bgp {
      group external-peers {
        type external;
        peer-as 17;
        neighbor 2001:db8:0:1:2a0:a502:0:19da;
      }
    }
    routing-options {
      router-id 1.1.1.1;
      autonomous-system 22;
    }
  }
}
B {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-E;
        encapsulation frame-relay;
```

```
        dlci 6;
        peer-unit 5;
        family inet6 {
            address 2001:db8:0:2::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 2 {
        family inet6 {
            address 2001:db8::2/128;
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:2:2a0:a502:0:5da;
        }
    }
    routing-options {
        router-id 2.2.2.2;
        autonomous-system 22;
    }
}
C {
    interfaces {
        lt-0/1/0 {
            unit 10 {
                description to-E;
                encapsulation frame-relay;
                dlci 10;
                peer-unit 9;
                family inet6 {
                    address 2001:db8:0:3::/64 {
                        eui-64;
                    }
                }
            }
        }
    }
    lo0 {
        unit 3 {
            family inet6 {
                address 2001:db8::3/128;
            }
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
```



```

        type external;
        peer-as 17;
        neighbor 2001:db8:0:3:2a0:a502:0:9da;
    }
}
routing-options {
    router-id 3.3.3.3;
    autonomous-system 22;
}
D {
    interfaces {
        lt-0/1/0 {
            unit 7 {
                description to-E;
                encapsulation frame-relay;
                dlci 7;
                peer-unit 21;
                family inet6 {
                    address 2001:db8:0:4::/64 {
                        eui-64;
                    }
                }
            }
        }
    }
    lo0 {
        unit 4 {
            family inet6 {
                address 2001:db8::4/128;
            }
        }
    }
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:4:2a0:a502:0:15da;
        }
    }
    routing-options {
        router-id 4.4.4.4;
        autonomous-system 79;
    }
}
E {
    interfaces {
        lt-0/1/0 {
            unit 5 {
                description to-B;
                encapsulation frame-relay;
                dlci 6;
                peer-unit 6;
                family inet6 {

```

```
        address 2001:db8:0:2::/64 {
            eui-64;
        }
    }
}
unit 9 {
    description to-C;
    encapsulation frame-relay;
    dlci 10;
    peer-unit 10;
    family inet6 {
        address 2001:db8:0:3::/64 {
            eui-64;
        }
    }
}
unit 21 {
    description to-D;
    encapsulation frame-relay;
    dlci 7;
    peer-unit 7;
    family inet6 {
        address 2001:db8:0:4::/64 {
            eui-64;
        }
    }
}
unit 25 {
    description to-A;
    encapsulation frame-relay;
    dlci 1;
    peer-unit 1;
    family inet6 {
        address 2001:db8:0:1::/64 {
            eui-64;
        }
    }
}
lo0 {
    unit 5 {
        family inet6 {
            address 2001:db8::5/128;
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 22;
            neighbor 2001:db8:0:1:2a0:a502:0:1da;
            neighbor 2001:db8:0:2:2a0:a502:0:6da;
            neighbor 2001:db8:0:3:2a0:a502:0:ada;
            neighbor 2001:db8:0:4:2a0:a502:0:7da {
```

```

        peer-as 79;
    }
}
}
routing-options {
    router-id 5.5.5.5;
    autonomous-system 17;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3167](#)
- [Verifying BGP Groups on page 3170](#)
- [Verifying BGP Summary Information on page 3170](#)
- [Checking the Routing Table on page 3170](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, run the **show bgp neighbor** command.

```

user@R1:E> show bgp neighbor
Peer: 2001:db8:0:1:2a0:a502:0:1da+54987 AS 22 Local:
2001:db8:0:1:2a0:a502:0:19da+179 AS 17
  Type: External   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Error: 'Open Message Error' Sent: 20 Recv: 0
  Peer ID: 1.1.1.1           Local ID: 5.5.5.5           Active Holdtime: 90
  Keepalive Interval: 30     Peer index: 0
  BFD: disabled, down
  Local Interface: lt-0/1/0.25
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet6.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync

```

```

Active prefixes:          0
Received prefixes:       0
Accepted prefixes:       0
Suppressed due to damping: 0
Advertised prefixes:     0
Last traffic (seconds): Received 7   Sent 18   Checked 81
Input messages:  Total 1611  Updates 1       Refreshes 0       Octets 30660
Output messages: Total 1594  Updates 0       Refreshes 0       Octets 30356
Output Queue[0]: 0

```

```

Peer: 2001:db8:0:2:2a0:a502:0:6da+179 AS 22 Local:
2001:db8:0:2:2a0:a502:0:5da+55502 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: Open Message Error
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Error: 'Open Message Error' Sent: 26 Recv: 0
Peer ID: 2.2.2.2      Local ID: 5.5.5.5      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 2
BFD: disabled, down
Local Interface: lt-0/1/0.5
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Accepted prefixes:       0
  Suppressed due to damping: 0
  Advertised prefixes:     0
Last traffic (seconds): Received 15   Sent 8   Checked 8
Input messages:  Total 1610  Updates 1       Refreshes 0       Octets 30601
Output messages: Total 1645  Updates 0       Refreshes 0       Octets 32417
Output Queue[0]: 0

```

```

Peer: 2001:db8:0:3:2a0:a502:0:ada+55983 AS 22 Local:
2001:db8:0:3:2a0:a502:0:9da+179 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 3.3.3.3      Local ID: 5.5.5.5      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 3
BFD: disabled, down
Local Interface: lt-0/1/0.9
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast

```

```

NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Accepted prefixes:       0
  Suppressed due to damping: 0
  Advertised prefixes:     0
Last traffic (seconds): Received 21   Sent 21   Checked 67
Input messages: Total 1610 Updates 1   Refreshes 0   Octets 30641
Output messages: Total 1587 Updates 0   Refreshes 0   Octets 30223
Output Queue[0]: 0

Peer: 2001:db8:0:4:2a0:a502:0:7da+49255 AS 79 Local:
2001:db8:0:4:2a0:a502:0:15da+179 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 4.4.4.4      Local ID: 5.5.5.5      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 1
BFD: disabled, down
Local Interface: lt-0/1/0.21
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 79)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Accepted prefixes:       0
  Suppressed due to damping: 0
  Advertised prefixes:     0
Last traffic (seconds): Received 6     Sent 17   Checked 25
Input messages: Total 1615 Updates 1   Refreshes 0   Octets 30736
Output messages: Total 1593 Updates 0   Refreshes 0   Octets 30337
Output Queue[0]: 0

```

Meaning IPv6 unicast network layer reachability information (NLRI) is being exchanged between the neighbors.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, run the **show bgp group** command.

```
user@R1:~> show bgp group
Group Type: External                               Local AS: 17
  Name: external-peers  Index: 0                   Flags: <>
  Holdtime: 0
  Total peers: 4      Established: 4
  2001:db8:0:1:2a0:a502:0:1da+54987
  2001:db8:0:2:2a0:a502:0:6da+179
  2001:db8:0:3:2a0:a502:0:ada+55983
  2001:db8:0:4:2a0:a502:0:7da+49255
  inet6.0: 0/0/0/0

Groups: 1  Peers: 4   External: 4   Internal: 0   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending
inet6.0      0         0         0           0        0         0         0
inet6.2      0         0         0           0        0         0         0
```

Meaning The group type is external, and the group has four peers.

Verifying BGP Summary Information

Purpose Verify that the BGP that the peer relationships are established.

Action From operational mode, run the **show bgp summary** command.

```
user@R1:~> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending
inet6.0      0         0         0           0        0         0         0
inet6.2      0         0         0           0        0         0         0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:db8:0:1:2a0:a502:0:1da      22     1617     1600      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
2001:db8:0:2:2a0:a502:0:6da      22     1616     1651      0      0
  12:06:56 Establ
    inet6.0: 0/0/0/0
2001:db8:0:3:2a0:a502:0:ada      22     1617     1594      0      0
  12:04:32 Establ
    inet6.0: 0/0/0/0
2001:db8:0:4:2a0:a502:0:7da      79     1621     1599      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
```

Meaning The Down peers: 0 output shows that the BGP peers are in the established state.

Checking the Routing Table

Purpose Verify that the inet6.0 routing table is populated with local and direct routes.

Action From operational mode, run the **show route** command.

```

user@R1:E> show route
inet6.0: 15 destinations, 18 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128    *[Direct/0] 12:41:18
                  > via lo0.5
2001:db8:0:1::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.25
2001:db8:0:1:2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
2001:db8:0:2::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
2001:db8:0:2:2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
2001:db8:0:3::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.9
2001:db8:0:3:2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
2001:db8:0:4::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.21
2001:db8:0:4:2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::/64         *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
                  [Direct/0] 14:40:02
                  > via lt-0/1/0.9
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.21
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.25
fe80::2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
fe80::2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
fe80::2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
fe80::2a0:a50f:fc56:1da/128
                  *[Direct/0] 12:41:18
                  > via lo0.5

```

Meaning The inet6.0 routing table contains local and direct routes. To populate the routing table with other types of routes, you must configure routing policies.

Related Documentation

- [Examples: Configuring Internal BGP Peering on page 3172](#)
- [BGP Configuration Overview](#)


Examples: Configuring Internal BGP Peering

- [Understanding Internal BGP Peering Sessions on page 3172](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3173](#)
- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3184](#)

Understanding Internal BGP Peering Sessions

When two BGP-enabled devices are in the same autonomous system (AS), the BGP session is called an *internal* BGP session, or IBGP session. BGP uses the same message types on IBGP and external BGP (EBGP) sessions, but the rules for when to send each message and how to interpret each message differ slightly. For this reason, some people refer to IBGP and EBGP as two separate protocols.

Figure 76: Internal and External BGP



In [Figure 76 on page 3172](#), Device Jackson, Device Memphis, and Device Biloxi have IBGP peer sessions with each other. Likewise, Device Miami and Device Atlanta have IBGP peer sessions between each other.

The purpose of IBGP is to provide a means by which EBGP route advertisements can be forwarded throughout the network. In theory, to accomplish this task you could redistribute all of your EBGP routes into an interior gateway protocol (IGP), such as OSPF or IS-IS. This, however, is not recommended in a production environment because of the large number of EBGP routes in the Internet and because of the way that IGPs operate. In short, with that many routes the IGP churns or crashes.

Generally, the loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.

While IBGP neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every

other device through neighbor peer relationships. The **neighbor** statement creates the mesh. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all IBGP devices in the AS. The full mesh need not be physical links. Rather, the configuration on each routing device must create a full mesh of peer sessions (using multiple **neighbor** statements).



NOTE: The requirement for a full mesh is waived if you configure a confederation or route reflection.

To understand the full-mesh requirement, consider that an IBGP-learned route cannot be readvertised to another IBGP peer. The reason for preventing the readvertisement of IBGP routes and requiring the full mesh is to avoid routing loops within an AS. The AS path attribute is the means by which BGP routing devices avoid loops. The path information is examined for the local AS number only when the route is received from an EBGP peer. Because the attribute is only modified across AS boundaries, this system works well. However, the fact that the attribute is only modified across AS boundaries presents an issue inside the AS. For example, suppose that routing devices A, B, and C are all in the same AS. Device A receives a route from an EBGP peer and sends the route to Device B, which installs it as the active route. The route is then sent to Device C, which installs it locally and sends it back to Device A. If Device A installs the route, a loop is formed within the AS. The routing devices are not able to detect the loop because the AS path attribute is not modified during these advertisements. Therefore, the BGP protocol designers decided that the only assurance of never forming a routing loop was to prevent an IBGP peer from advertising an IBGP-learned route within the AS. For route reachability, the IBGP peers are fully meshed.

IBGP supports multihop connections, so IBGP neighbors can be located anywhere within the AS and often do not share a link. A recursive route lookup resolves the loopback peering address to an IP forwarding next hop. The lookup service is provided by static routes, or an IGP such as OSPF, or BGP routes.

Example: Configuring Internal BGP Peer Sessions

This example shows how to configure internal BGP peer sessions.

- [Requirements on page 3173](#)
- [Overview on page 3173](#)
- [Configuration on page 3175](#)
- [Verification on page 3182](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

In this example, you configure internal BGP (IBGP) peer sessions. The loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address,

the IBGP peer session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peer session also goes up and down. Thus, if the device has link redundancy, the loopback interface provides fault tolerance in case the physical interface or one of the links goes down.

When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The **local-address** statement enables you to specify the source information in BGP update messages. If you omit the **local-address** statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally results in the egress interface address being the expected source of update messages. When this happens, the peer session is not established because a mismatch exists between the expected source address (the egress interface of the peer) and the actual source (the loopback interface of the peer). To make sure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.

Because IBGP supports multihop connections, IBGP neighbors can be located anywhere within the autonomous system (AS) and often do not share a link. A recursive route lookup resolves the loopback peer address to an IP forwarding next hop. In this example, this service is provided by OSPF. Although interior gateway protocol (IGP) neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every other device through neighbor peer relationships. The **neighbor** statement creates the mesh.




NOTE: The requirement for a full mesh is waived if you configure a confederation or route reflection.

After the BGP peers are established, BGP routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP routing information base (RIB) and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

Figure 77 on page 3175 shows a typical network with internal peer sessions.

Figure 77: Typical Network with IBGP Sessions

**Configuration**

- [Configuring Device A on page 3176](#)
- [Configuring Device B on page 3178](#)
- [Configuring Device C on page 3180](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```
set interfaces ge-0/1/0 unit 1 description to-B
set interfaces ge-0/1/0 unit 1 family inet address 10.10.10.1/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to B and C"
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.1
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17
```

Device B

```
set interfaces ge-0/1/0 unit 2 description to-A
set interfaces ge-0/1/0 unit 2 family inet address 10.10.10.2/30
set interfaces ge-0/1/1 unit 5 description to-C
set interfaces ge-0/1/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and C"
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.6.5
```

```
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.2
set protocols ospf area 0.0.0.0 interface ge-0/1/1.5
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

Device C

```
set interfaces ge-0/1/0 unit 6 description to-B
set interfaces ge-0/1/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and B"
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.6
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

Configuring Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 1]
user@A# set description to-B
user@A# set family inet address 10.10.10.1/30
```

```
[edit interfaces]
user@A# set lo0 unit 1 family inet address 192.168.6.5/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set description "connections to B and C"
user@A# set local-address 192.168.6.5
user@A# set export send-direct
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```

user@A# set interface lo0.1 passive
user@A# set interface ge-0/1/0.1

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 2]
user@A# set from protocol direct
user@A# set then accept

```

5. Configure the router ID and the AS number.

```

[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@A# show interfaces
ge-0/1/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to B and C";
    local-address 192.168.6.5;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
  }
}

```

```
}  
}  
ospf {  
  area 0.0.0.0 {  
    interface lo0.1 {  
      passive;  
    }  
    interface ge-0/1/0.1;  
  }  
}  
  
user@A# show routing-options  
router-id 192.168.6.5;  
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device B

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure internal BGP peer sessions on Device B:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 2]  
user@B# set description to-A  
user@B# set family inet address 10.10.10.2/30
```

```
[edit interfaces ge-0/1/1]  
user@B# set unit 5 description to-C  
user@B# set unit 5 family inet address 10.10.10.5/30
```

```
[edit interfaces]  
user@B# set lo0 unit 2 family inet address 192.163.6.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]  
user@B# set type internal  
user@B# set description "connections to A and C"  
user@B# set local-address 192.163.6.4  
user@B# set export send-direct  
user@B# set neighbor 192.168.40.4  
user@B# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]  
user@B# set interface lo0.2 passive  
user@B# set interface ge-0/1/0.2  
user@B# set interface ge-0/1/1.5
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@B# set from protocol direct
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
ge-0/1/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
ge-0/1/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}

user@B# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@B# show protocols
bgp {
  group internal-peers {
    type internal;
```

```
        description "connections to A and C";
        local-address 192.163.6.4;
        export send-direct;
        neighbor 192.168.40.4;
        neighbor 192.168.6.5;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.2 {
            passive;
        }
        interface ge-0/1/0.2;
        interface ge-0/1/1.5;
    }
}
```

```
user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device C

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device C:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 6]
user@C# set description to-B
user@C# set family inet address 10.10.10.6/30

[edit interfaces]
user@C# set lo0 unit 3 family inet address 192.168.40.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@C# set type internal
user@C# set description "connections to A and B"
user@C# set local-address 192.168.40.4
user@C# set export send-direct
user@C# set neighbor 192.163.6.4
user@C# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@C# set interface lo0.3 passive
user@C# set interface ge-0/1/0.6
```


4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@C# set from protocol direct
user@C# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
ge-0/1/0 {
  unit 6 {
    description to-B;
    family inet {
      address 10.10.10.6/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.40.4/32;
    }
  }
}

user@C# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@C# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to A and B";
    local-address 192.168.40.4;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.6.5;
  }
}
ospf {
```

```
area 0.0.0.0 {  
    interface lo0.3 {  
        passive;  
    }  
    interface ge-0/1/0.6;  
}  
}
```

```
user@C# show routing-options  
router-id 192.168.40.4;  
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3182](#)
- [Verifying BGP Groups on page 3183](#)
- [Verifying BGP Summary Information on page 3184](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3184](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From operational mode, enter the **show bgp neighbor** command.

```
user@A> show bgp neighbor  
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+58852 AS 17  
Type: Internal    State: Established    Flags: Sync  
Last State: OpenConfirm    Last Event: RecvKeepAlive  
Last Error: None  
Export: [ send-direct ]  
Options: Preference LocalAddress Refresh  
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170  
Number of flaps: 0  
Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90  
Keepalive Interval: 30    Peer index: 0  
BFD: disabled, down  
NLRI for restart configured on peer: inet-unicast  
NLRI advertised by peer: inet-unicast  
NLRI for this session: inet-unicast  
Peer supports Refresh capability (2)  
Restart time configured on the peer: 120  
Stale routes from peer are kept for: 300  
Restart time requested by this peer: 120  
NLRI that peer supports restart for: inet-unicast  
NLRI that restart is negotiated for: inet-unicast  
NLRI of received end-of-rib markers: inet-unicast  
NLRI of all end-of-rib markers sent: inet-unicast  
Peer supports 4 byte AS extension (peer-as 17)  
Peer does not support Addpath  
Table inet.0 Bit: 10000  
RIB State: BGP restart is complete  
Send state: in sync
```

```

Active prefixes:          0
Received prefixes:       3
Accepted prefixes:       3
Suppressed due to damping: 0
Advertised prefixes:     2
Last traffic (seconds): Received 25   Sent 19   Checked 67
Input messages:  Total 2420   Updates 4       Refreshes 0       Octets 46055
Output messages: Total 2411   Updates 2       Refreshes 0       Octets 45921
Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17   Local: 192.168.6.5+56466 AS 17
Type: Internal   State: Established   Flags: Sync
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ send-direct ]
Options: Preference LocalAddress Refresh
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4   Local ID: 192.168.6.5       Active Holdtime: 90
Keepalive Interval: 30   Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:       2
Accepted prefixes:       2
Suppressed due to damping: 0
Advertised prefixes:     2
Last traffic (seconds): Received 7   Sent 21   Checked 24
Input messages:  Total 2412   Updates 2       Refreshes 0       Octets 45867
Output messages: Total 2409   Updates 2       Refreshes 0       Octets 45883
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal   AS: 17                               Local AS: 17
Name: internal-peers   Index: 0                               Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2         Established: 2
192.163.6.4+179
192.168.40.4+179

```

```
inet.0: 0/5/5/0
```

Groups: 1	Peers: 2	External: 0	Internal: 2	Down peers: 0	Flaps: 0
Table	Tot Paths	Act Paths	Suppressed	History	Damp State
inet.0	5	0	0	0	0
					Pending
					0

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, enter the **show bgp summary** command.

```
user@A> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 5 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4 17 2441 2432 0 0 18:18:52
0/3/3/0 0/0/0/0
192.168.40.4 17 2432 2430 0 0 18:18:48
0/2/2/0 0/0/0/0
```

Verifying That BGP Routes Are Installed in the Routing Table

Purpose Verify that the export policy configuration is causing the BGP routes to be installed in the routing tables of the peers.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@A> show route protocol bgp
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
10.10.10.4/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
[BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.163.6.4/32 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.168.40.4/32 [BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
```

Example: Configuring Internal BGP Peering Sessions on Logical Systems

This example shows how to configure internal BGP peer sessions on logical systems.

- [Requirements on page 3185](#)
- [Overview on page 3185](#)

- [Configuration on page 3185](#)
- [Verification on page 3192](#)

Requirements

In this example, no special configuration beyond device initialization is required.


Overview

In this example, you configure internal BGP (IBGP) peering sessions.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

[Figure 78 on page 3185](#) shows a typical network with internal peer sessions.

Figure 78: Typical Network with IBGP Sessions



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-B
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-0/1/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-0/1/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
```

```

set logical-systems A routing-options autonomous-system 17
set logical-systems B interfaces lt-0/1/0 unit 2 description to-A
set logical-systems B interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-0/1/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-0/1/0 unit 5 description to-C
set logical-systems B interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-0/1/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17
set logical-systems C interfaces lt-0/1/0 unit 6 description to-B
set logical-systems C interfaces lt-0/1/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-0/1/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-0/1/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.


```

[edit logical-systems A interfaces lt-0/1/0 unit 1]
user@R1# set description to-B
user@R1# set encapsulation ethernet
user@R1# set peer-unit 2
user@R1# set family inet address 10.10.10.1/30

```

```

user@R1# set family inet address 192.168.6.5/32
user@R1# up
user@R1# up
[edit logical-systems A interfaces]
user@R1# set lo0 unit 1 family inet address 192.168.6.5/32
user@R1# exit
[edit]
user@R1# edit logical-systems B interfaces lt-0/1/0
[edit logical-systems B interfaces lt-0/1/0]
user@R1# set unit 2 description to-A
user@R1# set unit 2 encapsulation ethernet
user@R1# set unit 2 peer-unit 1
user@R1# set unit 2 family inet address 10.10.10.2/30
user@R1# set unit 5 description to-C
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 6
user@R1# set family inet address 10.10.10.5/30
user@R1# up
[edit logical-systems B interfaces]
user@R1# set lo0 unit 2 family inet address 192.163.6.4/32
user@R1# exit
[edit]
user@R1# edit logical-systems C interfaces lt-0/1/0 unit 6
[edit logical-systems C interfaces lt-0/1/0 unit 6]
set description to-B
set encapsulation ethernet
set peer-unit 5
set family inet address 10.10.10.6/30
user@R1# up
user@R1# up
[edit logical-systems C interfaces]
set lo0 unit 3 family inet address 192.168.40.4/32

```

2. Configure BGP.

On Logical System A, the **neighbor** statements are included for both Device B and Device C, even though Logical System A is not directly connected to Device C.

```

[edit logical-systems A protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.6.5
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.40.4

```

```

[edit logical-systems B protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.163.6.4
user@R1# set export send-direct
user@R1# set neighbor 192.168.40.4
user@R1# set neighbor 192.168.6.5

```

```

[edit logical-systems C protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.40.4
user@R1# set export send-direct

```

```
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface lt-0/1/0.1
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.2 passive
user@R1# set interface lt-0/1/0.2
user@R1# set interface lt-0/1/0.5
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.3 passive
user@R1# set interface lt-0/1/0.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit logical-systems A policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems B policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems C policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit logical-systems A routing-options]
user@R1# set router-id 192.168.6.5
user@R1# set autonomous-system 17
```

```
[edit logical-systems B routing-options]
user@R1# set router-id 192.163.6.4
user@R1# set autonomous-system 17
```

```
[edit logical-systems C routing-options]
user@R1# set router-id 192.168.40.4
user@R1# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show logical-systems
A {
```



```

interfaces {
  lt-0/1/0 {
    unit 1 {
      description to-B;
      encapsulation ethernet;
      peer-unit 2;
      family inet {
        address 10.10.10.1/30;
      }
    }
  }
  lo0 {
    unit 1 {
      family inet {
        address 192.168.6.5/32;
      }
    }
  }
}
protocols {
  bgp {
    group internal-peers {
      type internal;
      local-address 192.168.6.5;
      export send-direct;
      neighbor 192.163.6.4;
      neighbor 192.168.40.4;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.1 {
        passive;
      }
      interface lt-0/1/0.1;
    }
  }
}
policy-options {
  policy-statement send-direct {
    term 2 {
      from protocol direct;
      then accept;
    }
  }
}
routing-options {
  router-id 192.168.6.5;
  autonomous-system 17;
}
}
B {
  interfaces {
    lt-0/1/0 {
      unit 2 {
        description to-A;

```

```
        encapsulation ethernet;
        peer-unit 1;
        family inet {
            address 10.10.10.2/30;
        }
    }
    unit 5 {
        description to-C;
        encapsulation ethernet;
        peer-unit 6;
        family inet {
            address 10.10.10.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.163.6.4/32;
        }
    }
}
}
protocols {
    bgp {
        group internal-peers {
            type internal;
            local-address 192.163.6.4;
            export send-direct;
            neighbor 192.168.40.4;
            neighbor 192.168.6.5;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.2 {
                passive;
            }
            interface lt-0/1/0.2;
            interface lt-0/1/0.5;
        }
    }
}
policy-options {
    policy-statement send-direct {
        term 2 {
            from protocol direct;
            then accept;
        }
    }
}
routing-options {
    router-id 192.163.6.4;
    autonomous-system 17;
}
}
```

```

C {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-B;
        encapsulation ethernet;
        peer-unit 5;
        family inet {
          address 10.10.10.6/30;
        }
      }
    }
  }
  lo0 {
    unit 3 {
      family inet {
        address 192.168.40.4/32;
      }
    }
  }
}
protocols {
  bgp {
    group internal-peers {
      type internal;
      local-address 192.168.40.4;
      export send-direct;
      neighbor 192.163.6.4;
      neighbor 192.168.6.5;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.3 {
        passive;
      }
      interface lt-0/1/0.6;
    }
  }
}
policy-options {
  policy-statement send-direct {
    term 2 {
      from protocol direct;
      then accept;
    }
  }
}
routing-options {
  router-id 192.168.40.4;
  autonomous-system 17;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3192](#)
- [Verifying BGP Groups on page 3193](#)
- [Verifying BGP Summary Information on page 3193](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3194](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the operational mode, enter the **show bgp neighbor** command.

```
user@R1> show bgp neighbor logical-system A
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+58852 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-direct ]
  Options: <Preference LocalAddress Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 17)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        3
    Accepted prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      2
  Last traffic (seconds): Received 16    Sent 1    Checked 63
  Input messages: Total 15713 Updates 4    Refreshes 0    Octets 298622
  Output messages: Total 15690 Updates 2    Refreshes 0    Octets 298222
  Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17    Local: 192.168.6.5+56466 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
```

```

Export: [ send-direct ]
Options: <Preference LocalAddress Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:      0
  Received prefixes:    2
  Accepted prefixes:    2
  Suppressed due to damping: 0
  Advertised prefixes:  2
Last traffic (seconds): Received 15    Sent 22    Checked 68
Input messages: Total 15688 Updates 2    Refreshes 0    Octets 298111
Output messages: Total 15688 Updates 2    Refreshes 0    Octets 298184
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the operational mode, enter the **show bgp group** command.

```

user@A> show bgp group logical-system A
Group Type: Internal    AS: 17                      Local AS: 17
Name: internal-peers   Index: 0                    Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2          Established: 2
192.163.6.4+179
192.168.40.4+179
inet.0: 0/5/5/0

Groups: 1  Peers: 2   External: 0   Internal: 2   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History  Damp State   Pending
inet.0           5           0           0           0        0         0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary logical-system A

```

```

Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      5          0          0          0        0      0      0
Peer        AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4  17      15723   15700    0      0 4d 22:13:15
0/3/3/0      0/0/0/0
192.168.40.4 17      15698   15699    0      0 4d 22:13:11
0/2/2/0      0/0/0/0

```

Verifying That BGP Routes Are Installed in the Routing Table

Purpose Verify that the export policy configuration is working.

Action From the operational mode, enter the **show route protocol bgp** command.

```

user@A> show route protocol bgp logical-system A
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via 1t-0/1/0.1
10.10.10.4/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via 1t-0/1/0.1
                  [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via 1t-0/1/0.1
192.163.6.4/32     [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via 1t-0/1/0.1
192.168.40.4/32    [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via 1t-0/1/0.1

```

Related Documentation

- [Examples: Configuring External BGP Peering on page 3149](#)

BGP Path Attribute Configuration

- [Example: Configuring BGP Local Preference on page 3194](#)
- [Examples: Configuring BGP MED on page 3208](#)
- [Examples: Configuring BGP Local AS on page 3246](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3266](#)

Example: Configuring BGP Local Preference

- [Understanding the BGP Local Preference on page 3194](#)
- [Example: Configuring the Local Preference Value for BGP Routes on page 3195](#)

Understanding the BGP Local Preference

Internal BGP (IBGP) sessions use a metric called the *local preference*, which is carried in IBGP update packets in the path attribute LOCAL_PREF. When an autonomous system

(AS) has multiple routes to another AS, the local preference indicates the degree of preference for one route over the other routes. The route with the highest local preference value is preferred.

The LOCAL_PREF path attribute is always advertised to IBGP peers and to neighboring confederations. It is never advertised to external BGP (EBGP) peers. The default behavior is to not modify the LOCAL_PREF path attribute if it is present.

The LOCAL_PREF path attribute applies at export time only, when the routes are exported from the routing table into BGP.

If a BGP route is received without a LOCAL_PREF attribute, the route is stored in the routing table and advertised by BGP as if it were received with a LOCAL_PREF value of 100. A non-BGP route that is advertised by BGP is advertised with a LOCAL_PREF value of 100 by default.

Example: Configuring the Local Preference Value for BGP Routes

This example shows how to configure local preference in internal BGP (IBGP) peer sessions.

- [Requirements on page 3195](#)
- [Overview on page 3195](#)
- [Configuration on page 3196](#)
- [Verification on page 3206](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.


Overview

To change the local preference metric advertised in the path attribute, you must include the **local-preference** statement, specifying a value from 0 through 4,294,967,295 ($2^{32} - 1$).

There are several reasons you might want to prefer one path over another. For example, compared to other paths, one path might be less expensive to use, might have higher bandwidth, or might be more stable.

[Figure 79 on page 3196](#) shows a typical network with internal peer sessions and multiple exit points to a neighboring AS.

Figure 79: Typical Network with IBGP Sessions and Multiple Exit Points



To reach Device R4, Device R1 can take a path through either Device R2 or Device R3. By default, the local preference is 100 for either route. When the local preferences are equal, Junos OS has rules for breaking the tie and choosing a path. (See [“Understanding BGP Path Selection” on page 3332](#).) In this example, the active route is through Device R2 because the router ID of Device R2 is lower than the router ID of Device R3. The following example shows how to override the default behavior with an explicit setting for the local preference. The example configures a local preference of 300 on Device R3, thereby making Device R3 the preferred path to reach Device R4.

Configuration

- [Configuring Device R1 on page 3198](#)
- [Configuring Device R2 on page 3200](#)
- [Configuring Device R3 on page 3202](#)
- [Configuring Device R4 on page 3204](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
```



```

set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1

```

Device R2

```

set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1

```

Device R3

```

set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1

```

Device R4

```

set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3
set protocols bgp group external neighbor 24.24.24.2
set policy-options policy-statement send-direct term 1 from protocol direct

```

```
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1
```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24
```

```
[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show routing-options

```

```
autonomous-system 123;  
router-id 192.168.1.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]  
user@R2# set family inet address 12.12.12.21/24  
  
[edit interfaces fe-1/2/1 unit 4]  
user@R2# set family inet address 24.24.24.2/24  
  
[edit interfaces lo0 unit 2]  
user@R2# set family inet address 192.168.2.1/32
```
2. Configure BGP.

```
[edit protocols bgp group internal]  
user@R2# set type internal  
user@R2# set local-address 192.168.2.1  
user@R2# set export send-direct  
user@R2# set neighbor 192.168.1.1  
user@R2# set neighbor 192.168.3.1  
  
[edit protocols bgp group external]  
user@R2# set type external  
user@R2# set export send-direct  
user@R2# set peer-as 4  
user@R2# set neighbor 24.24.24.4
```
3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]  
user@R2# set interface lo0.2 passive  
user@R2# set interface fe-1/2/0.3  
user@R2# set interface fe-1/2/1.4
```
4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]  
user@R2# set from protocol direct  
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
```

```
}  
}  
ospf {  
  area 0.0.0.0 {  
    interface lo0.2 {  
      passive;  
    }  
    interface fe-1/2/0.3;  
    interface fe-1/2/1.4;  
  }  
}  
  
user@R2# show routing-options  
autonomous-system 123;  
router-id 192.168.2.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32
2. Configure BGP.

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4
3. Configure OSPF.

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5

```
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
```

```
user@R3# set from protocol direct
```

```
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
```

```
user@R3# set autonomous-system 123
```

```
user@R3# set router-id 192.168.3.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
```

```
fe-1/2/0 {
  unit 5 {
    family inet {
      address 13.13.13.3/24;
    }
  }
}
fe-1/2/1 {
  unit 6 {
    family inet {
      address 34.34.34.3/24;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.3.1/32;
    }
  }
}
```

```
user@R3# show policy-options
```

```
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
```

```
user@R3# show protocols
```

```
bgp {
  group internal {
    type internal;
    local-address 192.168.3.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.2.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 34.34.34.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.3 {
      passive;
    }
    interface fe-1/2/0.5;
    interface fe-1/2/1.6;
  }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
```
2. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
```



```
user@R4# set neighbor 34.34.34.3
user@R4# set neighbor 24.24.24.2
```

3. Configure a policy that accepts direct routes.



NOTE: Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```

4. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.4.1/32;
    }
  }
}

user@R4# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
```

```
user@R4# show protocols
bgp {
  group external {
    type external;
    export send-direct;
    peer-as 123;
    neighbor 34.34.34.3;
    neighbor 24.24.24.2;
  }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3206](#)
- [Altering the Local Preference to Change the Path Selection on page 3207](#)
- [Rechecking the Active Path From Device R1 to Device R4 on page 3207](#)

Checking the Active Path From Device R1 to Device R4

Purpose Verify that the active path from Device R1 to Device R4 goes through Device R2.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 11 destinations, 18 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32     [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32     [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32     *[BGP/170] 00:05:14, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
                  [BGP/170] 00:05:14, localpref 100, from 192.168.3.1
```

```
AS path: 4 I
> to 13.13.13.3 via fe-1/2/1.2
```

Meaning The asterisk (*) shows that the preferred path is through Device R2. In the default configuration, Device R2 has a lower router ID than Device R3. The router ID is controlling the path selection.

Altering the Local Preference to Change the Path Selection

Purpose Change the path so that it goes through Device R3.

Action From configuration mode, enter the **set local-preference 300** command.

```
[edit protocols bgp group internal]
user@R3# set local-preference 300
user@R3# commit
```

Rechecking the Active Path From Device R1 to Device R4

Purpose Verify that the active path from Device R1 to Device R4 goes through Device R3.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 11 destinations, 17 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32      *[BGP/170] 00:00:21, localpref 300, from 192.168.3.1
                  AS path: 4 I
                  > to 13.13.13.3 via fe-1/2/1.2
```

Meaning The asterisk (*) shows that the preferred path is through Device R3. In the altered configuration, Device R3 has a higher local preference than Device R2. The local preference is controlling the path selection.

Related Documentation

- [Examples: Configuring Internal BGP Peering on page 3172](#)
- [BGP Configuration Overview](#)

Examples: Configuring BGP MED

- [Understanding the MED Attribute on page 3208](#)
- [Example: Configuring the MED Attribute Directly on page 3210](#)
- [Example: Configuring the MED Using Route Filters on page 3223](#)
- [Example: Configuring the MED Using Communities on page 3236](#)
- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3236](#)

Understanding the MED Attribute

The BGP multiple exit discriminator (MED, or MULTI_EXIT_DISC) is a non-transitive attribute, meaning that it is not propagated throughout the Internet, but only to adjacent autonomous systems (ASs). The MED attribute is optional, meaning that it is not always sent with the BGP updates. The purpose of MED is to influence how other ASs enter your AS to reach a certain prefix.

The MED attribute has a value that is referred to as a *metric*. If all other factors in determining an exit point are equal, the exit point with the lowest metric is preferred.

If a MED is received over an external BGP link, it is propagated over internal links to other BGP-enabled devices within the AS.

BGP update messages include a MED metric if the route was learned from BGP and already had a MED metric associated with it, or if you configure the MED metric in the configuration file.

A MED metric is advertised with a route according to the following general rules:

- A more specific metric overrides a less specific metric. That is, a group-specific metric overrides a global BGP metric, and a peer-specific metric overrides a global BGP or group-specific metric.
- A metric defined with a routing policy overrides a metric defined with the **metric-out** statement.
- If any metric is defined, it overrides a metric received in a route.
- If the received route does not have an associated MED metric, and if you do not explicitly configure a metric value, no metric is advertised. When you do not explicitly configure a metric value, the MED value is equivalent to zero (0) when advertising an active route.

Because the AS path rather than the number of hops between hosts is the primary criterion for BGP route selection, an AS with multiple connections to a peer AS can have multiple equivalent AS paths. When the routing table contains two routes to the same host in a neighboring AS, an MED metric assigned to each route can determine which to include in the forwarding table. The MED metric you assign can force traffic through a particular exit point in an AS.

[Figure 80 on page 3209](#) illustrates how MED metrics are used to determine route selection.

Figure 80: Default MED Example




Figure 80 on page 3209 shows AS 1 and AS 2 connected by two separate BGP links to Routers C and D. Host E in AS 1 is located nearer to Router C. Host F, also in AS 1, is located nearer to Router D. Because the AS paths are equivalent, two routes exist for each host, one through Router C and one through Router D. To force all traffic destined for Host E through Router C, the network administrator for AS 2 assigns an MED metric for each router to Host E at its exit point. An MED metric of 10 is assigned to the route to Host E through Router C, and an MED metric of 20 is assigned to the route to Host E through Router D. BGP routers in AS 2 then select the route with the lower MED metric for the forwarding table.

By default, only the MEDs of routes that have the same peer ASs are compared. However, you can configure the routing table path selection options listed in Table 284 on page 3209 to compare MEDs in different ways. The MED options are not mutually exclusive and can be configured in combination or independently. For the MED options to take effect, you must configure them uniformly all through your network. The MED option or options you configure determine the route selected. Thus we recommend that you carefully evaluate your network for preferred routes before configuring the MED options.

Table 284: MED Options for Routing Table Path Selection

Option (Name)	Function	Use
Always comparing MEDs (always-compare-med)	Ensures that the MEDs for paths from peers in different ASs are always compared in the route selection process.	Useful when all enterprises participating in a network agree on a uniform policy for setting MEDs. For example, in a network shared by two ISPs, both must agree that a certain path is the better path to configure the MED values correctly.

Table 284: MED Options for Routing Table Path Selection (*continued*)

Option (Name)	Function	Use
Adding IGP cost to MED (med-plus-igp)	<p>Before comparing MED values for path selection, adds to the MED the cost of the IGP route to the BGP next-hop destination.</p> <p>This option replaces the MED value for the router, but does not affect the IGP metric comparison. As a result, when multiple routes have the same value after the MED-plus-IGP comparison, and route selection continues, the IGP route metric is also compared, even though it was added to the MED value and compared earlier in the selection process.</p>	Useful when the downstream AS requires the complete cost of a certain route that is received across multiple ASs.
Applying Cisco IOS nondeterministic behavior (cisco-non-deterministic)	<p>Specifies the nondeterministic behavior of the Cisco IOS software:</p> <ul style="list-style-type: none"> The active path is always first. All nonactive but eligible paths follow the active path and are maintained in the order in which they were received. Ineligible paths remain at the end of the list. When a new path is added to the routing table, path comparisons are made among all routes, including those paths that must never be selected because they lose the MED tie-breaking rule. 	We recommend that you do not configure this option, because the nondeterministic behavior sometimes prevents the system from properly comparing the MEDs between paths.

Example: Configuring the MED Attribute Directly

This example shows how to configure a multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3210](#)
- [Overview on page 3210](#)
- [Configuration on page 3212](#)
- [Verification on page 3222](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

To directly configure a MED metric to advertise in BGP update messages, include the **metric-out** statement:

```
metric-out (metric | minimum-igp offset | igp delay-med-update | offset);
```

metric is the primary metric on all routes sent to peers. It can be a value in the range from 0 through 4,294,967,295 ($2^{32} - 1$).


The following optional settings are also supported:

- **minimum-igp**—Sets the metric to the minimum metric value calculated in the interior gateway protocol (IGP) to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value.
- **igp**—Sets the metric to the most recent metric value calculated in the IGP to get to the BGP next hop.
- **delay-med-update**—Delays sending MED updates when the MED value increases. Include the **delay-med-update** statement when you configure the **igp** statement. The default interval to delay sending updates, unless the MED is lower or another attribute associated with the route has changed is 10 minutes. Include the **med-igp-update-interval minutes** statement at the [edit routing-options] hierarchy level to modify the default interval.
- **offset**—Specifies a value for **offset** to increase or decrease the metric that is used from the metric value calculated in the IGP. The metric value is offset by the value specified. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is increased if the **offset** value is positive. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is decreased if the **offset** value is negative.

offset can be a value in the range from -2^{31} through $2^{31} - 1$. Note that the adjusted metric can never go below 0 or above $2^{32} - 1$.

Figure 81 on page 3211 shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).

Figure 81: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 and a MED value of 20 to Device R2. This causes all of the devices in AS 123 to prefer the path through Device R2 to reach AS 4.

Configuration

- [Configuring Device R1 on page 3213](#)
- [Configuring Device R2 on page 3215](#)
- [Configuring Device R3 on page 3217](#)
- [Configuring Device R4 on page 3220](#)

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

Device R2

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```


Device R3

```

set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1

```

Device R4

```

set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 metric-out 30
set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1

```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24

[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24

[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32

```
2. Configure BGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}

user@R1# show policy-options
```

```

policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24

[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24

[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32

```
2. Configure BGP.


```

[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1

```

```
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1
```

```
[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}
```

```

    }
  }

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-1/2/0.3;
    interface fe-1/2/1.4;
  }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

[edit interfaces fe-1/2/1 unit 6]

```

```
user@R3# set family inet address 34.34.34.3/24
```

```
[edit interfaces lo0 unit 3]
```

```
user@R3# set family inet address 192.168.3.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
```

```
user@R3# set type internal
```

```
user@R3# set local-address 192.168.3.1
```

```
user@R3# set export send-direct
```

```
user@R3# set neighbor 192.168.1.1
```

```
user@R3# set neighbor 192.168.2.1
```

```
[edit protocols bgp group external]
```

```
user@R3# set type external
```

```
user@R3# set export send-direct
```

```
user@R3# set peer-as 4
```

```
user@R3# set neighbor 34.34.34.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R3# set interface lo0.3 passive
```

```
user@R3# set interface fe-1/2/0.5
```

```
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
```

```
user@R3# set from protocol direct
```

```
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
```

```
user@R3# set autonomous-system 123
```

```
user@R3# set router-id 192.168.3.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
```

```
fe-1/2/0 {
```

```
  unit 5 {
```

```
    family inet {
```

```
      address 13.13.13.3/24;
```

```
    }
```

```
  }
```

```
fe-1/2/1 {
```

```
  unit 6 {
```

```

        family inet {
            address 34.34.34.3/24;
        }
    }
}
lo0 {
    unit 3 {
        family inet {
            address 192.168.3.1/32;
        }
    }
}

user@R3# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R3# show protocols
bgp {
    group internal {
        type internal;
        local-address 192.168.3.1;
        export send-direct;
        neighbor 192.168.1.1;
        neighbor 192.168.2.1;
    }
    group external {
        type external;
        export send-direct;
        peer-as 4;
        neighbor 34.34.34.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.3 {
            passive;
        }
        interface fe-1/2/0.5;
        interface fe-1/2/1.6;
    }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24
```

```
[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24
```

```
[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32
```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
```

4. Configure a MED value of 30 for neighbor Device R3, and a MED value of 20 for neighbor Device R2.

```
[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 metric-out 30
user@R4# set neighbor 24.24.24.2 metric-out 20
```

This configuration causes autonomous system (AS) 123 (of which Device R1, Device R2, and Device R3 are members) to prefer the path through Device R2 to reach AS 4.

5. Configure the router ID and AS number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.4.1/32;
      address 44.44.44.44/32;
      address 144.144.144.144/32;
    }
  }
}

user@R4# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R4# show protocols
bgp {
  group external {
    type external;
    export send-direct;
    peer-as 123;
    neighbor 34.34.34.3 {
      metric-out 30;
    }
    neighbor 24.24.24.2 {
      metric-out 20;
    }
  }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3222](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3222](#)

Checking the Active Path From Device R1 to Device R4

Purpose Verify that the active path goes through Device R2.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32     *[BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:08:13, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.2.1/32      [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32     *[BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
```

Meaning The asterisk (*) shows that the preferred path is through Device R2. The reason for the path selection is listed as MED 20.

Verifying That Device R4 Is Sending Its Routes Correctly

Purpose Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

Action From operational mode, enter the **show route advertising-protocol bgp 24.24.24.2** command.

```
user@R4> show route advertising-protocol bgp 24.24.24.2
```

```
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref   AS path
* 24.24.24.0/24         Self              20                I
* 34.34.34.0/24         Self              20                I
* 44.44.44.44/32        Self              20                I
* 144.144.144.144/32    Self              20                I
* 192.168.4.1/32        Self              20                I
```

```
user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref   AS path
* 24.24.24.0/24         Self              30                I
* 34.34.34.0/24         Self              30                I
* 44.44.44.44/32        Self              30                I
* 144.144.144.144/32    Self              30                I
* 192.168.4.1/32        Self              30                I
```

Meaning The MED column shows that Device R4 is sending the correct MED values to its two external BGP (EBGP) neighbors.

Example: Configuring the MED Using Route Filters

This example shows how to configure a policy that uses route filters to modify the multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3223](#)
- [Overview on page 3223](#)
- [Configuration on page 3224](#)
- [Verification on page 3234](#)

Requirements


No special configuration beyond device initialization is required before you configure this example.

Overview

To configure a route-filter policy that modifies the advertised MED metric in BGP update messages, include the **metric** statement in the policy action.

[Figure 82 on page 3224](#) shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).

Figure 82: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 for all routes except 144.144.144.144. For 144.144.144.144, a MED value of 10 is advertised to Device 3. A MED value of 20 is advertised to Device R2, regardless of the route prefix.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

Device R2

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
```

```

set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1

```

Device R3

```

set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1

```

Device R4

```

set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 export med-10
set protocols bgp group external neighbor 34.34.34.3 export med-30
set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-options policy-statement med-10 then metric 10
set policy-options policy-statement med-10 then accept
set policy-options policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-options policy-statement med-30 then metric 30
set policy-options policy-statement med-30 then accept
set policy-options policy-statement send-direct term 1 from protocol direct

```

```
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1
```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24
```

```
[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24
```

```
[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24
```

```
[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1
```

```
[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```


Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
  }
}

```

```
}  
interface fe-1/2/0.3;  
interface fe-1/2/1.4;  
}  
}  
  
user@R2# show routing-options  
autonomous-system 123;  
router-id 192.168.2.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 5]  
user@R3# set family inet address 13.13.13.3/24
```

```
[edit interfaces fe-1/2/1 unit 6]  
user@R3# set family inet address 34.34.34.3/24
```

```
[edit interfaces lo0 unit 3]  
user@R3# set family inet address 192.168.3.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]  
user@R3# set type internal  
user@R3# set local-address 192.168.3.1  
user@R3# set export send-direct  
user@R3# set neighbor 192.168.1.1  
user@R3# set neighbor 192.168.2.1
```

```
[edit protocols bgp group external]  
user@R3# set type external  
user@R3# set export send-direct  
user@R3# set peer-as 4  
user@R3# set neighbor 34.34.34.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]  
user@R3# set interface lo0.3 passive  
user@R3# set interface fe-1/2/0.5  
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 5 {
    family inet {
      address 13.13.13.3/24;
    }
  }
}
fe-1/2/1 {
  unit 6 {
    family inet {
      address 34.34.34.3/24;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.3.1/32;
    }
  }
}

user@R3# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R3# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.3.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.2.1;
  }
  group external {
```

```
        type external;
        export send-direct;
        peer-as 4;
        neighbor 34.34.34.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.3 {
            passive;
        }
        interface fe-1/2/0.5;
        interface fe-1/2/1.6;
    }
}
```

```
user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24
```

```
[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24
```

```
[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32
```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
```

```

user@R4# set export send-direct
user@R4# set peer-as 123

```

4. Configure the two MED policies.

```

[edit policy-options]
set policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-statement med-10 then metric 10
set policy-statement med-10 then accept
set policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-statement med-30 then metric 30
set policy-statement med-30 then accept

```

5. Configure the two EBGP neighbors, applying the two MED policies to Device R3, and a MED value of 20 to Device R2.

```

[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 export med-10
user@R4# set neighbor 34.34.34.3 export med-30
user@R4# set neighbor 24.24.24.2 metric-out 20

```

6. Configure the router ID and autonomous system (AS) number.

```

[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.4.1/32;
      address 44.44.44.44/32;
      address 144.144.144.144/32;
    }
  }
}

```

```
user@R4# show policy-options
policy-statement med-10 {
  from {
    route-filter 144.144.144.144/32 exact;
  }
  then {
    metric 10;
    accept;
  }
}
policy-statement med-30 {
  from {
    route-filter 0.0.0.0/0 longer;
  }
  then {
    metric 30;
    accept;
  }
}
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R4# show protocols
bgp {
  group external {
    type external;
    export send-direct;
    peer-as 123;
    neighbor 24.24.24.2 {
      metric-out 20;
    }
    neighbor 34.34.34.3 {
      export [ med-10 med-30 ];
    }
  }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3234](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3235](#)

Checking the Active Path From Device R1 to Device R4

Purpose Verify that the active path goes through Device R2.

Action From operational mode, enter the **show route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32     *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:06:03, MED 10, localpref 100, from 192.168.3.1
                  AS path: 4 I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32     [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32     [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32     *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1

```

Meaning The output shows that the preferred path to the routes advertised by Device R4 is through Device R2 for all routes except 144.144.144.144/32. For 144.144.144.144/32, the preferred path is through Device R3.

Verifying That Device R4 Is Sending Its Routes Correctly

Purpose Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

Action From operational mode, enter the **show route advertising-protocol bgp 24.24.24.2** command.

```

user@R4> show route advertising-protocol bgp 24.24.24.2
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop      MED      Lclpref  AS path
* 24.24.24.0/24         Self        20             I
* 34.34.34.0/24         Self        20             I
* 44.44.44.44/32        Self        20             I
* 144.144.144.144/32    Self        20             I
* 192.168.4.1/32        Self        20             I

user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop      MED      Lclpref  AS path
* 24.24.24.0/24         Self        30             I

```

* 34.34.34.0/24	Self	30	I
* 44.44.44.44/32	Self	30	I
* 144.144.144.144/32	Self	10	I
* 192.168.4.1/32	Self	30	I

Meaning The MED column shows that Device R4 is sending the correct MED values to its two EBGp neighbors.

Example: Configuring the MED Using Communities

Set the multiple exit discriminator (MED) metric to 20 for all routes from a particular community.

```
[edit]
routing-options {
  router-id 10.0.0.1;
  autonomous-system 23;
}
policy-options {
  policy-statement from-otago {
    from community otago;
    then metric 20;
  }
  community otago members [56:2379 23:46944];
}
protocols {
  bgp {
    import from-otago;
    group 23 {
      type external;
      peer-as 56;
      neighbor 192.168.0.1 {
        traceoptions {
          file bgp-log-peer;
          flag packets;
        }
        log-updown;
      }
    }
  }
}
```

Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates

This example shows how to associate the multiple exit discriminator (MED) path attribute with the interior gateway protocol (IGP) metric, and configure a timer to delay update of the MED attribute.

- [Requirements on page 3237](#)
- [Overview on page 3237](#)
- [Configuration on page 3238](#)
- [Verification on page 3244](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

BGP can be configured to advertise the MED attribute for a route based on the IGP distance of its internal BGP (IBGP) route next-hop. The IGP metric enables internal routing to follow the shortest path according to the administrative setup. In some deployments, it might be ideal to communicate IGP shortest-path knowledge to external BGP (EBGP) peers in a neighboring autonomous system (AS). This allows those EBGP peers to forward traffic into your AS using the shortest paths possible.

Routes learned from an EBGP peer usually have a next hop on a directly connected interface, and thus the IGP value is equal to zero. Zero is the value advertised. The IGP metric is a nonzero value when a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the **multihop** command. In these scenarios, it might make sense to associate the MED value with the IGP metric by including the **metric-out minimum-igp** or **metric-out igp** option.

The drawback of associating the MED with the IGP metric is the risk of excessive route advertisements when there are IGP instabilities in the network. Configuring a delay for the MED update provides a mechanism to reduce route advertisements in such scenarios. The delay works by slowing down MED updates when the IGP metric for the next hop changes. The approach uses a timer to periodically advertise MED updates. When the timer expires, the MED attribute for routes with **metric-out igp delay-updates** configured is updated to the current IGP metric of the next hop. The BGP-enabled device sends out advertisements for routes for which the MED attribute has changed.

The **delay-updates** option identifies the BGP groups (or peers) for which the MED updates must be suppressed. The time for advertising MED updates is set to 10 minutes by default. You can increase the interval up to 600 minutes by including the **med-igp-update-interval** statement in the **routing-options** configuration.



NOTE: If you have nonstop active routing (NSR) enabled and a switchover occurs, the delayed MED updates might be advertised as soon as the switchover occurs.


When you configure the **metric-out igp** option, the IGP metric directly tracks the IGP cost to the IBGP peer. When the IGP cost goes down, so does the advertised MED value. Conversely, when the IGP cost goes up, the MED value goes up as well.

When you configure the **metric-out minimum-igp** option, the advertised MED value changes only when the IGP cost to the IBGP peer goes down. An increase in the IGP cost does not affect the MED value. The router monitors and remembers the lowest IGP cost until the routing process (rpd) is restarted. The BGP peer sends an update only if the MED is lower than the previously advertised value or another attribute associated with the route has changed, or if the BGP peer is responding to a refresh route request.

This example uses the **metric** statement in the OSPF configuration to demonstrate that when the IGP metric changes, the MED also changes after the configured delay interval. The OSPF metric can range from 1 through 65,535.

Figure 83 on page 3238 shows the sample topology.

Figure 83: Topology for Delaying the MED Update



In this example, the MED value advertised by Device R1 is associated with the IGP running in AS 1. The MED value advertised by Device R1 impacts the decisions of the neighboring AS (AS 2) when AS 2 is forwarding traffic into AS 1.

Configuration

- [Configuring Device R1 on page 3242](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1    set interfaces fe-1/2/0 unit 2 description R1->R2
              set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.1/30
              set interfaces fe-1/2/1 unit 7 description R1->R4
              set interfaces fe-1/2/1 unit 7 family inet address 172.16.0.1/30
              set interfaces lo0 unit 1 family inet address 192.168.0.1/32
              set protocols bgp group internal type internal
              set protocols bgp group internal local-address 192.168.0.1
              set protocols bgp group internal export send-direct
              set protocols bgp group internal neighbor 192.168.0.2
              set protocols bgp group internal neighbor 192.168.0.3
              set protocols bgp group external type external
              set protocols bgp group external metric-out igp delay-med-update
              set protocols bgp group external export send-direct
              set protocols bgp group external peer-as 2
              set protocols bgp group external neighbor 172.16.0.2
              set protocols ospf area 0.0.0.0 interface fe-1/2/0.2 metric 600
              set protocols ospf area 0.0.0.0 interface lo0.1 passive
              set policy-options policy-statement send-direct term 1 from protocol direct
              set policy-options policy-statement send-direct term 1 then accept
              set routing-options med-igp-update-interval 12
              set routing-options router-id 192.168.0.1
              set routing-options autonomous-system 1

Device R2    set interfaces fe-1/2/0 unit 1 description R2->R1
              set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
              set interfaces fe-1/2/1 unit 4 description R2->R3
              set interfaces fe-1/2/1 unit 4 family inet address 10.0.2.2/30
              set interfaces lo0 unit 2 family inet address 192.168.0.2/32
              set protocols bgp group internal type internal
              set protocols bgp group internal local-address 192.168.0.2
              set protocols bgp group internal export send-direct
              set protocols bgp group internal neighbor 192.168.0.1
              set protocols bgp group internal neighbor 192.168.0.3
              set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
              set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
              set protocols ospf area 0.0.0.0 interface lo0.2 passive
              set policy-options policy-statement send-direct term 1 from protocol direct
              set policy-options policy-statement send-direct term 1 then accept
              set routing-options router-id 192.168.0.2
              set routing-options autonomous-system 1

Device R3    set interfaces fe-1/2/0 unit 3 description R3->R2
              set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
              set interfaces fe-1/2/1 unit 5 description R3->R5
              set interfaces fe-1/2/1 unit 5 family inet address 172.16.0.5/30
              set interfaces lo0 unit 3 family inet address 192.168.0.3/32
              set protocols bgp group internal type internal
              set protocols bgp group internal local-address 192.168.0.3
              set protocols bgp group internal export send-direct
              set protocols bgp group internal neighbor 192.168.0.1
              set protocols bgp group internal neighbor 192.168.0.2
              set protocols bgp group external type external
              set protocols bgp group external export send-direct
              set protocols bgp group external peer-as 2
              set protocols bgp group external neighbor 172.16.0.6

```

```
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.3
set routing-options autonomous-system 1
```

Device R4

```
set interfaces fe-1/2/0 unit 8 description R4->R1
set interfaces fe-1/2/0 unit 8 family inet address 172.16.0.2/30
set interfaces fe-1/2/1 unit 9 description R4->R5
set interfaces fe-1/2/1 unit 9 family inet address 10.0.4.1/30
set interfaces fe-1/2/2 unit 13 description R4->R6
set interfaces fe-1/2/2 unit 13 family inet address 172.16.0.9/30
set interfaces lo0 unit 4 family inet address 192.168.0.4/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.4
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.5
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 172.16.0.10 peer-as 3
set protocols bgp group external neighbor 172.16.0.1 peer-as 1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.4
set routing-options autonomous-system 2
```

Device R5

```
set interfaces fe-1/2/0 unit 6 description R5->R3
set interfaces fe-1/2/0 unit 6 family inet address 172.16.0.6/30
set interfaces fe-1/2/1 unit 10 description R5->R4
set interfaces fe-1/2/1 unit 10 family inet address 10.0.4.2/30
set interfaces fe-1/2/2 unit 11 description R5->R8
set interfaces fe-1/2/2 unit 11 family inet address 172.16.0.13/30
set interfaces lo0 unit 5 family inet address 192.168.0.5/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.5
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.4
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 172.16.0.5 peer-as 1
set protocols bgp group external neighbor 172.16.0.14 peer-as 3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.10
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.5
set routing-options autonomous-system 2
```

Device R6

```
set interfaces fe-1/2/0 unit 14 description R6->R4
set interfaces fe-1/2/0 unit 14 family inet address 172.16.0.10/30
set interfaces fe-1/2/1 unit 15 description R6->R7
set interfaces fe-1/2/1 unit 15 family inet address 10.0.6.1/30
```

```

set interfaces lo0 unit 6 family inet address 192.168.0.6/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.6
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.7
set protocols bgp group internal neighbor 192.168.0.8
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.9 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.15
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.6
set routing-options autonomous-system 3

```

Device R7

```

set interfaces fe-1/2/0 unit 16 description R7->R6
set interfaces fe-1/2/0 unit 16 family inet address 10.0.6.2/30
set interfaces fe-1/2/1 unit 17 description R7->R8
set interfaces fe-1/2/1 unit 17 family inet address 10.0.7.2/30
set interfaces lo0 unit 7 family inet address 192.168.0.7/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.7
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.8
set protocols ospf area 0.0.0.0 interface fe-1/2/0.16
set protocols ospf area 0.0.0.0 interface fe-1/2/1.17
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.7
set routing-options autonomous-system 3

```

Device R8

```

set interfaces fe-1/2/0 unit 12 description R8->R5
set interfaces fe-1/2/0 unit 12 family inet address 172.16.0.14/30
set interfaces fe-1/2/1 unit 18 description R8->R7
set interfaces fe-1/2/1 unit 18 family inet address 10.0.7.1/30
set interfaces lo0 unit 8 family inet address 192.168.0.8/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.8
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.7
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.13 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.18
set protocols ospf area 0.0.0.0 interface lo0.8 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.8
set routing-options autonomous-system 3

```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 2]
user@R1# set description R1->R2
user@R1# set family inet address 10.0.0.1/30
```

```
[edit interfaces fe-1/2/1 unit 7]
user@R1# set description R1->R4
user@R1# set family inet address 172.16.0.1/30
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.0.1/32
```

2. Configure IBGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.0.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.0.2
user@R1# set neighbor 192.168.0.3
```

3. Configure EBGP.

```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set export send-direct
user@R1# set peer-as 2
user@R1# set neighbor 172.16.0.2
```

4. Associate the MED value with the IGP metric.

```
[edit protocols bgp group external]
user@R1# set metric-out igp delay-med-update
```

The default for the MED update is 10 minutes when you include the **delay-med-update** option. When you exclude the **delay-med-update** option, the MED update occurs immediately after the IGP metric changes.

5. (Optional) Configure the update interval for the MED update.

```
[edit routing-options]
user@R1# set med-igp-update-interval 12
```

You can configure the interval from 10 minutes through 600 minutes.

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.2 metric 600
user@R1# set interface lo0.1 passive
```

The **metric** statement is used here to demonstrate what happens when the IGP metric changes.

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

8. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set router-id 192.168.0.1
user@R1# set autonomous-system 1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 2 {
    description R1->R2;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-1/2/1 {
  unit 7 {
    description R1->R4;
    family inet {
      address 172.16.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
```

```
bgp {
  group internal {
    type internal;
    local-address 192.168.0.1;
    export send-direct;
    neighbor 192.168.0.2;
    neighbor 192.168.0.3;
  }
  group external {
    type external;
    metric-out igp delay-med-update;
    export send-direct;
    peer-as 2;
    neighbor 172.16.0.2;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.2 {
      metric 600;
    }
    interface lo0.1 {
      passive;
    }
  }
}
```

```
user@R1# show routing-options
med-igp-update-interval 12;
router-id 192.168.0.1;
autonomous-system 1;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration steps on the other devices in the topology, as needed for your network.

Verification

Confirm that the configuration is working properly.

- [Checking the BGP Advertisements on page 3244](#)
- [Verifying That the MED Value Changes When the OSPF Metric Changes on page 3245](#)
- [Testing the minimum-igp Setting on page 3245](#)

Checking the BGP Advertisements

Purpose Verify that Device R1 is advertising to Device R4 a BGP MED value that reflects the IGP metric.

Action From operational mode, enter the **show route advertising-protocol bgp** command.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix                Nexthop        MED      Lc1pref  AS path
* 10.0.0.0/30           Self           0         I         I
* 172.16.0.0/30         Self           0         I         I
```


* 172.16.0.4/30	Self	601	I
* 192.168.0.1/32	Self	0	I

Meaning The 601 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

Verifying That the MED Value Changes When the OSPF Metric Changes

Purpose Make sure that when you raise the OSPF metric to 700, the MED value is updated to reflect this change.

Action From configuration mode, enter the **set protocols ospf area 0 interface fe-1/2/0.2 metric 700** command.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 700
user@R1# commit
```

After waiting 12 minutes (the configured delay period), enter the **show route advertising-protocol bgp** command from operational mode.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix            Nexthop      MED      Lclpref    AS path
* 10.0.0.0/30       Self         0         I
* 172.16.0.0/30     Self         0         I
* 172.16.0.4/30     Self        701         I
* 192.168.0.1/32    Self         0         I
```

Meaning The 701 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

Testing the minimum-igp Setting

Purpose Change the configuration to use the **minimum-igp** statement instead of the **igp** statement. When you increase the OSPF metric, the MED value remains unchanged, but when you decrease the OSPF metric, the MED value reflects the new OSPF metric.

Action From configuration mode, delete the **igp** statement, add the **minimum-igp** statement, and increase the OSPF metric.

```
user@R1# delete protocols bgp group external metric-out igp
user@R1# set protocols bgp group external metric-out minimum-igp
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 800
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does not change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix            Nexthop      MED      Lclpref    AS path
* 10.0.0.0/30       Self         0         I
* 172.16.0.0/30     Self         0         I
* 172.16.0.4/30     Self        701         I
* 192.168.0.1/32    Self         0         I
```

From configuration mode, decrease the OSPF metric.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 20
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.0.0.0/30           Self             0                I
* 172.16.0.0/30         Self             0                I
* 172.16.0.4/30         Self             21               I
* 192.168.0.1/32        Self             0                I
```

Meaning When the **minimum-igp** statement is configured, the MED value changes only when a shorter path is available.

Related Documentation

- [Examples: Configuring External BGP Peering on page 3149](#)
- [BGP Configuration Overview](#)

Examples: Configuring BGP Local AS

- [Understanding the BGP Local AS Attribute on page 3246](#)
- [Example: Configuring a Local AS for EBGp Sessions on page 3251](#)
- [Example: Configuring a Private Local AS for EBGp Sessions on page 3261](#)

Understanding the BGP Local AS Attribute

When an Internet service provider (ISP) acquires a network that belongs to a different autonomous system (AS), there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. Sometimes customers do not want to or are not immediately able to modify their peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a *local AS*.

Using a local AS number permits the routing devices in an acquired network to appear to belong to the former AS.

For example, ISP A, with an AS of 200, acquires ISP B, with an AS of 250. ISP B has a customer, ISP C, that does not want to change its configuration. After ISP B becomes part of ISP A, a local AS number of 250 is configured for use in EBGp peer sessions with ISP C. Consequently, the local AS number of 250 is either prepended before or used instead of the global AS number of 200 in the AS path used to export routes to direct external peers in ISP C.

If the route is received from an internal BGP (IBGP) peer, the AS path includes the local AS number prepended before the global AS number.

The local AS number is used instead of the global AS number if the route is an external route, such as a static route or an interior gateway protocol (IGP) route that is imported into BGP. If the route is external and you want the global AS number to be included in the AS path, you can apply a routing policy that uses **as-path-expand** or **as-path-prepend**. Use the **as-path-expand** policy action to place the global AS number behind the local AS number. Use the **as-path-prepend** policy action to place the global AS number in front of the local AS number.

For example:

```

user@R2# show policy-options
policy-statement prepend-global {
  term 1 {
    from protocol static;
    then {
      as-path-prepend 200; # or use as-path-expand
      accept;
    }
  }
}

user@R2# show protocols bgp
group ext {
  export prepend-global;
  type external;
  local-as 250;
  neighbor 10.0.0.1 {
    peer-as 100;
  }
  neighbor 10.1.0.2 {
    peer-as 300;
  }
}

user@R2# show routing-options
static {
  route 1.1.1.1/32 next-hop 10.0.0.1;
}
autonomous-system 200;

user@R3# run show route 1.1.1.1 protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[BGP/170] 00:05:11, localpref 100
                   AS path: 200 250 I, validation-state: unverified
                   > to 10.1.0.1 via lt-1/2/0.4

```

In a Layer 3 VPN scenario, in which a provider edge (PE) device uses external BGP (EBGP) to peer with a customer edge (CE) device, the **local-as** statement behaves differently than in the non-VPN scenario. In the VPN scenario, the global AS number defined in the master instance is prepended to the AS path by default. To override this behavior, you can configure the **no-prepend-global-as** in the routing-instance BGP configuration on the PE device, as shown here:

```

user@R2# show routing-instances

```

```
red {  
  instance-type vrf;  
  interface fe-1/2/0.2;  
  route-distinguisher 2:1;  
  vrf-target target:2:1;  
  protocols {  
    bgp {  
      group toR1 {  
        type external;  
        peer-as 1;  
        local-as 200 no-prepend-global-as;  
        neighbor 10.1.1.1;  
      }  
    }  
  }  
}
```

The Junos operating system (Junos OS) implementation of the local AS attribute supports the following options:


- **Local AS with private option**—When you use the **private** option, the local AS is used during the establishment of the BGP session with an EBGP neighbor but is hidden in the AS path sent to other EBGP peers. Only the global AS is included in the AS path sent to external peers.

The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

For example, in [Figure 84 on page 3249](#), Router 1 and Router 2 are in AS 64496, Router 4 is in AS 64511, and Router 3 is in AS 64510. Router 2 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Router 3 still peers with Router 2 using its former AS (64497), Router 2 needs to be configured with a local AS of 64497 in order to maintain peering with Router 3. Configuring a local AS of 64497 permits Router 2 to add AS 64497 when advertising routes to Router 3. Router 3 sees an AS path of 64497 64496 for the prefix 10/8.

Figure 84: Local AS Configuration



To prevent Router 2 from adding the local AS number in its announcements to other peers, use the **local-as 64497 private** statement. This statement configures Router 2 to not include local AS 64497 when announcing routes to Router 1 and to Router 4. In this case, Router 4 sees an AS path of 64496 64510 for the prefix 10.222/16.

- **Local AS with alias option**—In Junos OS Release 9.5 and later, you can configure a local AS as an alias. During the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. If the local AS is used to connect with the EBGP neighbor, then only the local AS is prepended to the AS path when the BGP peer session is established. If the global AS is used to connect with the EBGP neighbor, then only the global AS is prepended to the AS path when the BGP peer session is established. The use of the **alias** option also means that

the local AS is not prepended to the AS path for any routes learned from that EBGp neighbor. Therefore, the local AS remains hidden from other external peers.

Configuring a local AS with the **alias** option is especially useful when you are migrating the routing devices in an acquired network to the new AS. During the migration process, some routing devices might be configured with the new AS while others remain configured with the former AS. For example, it is good practice to start by first migrating to the new AS any routing devices that function as route reflectors. However, as you migrate the route reflector clients incrementally, each route reflector has to peer with routing devices configured with the former AS, as well as peer with routing devices configured with the new AS. To establish local peer sessions, it can be useful for the BGP peers in the network to use both the local AS and the global AS. At the same time, you want to hide this local AS from external peers and use only the global AS in the AS path when exporting routes to another AS. In this kind of situation, configure the **alias** option.

Include the **alias** option to configure the local AS as an alias to the global AS configured at the **[edit routing-options]** hierarchy level. When you configure a local AS as an alias, during the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. The local AS is prepended to the AS path only when the peer session with an EBGp neighbor is established using that local AS. The local AS is hidden in the AS path sent to any other external peers. Only the global AS is prepended to the AS path when the BGP session is established using the global AS.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

- **Local AS with option not to prepend the global AS**—In Junos OS Release 9.6 and later, you can configure a local AS with the option not to prepend the global AS. Only the local AS is included in the AS path sent to external peers.

Use the **no-prepend-global-as** option when you want to strip the global AS number from outbound BGP updates in a virtual private network (VPN) scenario. This option is useful in a VPN scenario in which you want to hide the global AS from the VPN.

Include the **no-prepend-global-as** option to have the global AS configured at the **[edit routing-options]** hierarchy level removed from the AS path sent to external peers. When you use this option, only the local AS is included in the AS path for the routes sent to a customer edge (CE) device.

- **Number of loops option**—The local AS feature also supports specifying the number of times that detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

For the **loops number** statement, you can configure 1 through 10.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the `local-as` statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the `local-as` statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Example: Configuring a Local AS for EBGP Sessions

This example shows how to configure a local autonomous system (AS) for a BGP peer so that both the global AS and the local AS are used in BGP inbound and outbound updates.

- [Requirements on page 3251](#)
- [Overview on page 3251](#)
- [Configuration on page 3252](#)
- [Verification on page 3258](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Use the `local-as` statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The `local-as` statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.


This example shows how to use the `local-as` statement to configure a local AS. The `local-as` statement is supported for BGP at the global, group, and neighbor hierarchy levels.

When you configure the `local-as` statement, you must specify an AS number. You can specify a number from 1 through 4,294,967,295 in plain-number format. In Junos OS Release 9.1 and later, the range for AS numbers is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value from 0.0 through 65535.65535 in AS-dot notation format. Junos

OS continues to support 2-byte AS numbers. The 2-byte AS number range is 1 through 65,535 (this is a subset of the 4-byte range).

Figure 85 on page 3252 shows the sample topology.

Figure 85: Topology for Configuring the Local AS



In this example, Device R2 formerly belonged to AS 250 and now is in AS 200. Device R1 and Device R3 are configured to peer with AS 250 instead of with the new AS number (AS 200). Device R2 has the new AS number configured with the **autonomous-system 200** statement. To enable the peering sessions to work, the **local-as 250** statement is added in the BGP configuration. Because **local-as 250** is configured, Device R2 includes both the global AS (200) and the local AS (250) in its BGP inbound and outbound updates.

Configuration

- [Configuring Device R1 on page 3253](#)
- [Configuring Device R2 on page 3255](#)
- [Configuring Device R3 on page 3257](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 192.168.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 250
set protocols bgp group ext neighbor 10.0.0.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.1.0.0/30 next-hop 10.0.0.2
set routing-options autonomous-system 100
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 3 family inet address 10.1.0.1/30
set interfaces lo0 unit 2 family inet address 192.168.0.2/32
set protocols bgp group ext type external
```



```

set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext local-as 250
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options autonomous-system 200

```

Device R3

```

set interfaces fe-1/2/0 unit 4 family inet address 10.1.0.2/30
set interfaces lo0 unit 3 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 250
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.0.0.0/30 next-hop 10.1.0.1
set routing-options autonomous-system 300

```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.


```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 192.168.0.1/32

```
2. Configure external BGP (EBGP).


```

[edit protocols bgp group ext]
user@R1# set type external
user@R1# set export send-direct
user@R1# set export send-static
user@R1# set peer-as 250
user@R1# set neighbor 10.0.0.2

```
3. Configure the routing policy.


```

[edit policy-options]
user@R1# set policy-statement send-direct term 1 from protocol direct
user@R1# set policy-statement send-direct term 1 then accept
user@R1# set policy-statement send-static term 1 from protocol static
user@R1# set policy-statement send-static term 1 then accept

```

4. Configure a static route to the remote network between Device R2 and Device R3.

```
[edit routing-options]
user@R1# set static route 10.1.0.0/30 next-hop 10.0.0.2
```

5. Configure the global AS number.

```
[edit routing-options]
user@R1# set autonomous-system 100
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group ext {
    type external;
    export [ send-direct send-static ];
    peer-as 250;
    neighbor 10.0.0.2;
  }
}

user@R1# show routing-options
static {
```

```

    route 10.1.0.0/30 next-hop 10.0.0.2;
}
autonomous-system 100;

```

When you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.


```

[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 3 family inet address 10.1.0.1/30

user@R2# set lo0 unit 2 family inet address 192.168.0.2/32

```
2. Configure EBGP.


```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300

```
3. Configure the local autonomous system (AS) number.


```

[edit protocols bgp group ext]
user@R2# set local-as 250

```
4. Configure the global AS number.


```

[edit routing-options]
user@R2# set autonomous-system 200

```
5. Configure the routing policy.


```

[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
  unit 2 {

```

```
        family inet {
            address 10.0.0.2/30;
        }
    }
}
fe-1/2/1 {
    unit 3 {
        family inet {
            address 10.1.0.1/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.0.2/32;
        }
    }
}

user@R2# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}
policy-statement send-static {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R2# show protocols
bgp {
    group ext {
        type external;
        export [ send-direct send-static ];
        local-as 250;
        neighbor 10.0.0.1 {
            peer-as 100;
        }
        neighbor 10.1.0.2 {
            peer-as 300;
        }
    }
}

user@R2# show routing-options
autonomous-system 200;
```

When you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 4 family inet address 10.1.0.2/30

user@R3# set lo0 unit 3 family inet address 192.168.0.3/32
```
2. Configure EBGP.

```
[edit protocols bgp group ext]
user@R3# set type external
user@R3# set export send-direct
user@R3# set export send-static
user@R3# set peer-as 250
user@R3# set neighbor 10.1.0.1
```
3. Configure the global autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 300
```
4. Configure a static route to the remote network between Device R1 and Device R2.

```
[edit routing-options]
user@R3# set static route 10.0.0.0/30 next-hop 10.1.0.1
```
5. Configure the routing policy.

```
[edit policy-options]
user@R3# set policy-statement send-direct term 1 from protocol direct
user@R3# set policy-statement send-direct term 1 then accept
user@R3# set policy-statement send-static term 1 from protocol static
user@R3# set policy-statement send-static term 1 then accept
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 4 {
    family inet {
      address 10.1.0.2/30;
    }
  }
}
lo0 {
  unit 3 {
```

```
family inet {
    address 192.168.0.3/32;
}
}

user@R3# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}
policy-statement send-static {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show protocols
bgp {
    group ext {
        type external;
        export [ send-direct send-static ];
        peer-as 250;
        neighbor 10.1.0.1;
    }
}

user@R3# show routing-options
static {
    route 10.0.0.0/30 next-hop 10.1.0.1;
}
autonomous-system 300;
```

When you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Local and Global AS Settings on page 3258](#)
- [Checking the BGP Peering Sessions on page 3260](#)
- [Verifying the BGP AS Paths on page 3260](#)

Checking the Local and Global AS Settings

Purpose Make sure that Device R2 has the local and global AS settings configured.

Action From operational mode, enter the **show bgp neighbors** command.

```
user@R2> show bgp neighbors
Peer: 10.0.0.1+179 AS 100      Local: 10.0.0.2+61036 AS 250
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
```

```

Export: [ send-direct send-static ]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.1      Local ID: 192.168.0.2      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
BFD: disabled, down
Local Interface: fe-1/2/0.2
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 100)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        3
  Accepted prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      4
Last traffic (seconds): Received 6    Sent 14    Checked 47
Input messages: Total 258    Updates 3    Refreshes 0    Octets 4969
Output messages: Total 258    Updates 2    Refreshes 0    Octets 5037
Output Queue[0]: 0

Peer: 10.1.0.2+179 AS 300      Local: 10.1.0.1+52296 AS 250
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ send-direct send-static ]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.3      Local ID: 192.168.0.2      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 1
BFD: disabled, down
Local Interface: fe-1/2/1.3
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 300)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        3
  Accepted prefixes:        2

```

```

    Suppressed due to damping:    0
    Advertised prefixes:          4
    Last traffic (seconds): Received 19    Sent 26    Checked 9
    Input messages:  Total 256    Updates 3      Refreshes 0    Octets 4931
    Output messages: Total 256    Updates 2      Refreshes 0    Octets 4999
    Output Queue[0]: 0

```

Meaning The Local AS: 250 and Local System AS: 200 output shows that Device R2 has the expected settings. Additionally, the output shows that the options list includes LocalAS.

Checking the BGP Peering Sessions

Purpose Ensure that the sessions are established and that the local AS number 250 is displayed.

Action From operational mode, enter the **show bgp summary** command.

```

user@R1> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          4          2          0          0          0      0        0
Peer           AS        InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.2        250        232      233        0        4    1:42:37
2/4/4/0        0/0/0/0

```

```

user@R3> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          4          2          0          0          0      0        0
Peer           AS        InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.0.1        250        235      236        0        4    1:44:25
2/4/4/0        0/0/0/0

```

Meaning Device R1 and Device R3 appear to be peering with a device in AS 250, even though Device R2 is actually in AS 200.

Verifying the BGP AS Paths

Purpose Make sure that the routes are in the routing tables and that the AS paths show the local AS number 250.

Action From configuration mode, enter the **set route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30      [BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1
10.1.0.0/30      [BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.0.2/32   *[BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1

```



```

192.168.0.3/32      *[BGP/170] 01:46:40, localpref 100
                   AS path: 250 300 I
                   > to 10.0.0.2 via fe-1/2/0.1

user@R3> show route protocol bgp

inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30        [BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4
10.1.0.0/30        [BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4
192.168.0.1/32     *[BGP/170] 01:47:10, localpref 100
                   AS path: 250 100 I
                   > to 10.1.0.1 via fe-1/2/0.4
192.168.0.2/32     *[BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4

```

Meaning The output shows that Device R1 and Device R3 appear to have routes with AS paths that include AS 250, even though Device R2 is actually in AS 200.

Example: Configuring a Private Local AS for EBGp Sessions

This example shows how to configure a private local autonomous system (AS) number. The local AS is considered to be private because it is advertised to peers that use the local AS number for peering, but is hidden in the announcements to peers that can use the global AS number for peering.

- [Requirements on page 3261](#)
- [Overview on page 3261](#)
- [Configuration on page 3262](#)
- [Verification on page 3265](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Use the **local-as** statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The **local-as** statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.

When you use the **private** option, the local AS is used during the establishment of the BGP session with an external BGP (EBGP) neighbor, but is hidden in the AS path sent to other EBGp peers. Only the global AS is included in the AS path sent to external peers.


The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its

peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor, but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

Figure 86 on page 3262 shows the sample topology.

Figure 86: Topology for Configuring a Private Local AS



Device R1 is in AS 64496. Device R2 is in AS 64510. Device R3 is in AS 64511. Device R4 is in AS 64512. Device R1 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Device R3 still peers with Device R1, using its former AS, 64497, Device R1 needs to be configured with a local AS of 64497 in order to maintain peering with Device R3. Configuring a local AS of 64497 permits Device R1 to add AS 64497 when advertising routes to Device R3. Device R3 sees an AS path of 64497 64496 for the prefix 10.1.1.2/32, which is Device R2's loopback interface. Device R4, which is behind Device R3, sees an AS path of 64511 64497 64496 64510 to Device R2's loopback interface. To prevent Device R1 from adding the local AS number in its announcements to other peers, this example includes the **local-as 64497 private** statement. The **private** option configures Device R1 to not include the local AS 64497 when announcing routes to Device R2. Device R2 sees an AS path of 64496 64511 to Device R3 and an AS path of 64496 64511 64512 to Device R4. The **private** option in Device R1's configuration causes the AS number 64497 to be missing from the AS paths that Device R1 readvertises to Device R2.

Device R2 is hiding the private local AS from all the routers, except Device R3. The **private** option applies to the routes that Device R1 receives (learns) from Device R3 and that Device R1, in turn, readvertises to other routers. When these routes learned from Device R3 are readvertised by Device R1 to Device R2, the private local AS is missing from the AS path advertised to Device R2.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set interfaces fe-1/2/0 unit 3 family inet address 192.168.1.1/24
set interfaces fe-1/2/1 unit 5 family inet address 192.168.10.1/24
set interfaces lo0 unit 2 family inet address 10.1.1.1/32
set protocols bgp group external-AS64511 type external
set protocols bgp group external-AS64511 peer-as 64511
set protocols bgp group external-AS64511 local-as 64497
set protocols bgp group external-AS64511 local-as private
set protocols bgp group external-AS64511 neighbor 192.168.1.2
set protocols bgp group external-AS64510 type external
set protocols bgp group external-AS64510 peer-as 64510
set protocols bgp group external-AS64510 neighbor 192.168.10.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64496

```

Device R2

```

set interfaces fe-1/2/0 unit 6 family inet address 192.168.10.2/24
set interfaces lo0 unit 3 family inet address 10.1.1.2/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 64496
set protocols bgp group external neighbor 192.168.10.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64510

```

Device R3

```

set interfaces fe-1/2/0 unit 4 family inet address 192.168.1.2/24
set interfaces fe-1/2/1 unit 7 family inet address 192.168.5.1/24
set interfaces lo0 unit 4 family inet address 10.1.1.3/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 192.168.1.1 peer-as 64497
set protocols bgp group external neighbor 192.168.5.2 peer-as 64512
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64511

```

Device R4

```

set interfaces fe-1/2/0 unit 8 family inet address 192.168.5.2/24
set interfaces lo0 unit 5 family inet address 10.1.1.4/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 64511
set protocols bgp group external neighbor 192.168.5.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64512

```

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 192.168.1.1/24

[edit interfaces fe-1/2/1 unit 5]
user@R1# set family inet address 192.168.10.1/24

[edit interfaces lo0 unit 2]
user@R1# set family inet address 10.1.1.1/32
```
2. Configure the EBGP peering session with Device R2.

```
[edit protocols bgp group external-AS64510]
user@R1# set type external
user@R1# set peer-as 64510
user@R1# set neighbor 192.168.10.2
```
3. Configure the EBGP peering session with Device R3.

```
[edit protocols bgp group external-AS64511]
user@R1# set type external
user@R1# set peer-as 64511
user@R1# set local-as 64497
user@R1# set local-as private
user@R1# set neighbor 192.168.1.2
```
4. Configure the routing policy.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```
5. Configure the global autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 64496
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 192.168.1.1/24;
    }
  }
}
```

```

    }
  }
  fe-1/2/1 {
    unit 5 {
      family inet {
        address 192.168.10.1/24;
      }
    }
  }
  lo0 {
    unit 2 {
      family inet {
        address 10.1.1.1/32;
      }
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group external-AS64511 {
    type external;
    peer-as 64511;
    local-as 64497 private;
    neighbor 192.168.1.2;
  }
  group external-AS64510 {
    type external;
    peer-as 64510;
    neighbor 192.168.10.2;
  }
}

user@R1# show routing-options
autonomous-system 64496;

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the configuration as needed for the other devices in the topology.

Verification

Confirm that the configuration is working properly.

- [Checking Device R2's AS Paths on page 3266](#)
- [Checking Device R3's AS Paths on page 3266](#)

Checking Device R2's AS Paths

Purpose Make sure that Device R2 does not have AS 64497 in its AS paths to Device R3 and Device R4.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@R2> show route protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.3/32      *[BGP/170] 01:33:11, localpref 100
                  AS path: 64496 64511 I
                  > to 192.168.10.1 via fe-1/2/0.6
10.1.1.4/32      *[BGP/170] 01:33:11, localpref 100
                  AS path: 64496 64511 64512 I
                  > to 192.168.10.1 via fe-1/2/0.6
192.168.5.0/24  *[BGP/170] 01:49:15, localpref 100
                  AS path: 64496 64511 I
                  > to 192.168.10.1 via fe-1/2/0.6
```

Meaning Device R2's AS paths do not include AS 64497.

Checking Device R3's AS Paths

Purpose Make sure that Device R3 does not have AS 64497 in its AS path to Device R4.

Action From operational mode, enter the **show route protocol bgp** command.

```
user@R3> show route protocol bgp
inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2/32      *[BGP/170] 01:35:11, localpref 100
                  AS path: 64497 64496 64510 I
                  > to 192.168.1.1 via fe-1/2/0.4
10.1.1.4/32      *[BGP/170] 01:35:11, localpref 100
                  AS path: 64512 I
                  > to 192.168.5.2 via fe-1/2/1.7
192.168.5.0/24  [BGP/170] 01:51:15, localpref 100
                  AS path: 64512 I
                  > to 192.168.5.2 via fe-1/2/1.7
```

Meaning Device R3's route to Device R2 (prefix 10.1.1.2) includes both the local and the global AS configured on Device R1 (64497 and 64496, respectively).

Related Documentation

- [Examples: Configuring External BGP Peering on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring the Accumulated IGP Attribute for BGP

- [Understanding the Accumulated IGP Attribute for BGP on page 3267](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3267](#)

Understanding the Accumulated IGP Attribute for BGP

The interior gateway protocols (IGPs) are designed to handle routing within a single domain or an autonomous system (AS). Each link is assigned a particular value called a metric. The distance between the two nodes is calculated as a sum of all the metric values of links along the path. The IGP selects the shortest path between two nodes based on distance.

BGP is designed to provide routing over a large number of independent ASs with limited or no coordination among respective administrations. BGP does not use metrics in the path selection decisions.

The accumulated IGP (AIGP) metric attribute for BGP enables deployment in which a single administration can run several contiguous BGP ASs. Such deployments allow BGP to make routing decisions based on the IGP metric. In such networks, it is possible for BGP to select paths based on metrics as is done by IGPs. In this case, BGP chooses the shortest path between two nodes, even though the nodes might be in two different ASs.

The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. The Juniper Networks® Junos® operating system (Junos OS) currently supports the AIGP attribute for two BGP address families, **family inet labeled-unicast** and **family inet6 labeled-unicast**.

AIGP impacts the BGP best-route decision process. The AIGP attribute preference rule is applied after the local-preference rule. The AIGP distance is compared to break a tie. The BGP best-route decision process also impacts the way the interior cost rule is applied if the resolving next hop has an AIGP attribute. Without AIGP enabled, the interior cost of a route is based on the calculation of the metric to the next hop for the route. With AIGP enabled, the resolving AIGP distance is added to the interior cost.

The AIGP attribute is an optional non-transitive BGP path attribute and is specified in Internet draft draft-ietf-idr-aigp-06, *The Accumulated IGP Metric Attribute for BGP*.

Example: Configuring the Accumulated IGP Attribute for BGP

This example shows how to configure the accumulated IGP (AIGP) metric attribute for BGP.

- [Requirements on page 3267](#)
- [Overview on page 3268](#)
- [Configuration on page 3269](#)
- [Verification on page 3299](#)

Requirements

This example uses the following hardware and software components:

- Seven BGP-speaking devices.
- Junos OS Release 12.1 or later.

Overview

The AIGP attribute enables deployments in which a single administration can run several contiguous BGP autonomous systems (ASs). Such deployments allow BGP to make routing decisions based on the IGP metric. With AIGP enabled, BGP can select paths based on IGP metrics. This enables BGP to choose the shortest path between two nodes, even though the nodes might be in different ASs. The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. This example shows AIGP configured with MPLS label-switched paths.

To enable AIGP, you include the **aigp** statement in the BGP configuration on a protocol family basis. Configuring AIGP on a particular family enables sending and receiving of the AIGP attribute on that family. By default, AIGP is disabled. An AIGP-disabled neighbor does not send an AIGP attribute and silently discards a received AIGP attribute.

Junos OS supports AIGP for **family inet labeled-unicast** and **family inet6 labeled-unicast**. The **aigp** statement can be configured for a given family at the global BGP, group, or neighbor level.

By default, the value of the AIGP attribute for a local prefix is zero. An AIGP-enabled neighbor can originate an AIGP attribute for a given prefix by export policy, using the **aigp-originate** policy action. The value of the AIGP attribute reflects the IGP distance to the prefix. Alternatively, you can specify a value, by using the **aigp-originate distance distance** policy action. The configurable range is 0 through 4,294,967,295. Only one node needs to originate an AIGP attribute. The AIGP attribute is retained and readvertised if the neighbors are AIGP enabled with the **aigp** statement in the BGP configuration.

The policy action to originate the AIGP attribute has the following requirements:

- Neighbor must be AIGP enabled.
- Policy must be applied as an export policy.
- Prefix must have no current AIGP attribute.
- Prefix must export with next-hop self.
- Prefix must reside within the AIGP domain. Typically, a loopback IP address is the prefix to originate.


The policy is ignored if these requirements are not met.

Topology Diagram

[Figure 87 on page 3269](#) shows the topology used in this example. OSPF is used as the interior gateway protocol (IGP). Internal BGP (IBGP) is configured between Device PE1 and Device PE4. External BGP (EBGP) is configured between Device PE7 and Device PE1, between Device PE4 and Device PE3, and between Device PE4 and Device PE2. Devices PE4, PE2, and PE3 are configured for multihop. Device PE4 selects a path based on the AIGP value and then readvertises the AIGP value based on the AIGP and policy configuration. Device PE1 readvertises the AIGP value to Device PE7, which is in another administrative domain. Every device has two loopback interface addresses: 10.9.9.x is used for BGP peering and the router ID, and 10.100.1.x is used for the BGP next hop.

The network between Device PE1 and PE3 has IBGP peering and multiple OSPF areas. The external link to Device PE7 is configured to show that the AIGP attribute is readadvertised to a neighbor outside of the administrative domain, if that neighbor is AIGP enabled.

Figure 87: Advertisement of Multiple Paths in BGP



For origination of an AIGP attribute, the BGP next hop is required to be itself. If the BGP next hop remains unchanged, the received AIGP attribute is readadvertised, as is, to another AIGP neighbor. If the next hop changes, the received AIGP attribute is readadvertised with an increased value to another AIGP neighbor. The increase in value reflects the IGP distance to the previous BGP next hop. To demonstrate, this example uses loopback interface addresses for Device PE4's EBGP peering sessions with Device PE2 and Device PE3. Multihop is enabled on these sessions so that a recursive lookup is performed to determine the point-to-point interface. Because the next hop changes, the IGP distance is added to the AIGP distance.

Configuration

- [Configuring Device P1 on page 3275](#)
- [Configuring Device P2 on page 3278](#)
- [Configuring Device PE4 on page 3281](#)
- [Configuring Device PE1 on page 3286](#)
- [Configuring Device PE2 on page 3290](#)
- [Configuring Device PE3 on page 3294](#)
- [Configuring Device PE7 on page 3297](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device P1 set interfaces fe-1/2/0 unit 1 description P1-to-PE1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 1 family mpls
set interfaces fe-1/2/1 unit 4 description P1-to-P2
```

```
set interfaces fe-1/2/1 unit 4 family inet address 10.0.0.29/30
set interfaces fe-1/2/1 unit 4 family mpls
set interfaces fe-1/2/2 unit 8 description P1-to-PE4
set interfaces fe-1/2/2 unit 8 family inet address 10.0.0.17/30
set interfaces fe-1/2/2 unit 8 family mpls
set interfaces lo0 unit 3 family inet address 10.9.9.2/32
set interfaces lo0 unit 3 family inet address 10.100.1.2/32
set protocols rsvp interface fe-1/2/0.1
set protocols rsvp interface fe-1/2/2.8
set protocols rsvp interface fe-1/2/1.4
set protocols mpls label-switched-path P1-to-P2 to 10.9.9.3
set protocols mpls label-switched-path P1-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P1-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.1
set protocols mpls interface fe-1/2/2.8
set protocols mpls interface fe-1/2/1.4
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.2
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.1 interface fe-1/2/0.1 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.4 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/2.8 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.2 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.2 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.2 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.2 metric 1
set routing-options router-id 10.9.9.2
set routing-options autonomous-system 13979
```

Device P2

```
set interfaces fe-1/2/0 unit 3 description P2-to-PE1
set interfaces fe-1/2/0 unit 3 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 3 family mpls
set interfaces fe-1/2/1 unit 5 description P2-to-P1
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.30/30
set interfaces fe-1/2/1 unit 5 family mpls
set interfaces fe-1/2/2 unit 6 description P2-to-PE4
set interfaces fe-1/2/2 unit 6 family inet address 10.0.0.13/30
set interfaces fe-1/2/2 unit 6 family mpls
set interfaces lo0 unit 5 family inet address 10.9.9.3/32
set interfaces lo0 unit 5 family inet address 10.100.1.3/32
set protocols rsvp interface fe-1/2/1.5
set protocols rsvp interface fe-1/2/2.6
set protocols rsvp interface fe-1/2/0.3
set protocols mpls label-switched-path P2-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P2-to-P1 to 10.9.9.2
set protocols mpls label-switched-path P2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/1.5
set protocols mpls interface fe-1/2/2.6
set protocols mpls interface fe-1/2/0.3
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.3
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
```

```

set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.0 interface fe-1/2/2.6 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.3 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.3 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.3 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.3 metric 1
set routing-options router-id 10.9.9.3
set routing-options autonomous-system 13979

```

Device PE4

```

set interfaces fe-1/2/0 unit 7 description PE4-to-P2
set interfaces fe-1/2/0 unit 7 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 7 family mpls
set interfaces fe-1/2/1 unit 9 description PE4-to-P1
set interfaces fe-1/2/1 unit 9 family inet address 10.0.0.18/30
set interfaces fe-1/2/1 unit 9 family mpls
set interfaces fe-1/2/2 unit 10 description PE4-to-PE2
set interfaces fe-1/2/2 unit 10 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 10 family mpls
set interfaces fe-1/0/2 unit 12 description PE4-to-PE3
set interfaces fe-1/0/2 unit 12 family inet address 10.0.0.25/30
set interfaces fe-1/0/2 unit 12 family mpls
set interfaces lo0 unit 7 family inet address 10.9.9.4/32
set interfaces lo0 unit 7 family inet address 10.100.1.4/32
set protocols rsvp interface fe-1/2/0.7
set protocols rsvp interface fe-1/2/1.9
set protocols rsvp interface fe-1/2/2.10
set protocols rsvp interface fe-1/0/2.12
set protocols mpls label-switched-path PE4-to-PE2 to 10.9.9.5
set protocols mpls label-switched-path PE4-to-PE3 to 10.9.9.6
set protocols mpls label-switched-path PE4-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE4-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.7
set protocols mpls interface fe-1/2/1.9
set protocols mpls interface fe-1/2/2.10
set protocols mpls interface fe-1/0/2.12
set protocols bgp export next-hop
set protocols bgp export aigp
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.4
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.4
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external peer-as 7018
set protocols bgp group external neighbor 10.9.9.5
set protocols bgp group external neighbor 10.9.9.6
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.7 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.4 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.4 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.4 passive

```

```
set protocols ospf area 0.0.0.0 interface 10.100.1.4 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/2.10 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/0/2.12 metric 1
set policy-options policy-statement aigp term 10 from protocol static
set policy-options policy-statement aigp term 10 from route-filter 44.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 200
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.4
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.4/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.4/32
  exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 20 then accept
set routing-options static route 44.0.0.0/24 discard
set routing-options router-id 10.9.9.4
set routing-options autonomous-system 13979
```

Device PE1

```
set interfaces fe-1/2/0 unit 0 description PE1-to-P1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 2 description PE1-to-P2
set interfaces fe-1/2/1 unit 2 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 2 family mpls
set interfaces fe-1/2/2 unit 14 description PE1-to-PE7
set interfaces fe-1/2/2 unit 14 family inet address 10.0.0.9/30
set interfaces lo0 unit 1 family inet address 10.9.9.1/32
set interfaces lo0 unit 1 family inet address 10.100.1.1/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.2
set protocols rsvp interface fe-1/2/2.14
set protocols mpls label-switched-path PE1-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE1-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.2
set protocols mpls interface fe-1/2/2.14
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.1
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal export SET_EXPORT_ROUTES
set protocols bgp group internal vpn-apply-export
set protocols bgp group internal neighbor 10.9.9.4
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 7019
set protocols bgp group external neighbor 10.0.0.10
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.2 metric 1
set protocols ospf area 0.0.0.1 interface 10.9.9.1 passive
set protocols ospf area 0.0.0.1 interface 10.9.9.1 metric 1
```

```

set protocols ospf area 0.0.0.1 interface 10.100.1.1 passive
set protocols ospf area 0.0.0.1 interface 10.100.1.1 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.1
set routing-options autonomous-system 13979

```

Device PE2

```

set interfaces fe-1/2/0 unit 11 description PE2-to-PE4
set interfaces fe-1/2/0 unit 11 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 11 family mpls
set interfaces lo0 unit 9 family inet address 10.9.9.5/32 primary
set interfaces lo0 unit 9 family inet address 10.100.1.5/32
set protocols rsvp interface fe-1/2/0.11
set protocols mpls label-switched-path PE2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.11
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.5
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.2 interface 10.9.9.5 passive
set protocols ospf area 0.0.0.2 interface 10.9.9.5 metric 1
set protocols ospf area 0.0.0.2 interface 10.100.1.5 passive
set protocols ospf area 0.0.0.2 interface 10.100.1.5 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/0.11 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.5
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement aigp term 10 from route-filter 55.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 20
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement aigp term 20 from route-filter 99.0.0.0/24 exact
set policy-options policy-statement aigp term 20 then aigp-originate distance 30
set policy-options policy-statement aigp term 20 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 20 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.5
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.5/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.5/32
  exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.5
set policy-options policy-statement next-hop term 20 then accept

```

```
set routing-options static route 99.0.0.0/24 discard
set routing-options static route 55.0.0.0/24 discard
set routing-options router-id 10.9.9.5
set routing-options autonomous-system 7018
```

Device PE3

```
set interfaces fe-1/2/0 unit 13 description PE3-to-PE4
set interfaces fe-1/2/0 unit 13 family inet address 10.0.0.26/30
set interfaces fe-1/2/0 unit 13 family mpls
set interfaces lo0 unit 11 family inet address 10.9.9.6/32
set interfaces lo0 unit 11 family inet address 10.100.1.6/32
set protocols rsvp interface fe-1/2/0.13
set protocols mpls label-switched-path PE3-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.13
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.6
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.3 interface 10.9.9.6 passive
set protocols ospf area 0.0.0.3 interface 10.9.9.6 metric 1
set protocols ospf area 0.0.0.3 interface 10.100.1.6 passive
set protocols ospf area 0.0.0.3 interface 10.100.1.6 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/2/0.13 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.6
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.6/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.6/32
  exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 20 then accept
set routing-options router-id 10.9.9.6
set routing-options autonomous-system 7018
```

Device PE7

```
set interfaces fe-1/2/0 unit 15 description PE7-to-PE1
set interfaces fe-1/2/0 unit 15 family inet address 10.0.0.10/30
set interfaces lo0 unit 13 family inet address 10.9.9.7/32
set interfaces lo0 unit 13 family inet address 10.100.1.7/32
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.0.0.9
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
```

```

set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.7
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.7
set routing-options autonomous-system 7019

```

Configuring Device P1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P1:

1. Configure the interfaces.

```

[edit interfaces]
user@P1# set fe-1/2/0 unit 1 description P1-to-PE1
user@P1# set fe-1/2/0 unit 1 family inet address 10.0.0.2/30
user@P1# set fe-1/2/0 unit 1 family mpls
user@P1# set fe-1/2/1 unit 4 description P1-to-P2
user@P1# set fe-1/2/1 unit 4 family inet address 10.0.0.29/30
user@P1# set fe-1/2/1 unit 4 family mpls
user@P1# set fe-1/2/2 unit 8 description P1-to-PE4
user@P1# set fe-1/2/2 unit 8 family inet address 10.0.0.17/30
user@P1# set fe-1/2/2 unit 8 family mpls
user@P1# set lo0 unit 3 family inet address 10.9.9.2/32
user@P1# set lo0 unit 3 family inet address 10.100.1.2/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@P1# set rsvp interface fe-1/2/0.1
user@P1# set rsvp interface fe-1/2/2.8
user@P1# set rsvp interface fe-1/2/1.4
user@P1# set mpls label-switched-path P1-to-P2 to 10.9.9.3
user@P1# set mpls label-switched-path P1-to-PE1 to 10.9.9.1
user@P1# set mpls label-switched-path P1-to-PE4 to 10.9.9.4
user@P1# set mpls interface fe-1/2/0.1
user@P1# set mpls interface fe-1/2/2.8
user@P1# set mpls interface fe-1/2/1.4

```

3. Configure BGP.

```

[edit protocols bgp group internal]
user@P1# set type internal
user@P1# set local-address 10.9.9.2
user@P1# set neighbor 10.9.9.1
user@P1# set neighbor 10.9.9.3
user@P1# set neighbor 10.9.9.4

```

4. Enable AIGP.

```

[edit protocols bgp group internal]
user@P1# set family inet labeled-unicast aigp

```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P1# set area 0.0.0.1 interface fe-1/2/0.1 metric 1
user@P1# set area 0.0.0.1 interface fe-1/2/1.4 metric 1
user@P1# set area 0.0.0.0 interface fe-1/2/2.8 metric 1
user@P1# set area 0.0.0.0 interface 10.9.9.2 passive
user@P1# set area 0.0.0.0 interface 10.9.9.2 metric 1
user@P1# set area 0.0.0.0 interface 10.100.1.2 passive
user@P1# set area 0.0.0.0 interface 10.100.1.2 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P1# set router-id 10.9.9.2
user@P1# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show interfaces
fe-1/2/0 {
  unit 1 {
    description P1-to-PE1;
    family inet {
      address 10.0.0.2/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 4 {
    description P1-to-P2;
    family inet {
      address 10.0.0.29/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 8 {
    description P1-to-PE4;
    family inet {
      address 10.0.0.17/30;
    }
    family mpls;
  }
}
lo0 {
  unit 3 {
    family inet {
```



```

        address 10.9.9.2/32;
        address 10.100.1.2/32;
    }
}
}

user@P1# show protocols
rsvp {
    interface fe-1/2/0.1;
    interface fe-1/2/2.8;
    interface fe-1/2/1.4;
}
mpls {
    label-switched-path P1-to-P2 {
        to 10.9.9.3;
    }
    label-switched-path P1-to-PE1 {
        to 10.9.9.1;
    }
    label-switched-path P1-to-PE4 {
        to 10.9.9.4;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/2.8;
    interface fe-1/2/1.4;
}
bgp {
    group internal {
        type internal;
        local-address 10.9.9.2;
        family inet {
            labeled-unicast {
                aigp;
            }
        }
        neighbor 10.9.9.1;
        neighbor 10.9.9.3;
        neighbor 10.9.9.4;
    }
}
ospf {
    area 0.0.0.1 {
        interface fe-1/2/0.1 {
            metric 1;
        }
        interface fe-1/2/1.4 {
            metric 1;
        }
    }
    area 0.0.0.0 {
        interface fe-1/2/2.8 {
            metric 1;
        }
        interface 10.9.9.2 {
            passive;
            metric 1;
        }
    }
}

```

```
}  
interface 10.100.1.2 {  
    passive;  
    metric 1;  
}  
}  
}  
  
user@P1# show routing-options  
router-id 10.9.9.2;  
autonomous-system 13979;
```

Configuring Device P2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P2:

1. Configure the interfaces.

```
[edit interfaces]  
user@P2# set fe-1/2/0 unit 3 description P2-to-PE1  
user@P2# set fe-1/2/0 unit 3 family inet address 10.0.0.6/30  
user@P2# set fe-1/2/0 unit 3 family mpls  
user@P2# set fe-1/2/1 unit 5 description P2-to-P1  
user@P2# set fe-1/2/1 unit 5 family inet address 10.0.0.30/30  
user@P2# set fe-1/2/1 unit 5 family mpls  
user@P2# set fe-1/2/2 unit 6 description P2-to-PE4  
user@P2# set fe-1/2/2 unit 6 family inet address 10.0.0.13/30  
user@P2# set fe-1/2/2 unit 6 family mpls  
user@P2# set lo0 unit 5 family inet address 10.9.9.3/32  
user@P2# set lo0 unit 5 family inet address 10.100.1.3/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]  
user@P2# set rsvp interface fe-1/2/1.5  
user@P2# set rsvp interface fe-1/2/2.6  
user@P2# set rsvp interface fe-1/2/0.3  
user@P2# set mpls label-switched-path P2-to-PE1 to 10.9.9.1  
user@P2# set mpls label-switched-path P2-to-P1 to 10.9.9.2  
user@P2# set mpls label-switched-path P2-to-PE4 to 10.9.9.4  
user@P2# set mpls interface fe-1/2/1.5  
user@P2# set mpls interface fe-1/2/2.6  
user@P2# set mpls interface fe-1/2/0.3
```

3. Configure BGP.

```
[edit protocols bgp group internal]  
user@P2# set type internal  
user@P2# set local-address 10.9.9.3  
user@P2# set neighbor 10.9.9.1  
user@P2# set neighbor 10.9.9.2  
user@P2# set neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp group internal]
user@P2# set family inet labeled-unicast aigp
```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P2# set area 0.0.0.0 interface fe-1/2/2.6 metric 1
user@P2# set area 0.0.0.0 interface 10.9.9.3 passive
user@P2# set area 0.0.0.0 interface 10.9.9.3 metric 1
user@P2# set area 0.0.0.0 interface 10.100.1.3 passive
user@P2# set area 0.0.0.0 interface 10.100.1.3 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P2# set router-id 10.9.9.3
user@P2# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P2# show interfaces
fe-1/2/0 {
  unit 3 {
    description P2-to-PE1;
    family inet {
      address 10.0.0.6/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 5 {
    description P2-to-P1;
    family inet {
      address 10.0.0.30/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 6 {
    description P2-to-PE4;
    family inet {
      address 10.0.0.13/30;
    }
    family mpls;
  }
}
lo0 {
```

```
    unit 5 {
      family inet {
        address 10.9.9.3/32;
        address 10.100.1.3/32;
      }
    }
  }

user@P2# show protocols
rsvp {
  interface fe-1/2/1.5;
  interface fe-1/2/2.6;
  interface fe-1/2/0.3;
}
mpls {
  label-switched-path P2-to-PE1 {
    to 10.9.9.1;
  }
  label-switched-path P2-to-P1 {
    to 10.9.9.2;
  }
  label-switched-path P2-to-PE4 {
    to 10.9.9.4;
  }
  interface fe-1/2/1.5;
  interface fe-1/2/2.6;
  interface fe-1/2/0.3;
}
bgp {
  group internal {
    type internal;
    local-address 10.9.9.3;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
    neighbor 10.9.9.1;
    neighbor 10.9.9.2;
    neighbor 10.9.9.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/2.6 {
      metric 1;
    }
    interface 10.9.9.3 {
      passive;
      metric 1;
    }
    interface 10.100.1.3 {
      passive;
      metric 1;
    }
  }
}
```

```

}
user@P2# show routing-options
router-id 10.9.9.3;
autonomous-system 13979;

```

Configuring Device PE4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE4:

1. Configure the interfaces.

```

[edit interfaces]
user@PE4# set fe-1/2/0 unit 7 description PE4-to-P2
user@PE4# set fe-1/2/0 unit 7 family inet address 10.0.0.14/30
user@PE4# set fe-1/2/0 unit 7 family mpls
user@PE4# set fe-1/2/1 unit 9 description PE4-to-P1
user@PE4# set fe-1/2/1 unit 9 family inet address 10.0.0.18/30
user@PE4# set fe-1/2/1 unit 9 family mpls
user@PE4# set fe-1/2/2 unit 10 description PE4-to-PE2
user@PE4# set fe-1/2/2 unit 10 family inet address 10.0.0.21/30
user@PE4# set fe-1/2/2 unit 10 family mpls
user@PE4# set fe-1/0/2 unit 12 description PE4-to-PE3
user@PE4# set fe-1/0/2 unit 12 family inet address 10.0.0.25/30
user@PE4# set fe-1/0/2 unit 12 family mpls
user@PE4# set lo0 unit 7 family inet address 10.9.9.4/32
user@PE4# set lo0 unit 7 family inet address 10.100.1.4/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@PE4# set rsvp interface fe-1/2/0.7
user@PE4# set rsvp interface fe-1/2/1.9
user@PE4# set rsvp interface fe-1/2/2.10
user@PE4# set rsvp interface fe-1/0/2.12
user@PE4# set mpls label-switched-path PE4-to-PE2 to 10.9.9.5
user@PE4# set mpls label-switched-path PE4-to-PE3 to 10.9.9.6
user@PE4# set mpls label-switched-path PE4-to-P1 to 10.9.9.2
user@PE4# set mpls label-switched-path PE4-to-P2 to 10.9.9.3
user@PE4# set mpls interface fe-1/2/0.7
user@PE4# set mpls interface fe-1/2/1.9
user@PE4# set mpls interface fe-1/2/2.10
user@PE4# set mpls interface fe-1/0/2.12

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE4# set export next-hop
user@PE4# set export aigp
user@PE4# set group internal type internal
user@PE4# set group internal local-address 10.9.9.4
user@PE4# set group internal neighbor 10.9.9.1
user@PE4# set group internal neighbor 10.9.9.3

```

```
user@PE4# set group internal neighbor 10.9.9.2
user@PE4# set group external type external
user@PE4# set group external multihop ttl 2
user@PE4# set group external local-address 10.9.9.4
user@PE4# set group external peer-as 7018
user@PE4# set group external neighbor 10.9.9.5
user@PE4# set group external neighbor 10.9.9.6
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE4# set group external family inet labeled-unicast aigp
user@PE4# set group internal family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp term 10]
user@PE4# set from protocol static
user@PE4# set from route-filter 44.0.0.0/24 exact
user@PE4# set then aigp-originate distance 200
user@PE4# set then next-hop 10.100.1.4
user@PE4# set then accept
```

6. Enable the policies.

```
[edit policy-options policy-statement next-hop]
user@PE4# set term 10 from protocol bgp
user@PE4# set term 10 then next-hop 10.100.1.4
user@PE4# set term 10 then accept
user@PE4# set term 20 from protocol direct
user@PE4# set term 20 from route-filter 10.9.9.4/32 exact
user@PE4# set term 20 from route-filter 10.100.1.4/32 exact
user@PE4# set term 20 then next-hop 10.100.1.4
user@PE4# set term 20 then accept
```

7. Configure a static route.

```
[edit routing-options]
user@PE4# set static route 44.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@PE4# set area 0.0.0.0 interface fe-1/2/1.9 metric 1
user@PE4# set area 0.0.0.0 interface fe-1/2/0.7 metric 1
user@PE4# set area 0.0.0.0 interface 10.9.9.4 passive
user@PE4# set area 0.0.0.0 interface 10.9.9.4 metric 1
user@PE4# set area 0.0.0.0 interface 10.100.1.4 passive
user@PE4# set area 0.0.0.0 interface 10.100.1.4 metric 1
user@PE4# set area 0.0.0.2 interface fe-1/2/2.10 metric 1
user@PE4# set area 0.0.0.3 interface fe-1/0/2.12 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
```

```

user@PE4# set router-id 10.9.9.4
user@PE4# set autonomous-system 13979

```

10. If you are done configuring the device, commit the configuration.

```

user@PE4# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE4# show interfaces
fe-1/0/2 {
  unit 12 {
    description PE4-to-PE3;
    family inet {
      address 10.0.0.25/30;
    }
    family mpls;
  }
}
fe-1/2/0 {
  unit 7 {
    description PE4-to-P2;
    family inet {
      address 10.0.0.14/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 9 {
    description PE4-to-P1;
    family inet {
      address 10.0.0.18/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 10 {
    description PE4-to-PE2;
    family inet {
      address 10.0.0.21/30;
    }
    family mpls;
  }
}
lo0 {
  unit 7 {
    family inet {
      address 10.9.9.4/32;
      address 10.100.1.4/32;
    }
  }
}

```

```
    }  
  }  
  
user@PE4# show policy-options  
policy-statement aigp {  
  term 10 {  
    from {  
      protocol static;  
      route-filter 44.0.0.0/24 exact;  
    }  
    then {  
      aigp-originate distance 200;  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
}  
  
policy-statement next-hop {  
  term 10 {  
    from protocol bgp;  
    then {  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
  term 20 {  
    from {  
      protocol direct;  
      route-filter 10.9.9.4/32 exact;  
      route-filter 10.100.1.4/32 exact;  
    }  
    then {  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
}  
  
user@PE4# show protocols  
rsvp {  
  interface fe-1/2/0.7;  
  interface fe-1/2/1.9;  
  interface fe-1/2/2.10;  
  interface fe-1/0/2.12;  
}  
mpls {  
  label-switched-path PE4-to-PE2 {  
    to 10.9.9.5;  
  }  
  label-switched-path PE4-to-PE3 {  
    to 10.9.9.6;  
  }  
  label-switched-path PE4-to-P1 {  
    to 10.9.9.2;  
  }  
  label-switched-path PE4-to-P2 {  
    to 10.9.9.3;  
  }  
}
```



```

    }
    interface fe-1/2/0.7;
    interface fe-1/2/1.9;
    interface fe-1/2/2.10;
    interface fe-1/0/2.12;
  }
  bgp {
    export [ next-hop aigp ];
    group internal {
      type internal;
      local-address 10.9.9.4;
      family inet {
        labeled-unicast {
          aigp;
        }
      }
      neighbor 10.9.9.1;
      neighbor 10.9.9.3;
      neighbor 10.9.9.2;
    }
    group external {
      type external;
      multihop {
        ttl 2;
      }
      local-address 10.9.9.4;
      family inet {
        labeled-unicast {
          aigp;
        }
      }
      peer-as 7018;
      neighbor 10.9.9.5;
      neighbor 10.9.9.6;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface fe-1/2/1.9 {
        metric 1;
      }
      interface fe-1/2/0.7 {
        metric 1;
      }
      interface 10.9.9.4 {
        passive;
        metric 1;
      }
      interface 10.100.1.4 {
        passive;
        metric 1;
      }
    }
    area 0.0.0.2 {
      interface fe-1/2/2.10 {
        metric 1;
      }
    }
  }

```

```
    }  
  }  
  area 0.0.0.3 {  
    interface fe-1/0/2.12 {  
      metric 1;  
    }  
  }  
}  
  
user@PE4# show routing-options  
static {  
  route 44.0.0.0/24 discard;  
}  
router-id 10.9.9.4;  
autonomous-system 13979;
```

Configuring Device PE1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.

```
[edit interfaces]  
user@PE1# set fe-1/2/0 unit 0 description PE1-to-P1  
user@PE1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30  
user@PE1# set fe-1/2/0 unit 0 family mpls  
user@PE1# set fe-1/2/1 unit 2 description PE1-to-P2  
user@PE1# set fe-1/2/1 unit 2 family inet address 10.0.0.5/30  
user@PE1# set fe-1/2/1 unit 2 family mpls  
user@PE1# set fe-1/2/2 unit 14 description PE1-to-PE7  
user@PE1# set fe-1/2/2 unit 14 family inet address 10.0.0.9/30  
user@PE1# set lo0 unit 1 family inet address 10.9.9.1/32  
user@PE1# set lo0 unit 1 family inet address 10.100.1.1/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]  
user@PE1# set rsvp interface fe-1/2/0.0  
user@PE1# set rsvp interface fe-1/2/1.2  
user@PE1# set rsvp interface fe-1/2/2.14  
user@PE1# set mpls label-switched-path PE1-to-P1 to 10.9.9.2  
user@PE1# set mpls label-switched-path PE1-to-P2 to 10.9.9.3  
user@PE1# set mpls interface fe-1/2/0.0  
user@PE1# set mpls interface fe-1/2/1.2  
user@PE1# set mpls interface fe-1/2/2.14
```

3. Configure BGP.

```
[edit protocols bgp]  
user@PE1# set group internal type internal  
user@PE1# set group internal local-address 10.9.9.1  
user@PE1# set group internal export SET_EXPORT_ROUTES  
user@PE1# set group internal vpn-apply-export
```

```

user@PE1# set group internal neighbor 10.9.9.4
user@PE1# set group internal neighbor 10.9.9.2
user@PE1# set group internal neighbor 10.9.9.3
user@PE1# set group external type external
user@PE1# set group external export SET_EXPORT_ROUTES
user@PE1# set group external peer-as 7019
user@PE1# set group external neighbor 10.0.0.10

```

4. Enable AIGP.

```

[edit protocols bgp]
user@PE1# set group internal family inet labeled-unicast aigp
user@PE1# set group external family inet labeled-unicast aigp

```

5. Enable the policies.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE1# set from protocol direct
user@PE1# set from protocol bgp
user@PE1# set then next-hop 10.100.1.1
user@PE1# set then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.1]
user@PE1# set interface fe-1/2/0.0 metric 1
user@PE1# set interface fe-1/2/1.2 metric 1
user@PE1# set interface 10.9.9.1 passive
user@PE1# set interface 10.9.9.1 metric 1
user@PE1# set interface 10.100.1.1 passive
user@PE1# set interface 10.100.1.1 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE1# set router-id 10.9.9.1
user@PE1# set autonomous-system 13979

```

8. If you are done configuring the device, commit the configuration.

```

user@PE1# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE1# show interfaces
fe-1/2/0 {
  unit 0 {
    description PE1-to-P1;
    family inet {
      address 10.0.0.1/30;
    }
    family mpls;
  }
}
fe-1/2/1 {

```

```
    unit 2 {
      description PE1-to-P2;
      family inet {
        address 10.0.0.5/30;
      }
      family mpls;
    }
  }
  fe-1/2/2 {
    unit 14 {
      description PE1-to-PE7;
      family inet {
        address 10.0.0.9/30;
      }
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.9.9.1/32;
      address 10.100.1.1/32;
    }
  }
}

user@PE1# show policy-options
policy-statement SET_EXPORT_ROUTES {
  term 10 {
    from protocol [ direct bgp ];
    then {
      next-hop 10.100.1.1;
      accept;
    }
  }
}

user@PE1# show protocols
rsvp {
  interface fe-1/2/0.0;
  interface fe-1/2/1.2;
  interface fe-1/2/2.14;
}
mpls {
  label-switched-path PE1-to-P1 {
    to 10.9.9.2;
  }
  label-switched-path PE1-to-P2 {
    to 10.9.9.3;
  }
  interface fe-1/2/0.0;
  interface fe-1/2/1.2;
  interface fe-1/2/2.14;
}
bgp {
  group internal {
    type internal;
    local-address 10.9.9.1;
```

```
family inet {
  labeled-unicast {
    aigp;
  }
}
export SET_EXPORT_ROUTES;
vpn-apply-export;
neighbor 10.9.9.4;
neighbor 10.9.9.2;
neighbor 10.9.9.3;
}
group external {
  type external;
  family inet {
    labeled-unicast {
      aigp;
    }
  }
  export SET_EXPORT_ROUTES;
  peer-as 7019;
  neighbor 10.0.0.10;
}
}
ospf {
  area 0.0.0.1 {
    interface fe-1/2/0.0 {
      metric 1;
    }
    interface fe-1/2/1.2 {
      metric 1;
    }
    interface 10.9.9.1 {
      passive;
      metric 1;
    }
    interface 10.100.1.1 {
      passive;
      metric 1;
    }
  }
}
}
```

```
user@PE1# show routing-options
router-id 10.9.9.1;
autonomous-system 13979;
```

Configuring Device PE2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE2:

1. Configure the interfaces.

```
[edit interfaces]
user@PE2# set fe-1/2/0 unit 11 description PE2-to-PE4
user@PE2# set fe-1/2/0 unit 11 family inet address 10.0.0.22/30
user@PE2# set fe-1/2/0 unit 11 family mpls
user@PE2# set lo0 unit 9 family inet address 10.9.9.5/32 primary
user@PE2# set lo0 unit 9 family inet address 10.100.1.5/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE2# set rsvp interface fe-1/2/0.11
user@PE2# set mpls label-switched-path PE2-to-PE4 to 10.9.9.4
user@PE2# set mpls interface fe-1/2/0.11
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE2# set group external type external
user@PE2# set group external multihop ttl 2
user@PE2# set group external local-address 10.9.9.5
user@PE2# set group external export next-hop
user@PE2# set group external export aigp
user@PE2# set group external export SET_EXPORT_ROUTES
user@PE2# set group external vpn-apply-export
user@PE2# set group external peer-as 13979
user@PE2# set group external neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE2# set group external family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.0/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 99.0.0.0/24 exact
user@PE2# set term 20 then aigp-originate distance 30
user@PE2# set term 20 then next-hop 10.100.1.5
user@PE2# set term 20 then accept
```

6. Enable the policies.

```
[edit policy-options]
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
direct
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
static
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.5
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE2# set policy-statement next-hop term 10 from protocol bgp
user@PE2# set policy-statement next-hop term 10 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 10 then accept
user@PE2# set policy-statement next-hop term 20 from protocol direct
user@PE2# set policy-statement next-hop term 20 from route-filter 10.9.9.5/32
exact
user@PE2# set policy-statement next-hop term 20 from route-filter 10.100.1.5/32
exact
user@PE2# set policy-statement next-hop term 20 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 20 then accept
```

7. Enable some static routes.

```
[edit routing-options]
user@PE2# set static route 99.0.0.0/24 discard
user@PE2# set static route 55.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf area 0.0.0.2]
user@PE2# set interface 10.9.9.5 passive
user@PE2# set interface 10.9.9.5 metric 1
user@PE2# set interface 10.100.1.5 passive
user@PE2# set interface 10.100.1.5 metric 1
user@PE2# set interface fe-1/2/0.11 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE2# set router-id 10.9.9.5
user@PE2# set autonomous-system 7018
```

10. If you are done configuring the device, commit the configuration.

```
user@PE2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
fe-1/2/0 {
  unit 11 {
    description PE2-to-PE4;
```

```
        family inet {
            address 10.0.0.22/30;
        }
        family mpls;
    }
}
lo0 {
    unit 9 {
        family inet {
            address 10.9.9.5/32 {
                primary;
            }
            address 10.100.1.5/32;
        }
    }
}

user@PE2# show policy-options
policy-statement SET_EXPORT_ROUTES {
    term 10 {
        from protocol [ direct static bgp ];
        then {
            next-hop 10.100.1.5;
            accept;
        }
    }
}

policy-statement aigp {
    term 10 {
        from {
            route-filter 55.0.0.0/24 exact;
        }
        then {
            aigp-originate distance 20;
            next-hop 10.100.1.5;
            accept;
        }
    }
    term 20 {
        from {
            route-filter 99.0.0.0/24 exact;
        }
        then {
            aigp-originate distance 30;
            next-hop 10.100.1.5;
            accept;
        }
    }
}

policy-statement next-hop {
    term 10 {
        from protocol bgp;
        then {
            next-hop 10.100.1.5;
            accept;
        }
    }
}
```



```

    }
    term 20 {
      from {
        protocol direct;
        route-filter 10.9.9.5/32 exact;
        route-filter 10.100.1.5/32 exact;
      }
      then {
        next-hop 10.100.1.5;
        accept;
      }
    }
  }
}

user@PE2# show protocols
rsvp {
  interface fe-1/2/0.11;
}
mpls {
  label-switched-path PE2-to-PE4 {
    to 10.9.9.4;
  }
  interface fe-1/2/0.11;
}
bgp {
  group external {
    type external;
    multihop {
      ttl 2;
    }
    local-address 10.9.9.5;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
    export [ next-hop aigp SET_EXPORT_ROUTES ];
    vpn-apply-export;
    peer-as 13979;
    neighbor 10.9.9.4;
  }
}
ospf {
  area 0.0.0.2 {
    interface 10.9.9.5 {
      passive;
      metric 1;
    }
    interface 10.100.1.5 {
      passive;
      metric 1;
    }
    interface fe-1/2/0.11 {
      metric 1;
    }
  }
}

```

```
}
user@PE2# show routing-options
static {
    route 99.0.0.0/24 discard;
    route 55.0.0.0/24 discard;
}
router-id 10.9.9.5;
autonomous-system 7018;
```

Configuring Device PE3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE3:

1. Configure the interfaces.

```
[edit interfaces]
user@PE3# set fe-1/2/0 unit 13 description PE3-to-PE4
user@PE3# set fe-1/2/0 unit 13 family inet address 10.0.0.26/30
user@PE3# set fe-1/2/0 unit 13 family mpls
user@PE3# set lo0 unit 11 family inet address 10.9.9.6/32
user@PE3# set lo0 unit 11 family inet address 10.100.1.6/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE3# set rsvp interface fe-1/2/0.13
user@PE3# set mpls label-switched-path PE3-to-PE4 to 10.9.9.4
user@PE3# set mpls interface fe-1/2/0.13
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@PE3# set type external
user@PE3# set multihop ttl 2
user@PE3# set local-address 10.9.9.6
user@PE3# set export next-hop
user@PE3# set export SET_EXPORT_ROUTES
user@PE3# set vpn-apply-export
user@PE3# set peer-as 13979
user@PE3# set neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp group external]
user@PE3# set family inet labeled-unicast aigp
```

5. Enable the policies.

```
[edit policy-options]
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
    direct
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
    static
```

```

user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.6
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE3# set policy-statement next-hop term 10 from protocol bgp
user@PE3# set policy-statement next-hop term 10 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 10 then accept
user@PE3# set policy-statement next-hop term 20 from protocol direct
user@PE3# set policy-statement next-hop term 20 from route-filter 10.9.9.6/32
exact
user@PE3# set policy-statement next-hop term 20 from route-filter 10.100.1.6/32
exact
user@PE3# set policy-statement next-hop term 20 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 20 then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.3]
user@PE3# set interface 10.9.9.6 passive
user@PE3# set interface 10.9.9.6 metric 1
user@PE3# set interface 10.100.1.6 passive
user@PE3# set interface 10.100.1.6 metric 1
user@PE3# set interface fe-1/2/0.13 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE3# set router-id 10.9.9.6
user@PE3# set autonomous-system 7018

```

8. If you are done configuring the device, commit the configuration.

```

user@PE3# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE3# show interfaces
fe-1/2/0 {
  unit 13 {
    description PE3-to-PE4;
    family inet {
      address 10.0.0.26/30;
    }
    family mpls;
  }
}
lo0 {
  unit 11 {
    family inet {
      address 10.9.9.6/32;
      address 10.100.1.6/32;
    }
  }
}

```

```
    }
  }
user@PE3# show policy-options
policy-statement SET_EXPORT_ROUTES {
  term 10 {
    from protocol [ direct static bgp ];
    then {
      next-hop 10.100.1.6;
      accept;
    }
  }
}
policy-statement next-hop {
  term 10 {
    from protocol bgp;
    then {
      next-hop 10.100.1.6;
      accept;
    }
  }
  term 20 {
    from {
      protocol direct;
      route-filter 10.9.9.6/32 exact;
      route-filter 10.100.1.6/32 exact;
    }
    then {
      next-hop 10.100.1.6;
      accept;
    }
  }
}
user@PE3# show protocols
rsvp {
  interface fe-1/2/0.13;
}
mpls {
  label-switched-path PE3-to-PE4 {
    to 10.9.9.4;
  }
  interface fe-1/2/0.13;
}
bgp {
  group external {
    type external;
    multihop {
      ttl 2;
    }
    local-address 10.9.9.6;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
  }
  export [ next-hop SET_EXPORT_ROUTES ];
```

```

        vpn-apply-export;
        peer-as 13979;
        neighbor 10.9.9.4;
    }
}
ospf {
    area 0.0.0.3 {
        interface 10.9.9.6 {
            passive;
            metric 1;
        }
        interface 10.100.1.6 {
            passive;
            metric 1;
        }
        interface fe-1/2/0.13 {
            metric 1;
        }
    }
}

user@PE3# show routing-options
router-id 10.9.9.6;
autonomous-system 7018;

```

Configuring Device PE7

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE7:

1. Configure the interfaces.

```

[edit interfaces]
user@PE7# set fe-1/2/0 unit 15 description PE7-to-PE1
user@PE7# set fe-1/2/0 unit 15 family inet address 10.0.0.10/30
user@PE7# set lo0 unit 13 family inet address 10.9.9.7/32
user@PE7# set lo0 unit 13 family inet address 10.100.1.7/32

```

2. Configure BGP.

```

[edit protocols bgp group external]
user@PE7# set type external
user@PE7# set export SET_EXPORT_ROUTES
user@PE7# set peer-as 13979
user@PE7# set neighbor 10.0.0.9

```

3. Enable AIGP.

```

[edit protocols bgp group external]
user@PE7# set family inet labeled-unicast aigp

```

4. Configure the routing policy.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE7# set from protocol direct

```

```
user@PE7# set from protocol bgp
user@PE7# set then next-hop 10.100.1.7
user@PE7# set then accept
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE7# set router-id 10.9.9.7
user@PE7# set autonomous-system 7019
```

6. If you are done configuring the device, commit the configuration.

```
user@PE7# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE7# show interfaces
interfaces {
  fe-1/2/0 {
    unit 15 {
      description PE7-to-PE1;
      family inet {
        address 10.0.0.10/30;
      }
    }
  }
  lo0 {
    unit 13 {
      family inet {
        address 10.9.9.7/32;
        address 10.100.1.7/32;
      }
    }
  }
}

user@PE7# show policy-options
policy-statement SET_EXPORT_ROUTES {
  term 10 {
    from protocol [ direct bgp ];
    then {
      next-hop 10.100.1.7;
      accept;
    }
  }
}

user@PE7# show protocols
bgp {
  group external {
    type external;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
  }
}
```

```

    }
  }
  export SET_EXPORT_ROUTES;
  peer-as 13979;
  neighbor 10.0.0.9;
}
}

user@PE7# show routing-options
router-id 10.9.9.7;
autonomous-system 7019;

```

Verification

Confirm that the configuration is working properly.

- [Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2 on page 3299](#)
- [Checking the IGP Metric on page 3299](#)
- [Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute on page 3300](#)
- [Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE1 on page 3300](#)
- [Verifying the Resolving AIGP Metric on page 3301](#)
- [Verifying the Presence of AIGP Attributes in BGP Updates on page 3304](#)

Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2

Purpose Make sure that the AIGP policy on Device PE2 is working.

Action

```

user@PE4> show route receive-protocol bgp 10.9.9.5 extensive
* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 299888
  Nexthop: 10.100.1.5
  AS path: 7018 I
  AIGP: 20

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 299888
  Nexthop: 10.100.1.5
  AS path: 7018 I
  AIGP: 30

```

Meaning On Device PE2, the **aigp-originate** statement is configured with a distance of 20 (**aigp-originate distance 20**). This statement is applied to route 55.0.0.0/24. Likewise, the **aigp-originate distance 30** statement is applied to route 99.0.0.0/24. Thus, when Device PE4 receives these routes, the AIGP attribute is attached with the configured metrics.

Checking the IGP Metric

Purpose From Device PE4, check the IGP metric to the BGP next hop 10.100.1.5.

Action user@PE4> show route 10.100.1.5
inet.0: 30 destinations, 40 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.1.5/32 * [OSPF/10] 05:35:50, metric 2
 > to 10.0.0.22 via fe-1/2/2.10
 [BGP/170] 03:45:07, localpref 100, from 10.9.9.5
 AS path: 7018 I
 > to 10.0.0.22 via fe-1/2/2.10

Meaning The IGP metric for this route is 2.

Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute

Purpose Make sure that Device PE4 adds the IGP metric to the AIGP attribute when it readvertises routes to its IBGP neighbor, Device PE1.

Action user@PE4> show route advertising-protocol bgp 10.9.9.1 extensive

* 55.0.0.0/24 (1 entry, 1 announced)
BGP group internal type Internal
Route Label: 300544
Nexthop: 10.100.1.4
Flags: Nexthop Change
Localpref: 100
AS path: [13979] 7018 I
AIGP: 22

* 99.0.0.0/24 (1 entry, 1 announced)
BGP group internal type Internal
Route Label: 300544
Nexthop: 10.100.1.4
Flags: Nexthop Change
Localpref: 100
AS path: [13979] 7018 I
AIGP: 32

Meaning The IGP metric is added to the AIGP metric ($20 + 2 = 22$ and $30 + 2 = 32$), because the next hop is changed for these routes.

Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGP Neighbor PE1

Purpose Make sure that the AIGP policy on Device PE1 is working.

Action user@PE7> show route receive-protocol bgp 10.0.0.9 extensive

```
* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 35
```

Meaning The 44.0.0.0/24 route is originated at Device PE4. The 55.0.0.0/24 and 99.0.0.0/24 routes are originated at Device PE2. The IGP distances are added to the configured AIGP distances.

Verifying the Resolving AIGP Metric

Purpose Confirm that if the prefix is resolved through recursion and the recursive next hops have AIGP metrics, the prefix has the sum of the AIGP values that are on the recursive BGP next hops.

Action 1. Add a static route to 66.0.0.0/24.

```
[edit routing-options]
user@PE2# set static route 66.0.0.0/24 discard
```

2. Delete the existing terms in the **aigp** policy statement on Device PE2.

```
[edit policy-options policy-statement aigp]
user@PE2# delete term 10
user@PE2# delete term 20
```

3. Configure a recursive route lookup for the route to 66.0.0.0.

The policy shows the AIGP metric for prefix 66.0.0.0/24 (none) and its recursive next hop. Prefix 66.0.0.0/24 is resolved by 55.0.0.1. Prefix 66.0.0.0/24 does not have its own AIGP metric being originated, but its recursive next hop, 55.0.0.1, has an AIGP value.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.1/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 66.0.0.0/24 exact
user@PE2# set term 20 then next-hop 55.0.0.1
```

user@PE2# set term 20 then accept

4. On Device PE4, run the **show route 55.0.0.0 extensive** command.

The value of Metric2 is the IGP metric to the BGP next hop. When Device PE4 readvertises these routes to its IBGP peer, Device PE1, the AIGP metric is the sum of AIGP + its Resolving AIGP metric + Metric2.

Prefix 55.0.0.0 shows its own IGP metric 20, as defined and advertised by Device PE2. It does not show a resolving AIGP value because it does not have a recursive BGP next hop. The value of Metric2 is 2.

```
user@PE4> show route 55.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 55.0.0.0/24 -> {indirect(262151)}
Page 0 idx 0 Type 1 val 928d1b8
  Flags: Nexthop Change
  Nexthop: 10.100.1.4
  Localpref: 100
  AS path: [13979] 7018 I
  Communities:
  AIGP: 22
Path 55.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x925da38
    Next-hop reference count: 4
    Source: 10.9.9.5
    Next hop type: Router, Next hop index: 1004
    Next hop: 10.0.0.22 via fe-1/2/2.10, selected
    Label operation: Push 299888
    Label TTL action: prop-ttl
    Protocol next hop: 10.100.1.5
    Push 299888
    Indirect next hop: 93514d8 262151
    State: <Active Ext>
    Local AS: 13979 Peer AS: 7018
    Age: 22:03:26 Metric2: 2
    AIGP: 20
    Task: BGP_7018.10.9.9.5+58560
    Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
    AS path: 7018 I
    Accepted
    Route Label: 299888
    Localpref: 100
    Router ID: 10.9.9.5
    Indirect next hops: 1
      Protocol next hop: 10.100.1.5 Metric: 2
      Push 299888
      Indirect next hop: 93514d8 262151
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.0.22 via fe-1/2/2.10
      10.100.1.5/32 Originating RIB: inet.0
        Metric: 2 Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.0.0.22 via fe-1/2/2.10
```

5. On Device PE4, run the **show route 66.0.0.0 extensive** command.

Prefix 66.0.0.0/24 shows the Resolving AIGP, which is the sum of its own AIGP metric and its recursive BGP next hop:

66.0.0.1 = 0, 55.0.0.1 = 20, 0+20 = 20

```

user@PE4> show route 66.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
66.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kerne1 66.0.0.0/24 -> {indirect(262162)}
Page 0 idx 0 Type 1 val 928cefc
  Flags: Nexthop Change
  Nexthop: 10.100.1.4
  Localpref: 100
  AS path: [13979] 7018 I
  Communities:
Path 66.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x925d4e0
    Next-hop reference count: 4
    Source: 10.9.9.5
    Next hop type: Router, Next hop index: 1006
    Next hop: 10.0.0.22 via fe-1/2/2.10, selected
    Label operation: Push 299888, Push 299888(top)
    Label TTL action: prop-ttl, prop-ttl(top)
    Protocol next hop: 55.0.0.1
    Push 299888
    Indirect next hop: 9353e88 262162
    State: <Active Ext>
    Local AS: 13979 Peer AS: 7018
    Age: 31:42 Metric2:2
    Resolving-AIGP: 20
    Task: BGP_7018.10.9.9.5+58560
    Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
    AS path: 7018 I
    Accepted
    Route Label: 299888
    Localpref: 100
    Router ID: 10.9.9.5
    Indirect next hops: 1
      Protocol next hop: 55.0.0.1 Metric: 2 AIGP: 20
      Push 299888
      Indirect next hop: 9353e88 262162
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.0.22 via fe-1/2/2.10
      55.0.0.0/24 Originating RIB: inet.0
        Metric: 2 Node path count: 1
        Indirect nexthops: 1
          Protocol Nexthop: 10.100.1.5 Metric: 2 Push 299888
          Indirect nexthop: 93514d8 262151
          Indirect path forwarding nexthops: 1
            Nexthop: 10.0.0.22 via fe-1/2/2.10
          10.100.1.5/32 Originating RIB: inet.0
            Metric: 2 Node path count: 1
            Forwarding nexthops: 1
              Nexthop: 10.0.0.22 via fe-1/2/2.10

```

Verifying the Presence of AIGP Attributes in BGP Updates

Purpose If the AIGP attribute is not enabled under BGP (or the **group** or **neighbor** hierarchies), the AIGP attribute is silently discarded. Enable **traceoptions** and include the **packets** flag in the **detail** option in the configuration to confirm the presence of the AIGP attribute in transmitted or received BGP updates. This is useful when debugging AIGP issues.

Action 1. Configure Device PE2 and Device PE4 for **traceoptions**.

```
user@host> show protocols bgp
  traceoptions {
    file bgp size 1m files 5;
    flag packets detail;
  }
```

2. Check the **traceoptions** file on Device PE2.

The following sample shows Device PE2 advertising prefix 99.0.0.0/24 to Device PE4 (10.9.9.4) with an AIGP metric of 20:

```
user@PE2> show log bgp
Mar 22 09:27:18.982150 BGP SEND 10.9.9.5+49652 -> 10.9.9.4+179
Mar 22 09:27:18.982178 BGP SEND message type 2 (Update) length 70
Mar 22 09:27:18.982198 BGP SEND Update PDU length 70
Mar 22 09:27:18.982248 BGP SEND flags 0x40 code Origin(1): IGP
Mar 22 09:27:18.982273 BGP SEND flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:27:18.982295 BGP SEND flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:27:18.982316 BGP SEND flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:27:18.982341 BGP SEND      nhop 10.100.1.5 len 4
Mar 22 09:27:18.982372 BGP SEND    99.0.0.0/24 (label 301664)
Mar 22 09:27:33.665412 bgp_send: sending 19 bytes to abcd::10:255:170:84
(External AS 13979)
```

3. Verify that the route was received on Device PE4 using the **show route receive-protocol** command.

AIGP is not enabled on Device PE4, so the AIGP attribute is silently discarded for prefix 99.0.0.0/24 and does not appear in the following output:

```
user@PE4> show route receive-protocol bgp 10.9.9.5 extensive | find 55.0.0.0
* 99.0.0.0/24 (2 entries, 1 announced)
  Accepted
  Route Label: 301728
  Nexthop: 10.100.1.5
  AS path: 7018 I
```

4. Check the **traceoptions** file on Device PE4.

The following output from the **traceoptions** log shows that the 99.0.0.0/24 prefix was received with the AIGP attribute attached:

```
user@PE4> show log bgp
Mar 22 09:41:39.650295 BGP RECV 10.9.9.5+64690 -> 10.9.9.4+179
Mar 22 09:41:39.650331 BGP RECV message type 2 (Update) length 70
Mar 22 09:41:39.650350 BGP RECV Update PDU length 70
Mar 22 09:41:39.650370 BGP RECV flags 0x40 code Origin(1): IGP
Mar 22 09:41:39.650394 BGP RECV flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:41:39.650415 BGP RECV flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:41:39.650436 BGP RECV flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:41:39.650459 BGP RECV      nhop 10.100.1.5 len 4
```

```

Mar 22 09:41:39.650495 BGP RECV    99.0.0.0/24 (label 301728)
Mar 22 09:41:39.650574 bgp_rcv_nlri: 99.0.0.0/24
Mar 22 09:41:39.650607 bgp_rcv_nlri: 99.0.0.0/24 belongs to meshgroup
Mar 22 09:41:39.650629 bgp_rcv_nlri: 99.0.0.0/24 qualified bnp->ribact 0x0
12afcb 0x0

```

Meaning Performing this verification helps with AIGP troubleshooting and debugging issues. It enables you to verify which devices in your network send and receive AIGP attributes.

- Related Documentation**
- [Understanding BGP Path Selection on page 3332](#)
 - [Examples: Configuring Internal BGP Peering on page 3172](#)

BGP Policy Configuration

- [Example: Configuring BGP Interactions with IGP on page 3305](#)
- [Example: Configuring BGP Route Advertisement on page 3309](#)
- [Example: Configuring EBGp Multihop on page 3317](#)
- [Example: Configuring BGP Route Preference \(Administrative Distance\) on page 3326](#)
- [Example: Configuring BGP Path Selection on page 3332](#)
- [Example: Removing Private AS Numbers on page 3342](#)

Example: Configuring BGP Interactions with IGPs

- [Understanding Routing Policies on page 3305](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 3306](#)

Understanding Routing Policies

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration.

Once a policy is created and named, it must be applied before it is active. You apply routing policies using the **import** and **export** statements at the **protocols>protocol-name** level in the configuration hierarchy.

In the **import** statement, you list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

Example: Injecting OSPF Routes into the BGP Routing Table

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 3306](#)
- [Overview on page 3306](#)
- [Configuration on page 3306](#)
- [Verification on page 3308](#)
- [Troubleshooting on page 3308](#)

Requirements

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3150](#).
- Configure interior gateway protocol (IGP) sessions between peers.

Overview

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

Configuration

- [Configuring the Routing Policy on page 3306](#)
- [Configuring Tracing for the Routing Policy on page 3307](#)

Configuring the Routing Policy

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```

- Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

- Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```

- Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

- Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

Results Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    from {
      protocol ospf;
      area 0.0.0.1;
    }
    then accept;
  }
}

user@host# show protocols bgp
export injectpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Tracing for the Routing Policy

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Include a trace action in the policy.

```
[edit policy-options policy-statement injectpolicy] term injectterm1]
user@host# then trace
```

2. Configure the tracing file for the output.

```
[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy
```

Results Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    then {
      trace;
    }
  }
}

user@host# show routing-options
traceoptions {
  file ospf-bgp-policy-log size 5m files 5;
  flag policy;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying That the Expected BGP Routes Are Present

Purpose Verify the effect of the export policy.

Action From operational mode, enter the **show route** command.

Troubleshooting

- [Using the show log Command to Examine the Actions of the Routing Policy on page 3308](#)

Using the show log Command to Examine the Actions of the Routing Policy

Problem The routing table contains unexpected routes, or routes are missing from the routing table.

Solution If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring BGP Route Advertisement

- [Understanding Route Advertisement on page 3309](#)
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3313](#)

Understanding Route Advertisement

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes they should advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table. For information about routing policy, see the *Routing Policy Configuration Guide*.

When configuring BGP routing policy, you can perform the following tasks:

- [Applying Routing Policy on page 3309](#)
- [Setting BGP to Advertise Inactive Routes on page 3310](#)
- [Configuring BGP to Advertise the Best External Route to Internal Peers on page 3310](#)
- [Configuring How Often BGP Exchanges Routes with the Routing Table on page 3312](#)
- [Disabling Suppression of Route Advertisements on page 3313](#)

Applying Routing Policy

You define routing policy at the **[edit policy-options]** hierarchy level. To apply policies you have defined for BGP, include the **import** and **export** statements within the BGP configuration. For information about defining policy, see the *Routing Policy Configuration Guide*.

You can apply policies as follows:

- BGP global **import** and **export** statements—Include these statements at the **[edit protocols bgp]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level).
- Group **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp group group-name]** hierarchy level).
- Peer **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name neighbor address]** hierarchy level (for routing instances, include

these statements at the `[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]` hierarchy level).

A peer-level **import** or **export** statement overrides a group **import** or **export** statement. A group-level **import** or **export** statement overrides a global BGP **import** or **export** statement.

To apply policies, see the following sections:

- [Applying Policies to Routes Being Imported into the Routing Table from BGP on page 3310](#)
- [Applying Policies to Routes Being Exported from the Routing Table into BGP on page 3310](#)

Applying Policies to Routes Being Imported into the Routing Table from BGP

To apply policy to routes being imported into the routing table from BGP, include the **import** statement, listing the names of one or more policies to be evaluated:

```
import [ policy-names ];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices.

Applying Policies to Routes Being Exported from the Routing Table into BGP

To apply policy to routes being exported from the routing table into BGP, include the **export** statement, listing the names of one or more policies to be evaluated:

```
export [ policy-names ];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP.

Setting BGP to Advertise Inactive Routes

By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. To have the routing table export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route, include the **advertise-inactive** statement:

```
advertise-inactive;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring BGP to Advertise the Best External Route to Internal Peers

In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this

behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.

In Junos OS Release 9.3 and later, you can configure BGP to advertise the best external route into an internal BGP (IBGP) mesh group, a route reflector cluster, or an autonomous system (AS) confederation, even when the best route is an internal route.



NOTE: In order to configure the `advertise-external` statement on a route reflector, you must disable intracluster reflection with the `no-client-reflect` statement.

When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.

In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external.

You can also configure BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external route with an AS path worse (that is, longer) than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes. For more information, see the *Routing Policy Configuration Guide*.

To configure BGP to advertise the best external path to internal peers, include the `advertise-external` statement:

```
advertise-external;
```



NOTE: The `advertise-external` statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.

For a complete list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To configure BGP to advertise the best external path only if the route selection process reaches the point where the MED value is evaluated, include the `conditional` statement:

```
advertise-external {
  conditional;
```

```
}
```

For a complete list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring How Often BGP Exchanges Routes with the Routing Table

BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. To guard against this, you can delay the time between when BGP and the routing table exchange route information.

To configure how often BGP and the routing table exchange route information, include the **out-delay** statement:

```
out-delay seconds;
```

By default, the routing table retains some of the route information learned from BGP. To have the routing table retain all or none of this information, include the **keep** statement:

```
keep (all | none);
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

The routing table can retain the route information learned from BGP in one of the following ways:

- Default (omit the **keep** statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.
- **keep all**—Keep all route information that was learned from BGP.
- **keep none**—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure **keep none** for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.

- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.

Disabling Suppression of Route Advertisements

Junos OS does not advertise the routes learned from one EBGp peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGp peers that are in the same AS as the originating peer, regardless of the routing instance. You can modify this behavior by including the **advertise-peer-as** statement in the configuration. To disable the default advertisement suppression, include the **advertise-peer-as** statement:

advertise-peer-as;



NOTE: The route suppression default behavior is disabled if the **as-override** statement is included in the configuration.

If you include the **advertise-peer-as** statement in the configuration, BGP advertises the route regardless of this check.

To restore the default behavior, include the **no-advertise-peer-as** statement in the configuration:

no-advertise-peer-as;

If you include both the **as-override** and **no-advertise-peer-as** statements in the configuration, the **no-advertise-peer-as** statement is ignored. You can include these statements at multiple hierarchy levels.

For a list of hierarchy levels at which you can include these statements, see the statement summary section for these statements.

Example: Configuring BGP Prefix-Based Outbound Route Filtering

This example shows how to configure a Juniper Networks router to accept route filters from remote peers and perform outbound route filtering using the received filters.

- [Requirements on page 3313](#)
- [Overview on page 3314](#)
- [Configuration on page 3314](#)
- [Verification on page 3316](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

Overview

You can configure a BGP peer to accept route filters from remote peers and perform outbound route filtering using the received filters. By filtering out unwanted updates, the sending peer saves resources needed to generate and transmit updates, and the receiving peer saves resources needed to process updates. This feature can be useful, for example, in a virtual private network (VPN) in which subsets of customer edge (CE) devices are not capable of processing all the routes in the VPN. The CE devices can use prefix-based outbound route filtering to communicate to the provider edge (PE) routing device to transmit only a subset of routes, such as routes to the main data centers only.

The maximum number of prefix-based outbound route filters that a BGP peer can accept is 5000. If a remote peer sends more than 5000 outbound route filters to a peer address, the additional filters are discarded, and a system log message is generated.


You can configure interoperability for the routing device as a whole or for specific BGP groups or peers only.

Topology

In the sample network, Device CE1 is a router from another vendor. The configuration shown in this example is on Juniper Networks Router PE1.

Figure 88 on page 3314 shows the sample network.

Figure 88: BGP Prefix-Based Outbound Route Filtering



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
PE1 set protocols bgp group cisco-peers type external
    set protocols bgp group cisco-peers description "to CE1"
    set protocols bgp group cisco-peers local-address 192.168.165.58
    set protocols bgp group cisco-peers peer-as 35
    set protocols bgp group cisco-peers outbound-route-filter bgp-orf-cisco-mode
    set protocols bgp group cisco-peers outbound-route-filter prefix-based accept inet
    set protocols bgp group cisco-peers neighbor 192.168.165.56
    set routing-options autonomous-system 65500
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.

```
[edit routing-options]
user@PE1# set autonomous-system 65500
```

2. Configure external peering with Device CE1.

```
[edit protocols bgp group cisco-peers]
user@PE1# set type external
user@PE1# set description "to CE1"
user@PE1# set local-address 192.168.165.58
user@PE1# set peer-as 35
user@PE1# set neighbor 192.168.165.56
```

3. Configure Router PE1 to accept IPv4 route filters from Device CE1 and perform outbound route filtering using the received filters.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter prefix-based accept inet
```

4. (Optional) Enable interoperability with routing devices that use the vendor-specific compatibility code of 130 for outbound route filters and the code type of 128.

The IANA standard code is 3, and the standard code type is 64.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter bgp-orf-cisco-mode
```

Results From configuration mode, confirm your configuration by entering the **show protocols** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show protocols
group cisco-peers {
  type external;
  description "to CE1";
  local-address 192.168.165.58;
  peer-as 35;
  outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
      accept {
        inet;
      }
    }
  }
  neighbor 192.168.165.56;
}

user@PE1# show routing-options
autonomous-system 65500;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the Outbound Route Filter on page 3316](#)
- [Verifying the BGP Neighbor Mode on page 3316](#)

Verifying the Outbound Route Filter

Purpose Display information about the prefix-based outbound route filter received from Device CE1.

Action From operational mode, enter the [show bgp neighbor orf detail](#) command.

```
user@PE1> show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56 Type: External
Group: cisco-peers

inet-unicast
Filter updates rcv:          4 Immediate:          0
Filter: prefix-based         receive
Updates rcv:                 4
Received filter entries:
  seq 10 2.2.0.0/16 deny minlen 0 maxlen 0
  seq 20 3.3.0.0/16 deny minlen 24 maxlen 0
  seq 30 4.4.0.0/16 deny minlen 0 maxlen 28
  seq 40 5.5.0.0/16 deny minlen 24 maxlen 28
```

Verifying the BGP Neighbor Mode

Purpose Verify that the **bgp-orf-cisco-mode** setting is enabled for the peer by making sure that the **ORFCiscoMode** option is displayed in the **show bgp neighbor** command output.

Action From operational mode, enter the [show bgp neighbor](#) command.

```
user@PE1> show bgp neighbor
Peer: 192.168.165.56 AS 35          Local: 192.168.165.58 AS 65500
Type: External   State: Active     Flags: <>
Last State: Idle   Last Event: Start
Last Error: None
Export: [ adv_stat ]
Options: <Preference LocalAddress AddressFamily PeerAS Refresh>
Options: <ORF ORFCiscoMode>
Address families configured: inet-unicast
Local Address: 192.168.165.58 Holdtime: 90 Preference: 170
Number of flaps: 0
Trace options: detail open detail refresh
Trace file: /var/log/orf size 5242880 files 20
```

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring EBGP Multihop

- [Understanding BGP Multihop on page 3317](#)
- [Example: Configuring EBGP Multihop Sessions on page 3317](#)

Understanding BGP Multihop

When external BGP (EBGP) peers are not directly connected to each other, they must cross one or more non-BGP routers to reach each other. Configuring multihop EBGP enables the peers to pass through the other routers to form peer relationships and exchange update messages. This type of configuration is typically used when a Juniper Networks routing device needs to run EBGP with a third-party router that does not allow direct connection of the two EBGP peers. EBGP multihop enables a neighbor connection between two EBGP peers that do not have a direct connection.

Example: Configuring EBGP Multihop Sessions

This example shows how to configure an external BGP (EBGP) peer that is more than one hop away from the local router. This type of session is called a *multihop* BGP session.

- [Requirements on page 3317](#)
- [Overview on page 3317](#)
- [Configuration on page 3318](#)
- [Verification on page 3324](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

The configuration to enable multihop EBGP sessions requires connectivity between the two EBGP peers. This example uses static routes to provide connectivity between the devices.


Unlike directly connected EBGP sessions in which physical address are typically used in the **neighbor** statements, you must use loopback interface addresses for multihop EBGP by specifying the loopback interface address of the indirectly connected peer. In this way, EBGP multihop is similar to internal BGP (IBGP).

Finally, you must add the **multihop** statement. Optionally, you can set a maximum time-to-live (TTL) value with the **ttl** statement. The TTL is carried in the IP header of BGP packets. If you do not specify a TTL value, the system's default maximum TTL value is used. The default TTL value is 64 for multihop EBGP sessions. Another option is to retain the BGP next-hop value for route advertisements by including the **no-nexthop-change** statement.

[Figure 89 on page 3318](#) shows a typical EBGP multihop network.

Device C and Device E have an established EBGP session. Device D is not a BGP-enabled device. All of the devices have connectivity via static routes.

Figure 89: Typical Network with EBGp Multihop Sessions



8041621

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device C

```

set interfaces fe-1/2/0 unit 9 description to-D
set interfaces fe-1/2/0 unit 9 family inet address 10.10.10.9/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.40.4
set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 18
set protocols bgp group external-peers neighbor 192.168.6.7
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.10.10.14/32 next-hop 10.10.10.10
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.10
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17

```

Device D

```

set interfaces fe-1/2/0 unit 10 description to-C
set interfaces fe-1/2/0 unit 10 family inet address 10.10.10.10/30
set interfaces fe-1/2/1 unit 13 description to-E
set interfaces fe-1/2/1 unit 13 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.6.6/32
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.9
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.14
set routing-options router-id 192.168.6.6

```

Device E

```

set interfaces fe-1/2/0 unit 14 description to-D
set interfaces fe-1/2/0 unit 14 family inet address 10.10.10.14/30
set interfaces lo0 unit 5 family inet address 192.168.6.7/32
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.6.7
set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 17
set protocols bgp group external-peers neighbor 192.168.40.4
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept

```

```

set routing-options static route 10.10.10.8/30 next-hop 10.10.10.13
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.13
set routing-options router-id 192.168.6.7
set routing-options autonomous-system 18

```

Device C

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device C:

1. Configure the interface to the directly connected device (to-D), and configure the loopback interface.

```

[edit interfaces fe-1/2/0 unit 9]
user@C# set description to-D
user@C# set family inet address 10.10.10.9/30

```

```

[edit interfaces lo0 unit 3]
user@C# set family inet address 192.168.40.4/32

```

2. Configure an EBGP session with Device E.

The **neighbor** statement points to the loopback interface on Device E.

```

[edit protocols bgp group external-peers]
user@C# set type external
user@C# set local-address 192.168.40.4
user@C# set export send-static
user@C# set peer-as 18
user@C# set neighbor 192.168.6.7

```

3. Configure the multihop statement to enable Device C and Device E to become EBGP peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```

[edit protocols bgp group external-peers]
user@C# set multihop ttl 2

```

4. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```

[edit routing-options]
user@C# set static route 10.10.10.14/32 next-hop 10.10.10.10
user@C# set static route 192.168.6.7/32 next-hop 10.10.10.10

```

5. Configure the local router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17

```

6. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@C# set from protocol static
user@C# set then accept
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
fe-1/2/0 {
  unit 9 {
    description to-D;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.40.4/32;
    }
  }
}

user@C# show protocols
bgp {
  group external-peers {
    type external;
    multihop {
      ttl 2;
    }
    local-address 192.168.40.4;
    export send-static;
    peer-as 18;
    neighbor 192.168.6.7;
  }
}

user@C# show policy-options
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@C# show routing-options
static {
  route 10.10.10.14/32 next-hop 10.10.10.10;
  route 192.168.6.7/32 next-hop 10.10.10.10;
}
```

```
router-id 192.168.40.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.
Repeat these steps for all BFD sessions in the topology.

Configuring Device D

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device D:

1. Set the CLI to Device D.

```
user@host> set cli logical-system D
```

2. Configure the interfaces to the directly connected devices, and configure a loopback interface.

```
[edit interfaces fe-1/2/0 unit 10]
user@D# set description to-C
user@D# set family inet address 10.10.10.10/30
```

```
[edit interfaces fe-1/2/1 unit 13]
user@D# set description to-E
user@D# set family inet address 10.10.10.13/30
```

```
[edit interfaces lo0 unit 4]
user@D# set family inet address 192.168.6.6/32
```

3. Configure connectivity to the other devices using static routes to the loopback interface addresses.

On Device D, you do not need static routes to the physical addresses because Device D is directly connected to Device C and Device E.

```
[edit routing-options]
user@D# set static route 192.168.40.4/32 next-hop 10.10.10.9
user@D# set static route 192.168.6.7/32 next-hop 10.10.10.14
```

4. Configure the local router ID.

```
[edit routing-options]
user@D# set router-id 192.168.6.6
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
fe-1/2/0 {
  unit 10 {
    description to-C;
    family inet {
      address 10.10.10.10/30;
```

```
    }  
  }  
}  
fe-1/2/1 {  
  unit 13 {  
    description to-E;  
    family inet {  
      address 10.10.10.13/30;  
    }  
  }  
}  
lo0 {  
  unit 4 {  
    family inet {  
      address 192.168.6.6/32;  
    }  
  }  
}  
}  
  
user@D# show protocols  
  
user@D# show routing-options  
static {  
  route 192.168.40.4/32 next-hop 10.10.10.9;  
  route 192.168.6.7/32 next-hop 10.10.10.14;  
}  
router-id 192.168.6.6;
```

If you are done configuring the device, enter **commit** from configuration mode.
Repeat these steps for all BFD sessions in the topology.

Configuring Device E

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E:

1. Set the CLI to Device E.
2. Configure the interface to the directly connected device (to-D), and configure the loopback interface.

```
[edit interfaces fe-1/2/0 unit 14]  
user@E# set description to-D  
user@E# set family inet address 10.10.10.14/30
```

```
[edit interfaces lo0 unit 5]  
user@E# set family inet address 192.168.6.7/32
```

3. Configure an EBGP session with Device E.

The **neighbor** statement points to the loopback interface on Device C.

```
[edit protocols bgp group external-peers]
```

```

user@E# set local-address 192.168.6.7
user@E# set export send-static
user@E# set peer-as 17
user@E# set neighbor 192.168.40.4

```

4. Configure the **multihop** statement to enable Device C and Device E to become EBGp peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```

[edit protocols bgp group external-peers]
user@E# set multihop ttl 2

```

5. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```

[edit routing-options]
user@E# set static route 10.10.10.8/30 next-hop 10.10.10.13
user@E# set static route 192.168.40.4/32 next-hop 10.10.10.13

```

6. Configure the local router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@E# set router-id 192.168.6.7
user@E# set autonomous-system 18

```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-static term 1]
user@E# set from protocol static
user@E# set then accept

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@E# show interfaces
fe-1/2/0 {
  unit 14 {
    description to-D;
    family inet {
      address 10.10.10.14/30;
    }
  }
}
lo0 {
  unit 5 {
    family inet {
      address 192.168.6.7/32;
    }
  }
}

```

```
}
user@E# show protocols
bgp {
  group external-peers {
    multihop {
      ttl 2;
    }
    local-address 192.168.6.7;
    export send-static;
    peer-as 17;
    neighbor 192.168.40.4;
  }
}

user@E# show policy-options
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@E# show routing-options
static {
  route 10.10.10.8/30 next-hop 10.10.10.13;
  route 192.168.40.4/32 next-hop 10.10.10.13;
}
router-id 192.168.6.7;
autonomous-system 18;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 3324](#)
- [Verifying That BGP Sessions Are Established on page 3325](#)
- [Viewing Advertised Routes on page 3325](#)

Verifying Connectivity

Purpose Make sure that Device C can ping Device E, specifying the loopback interface address as the source of the ping request.

The loopback interface address is the source address that BGP will use.

Action From operational mode, enter the **ping 10.10.10.14 source 192.168.40.4** command from Device C, and enter the **ping 10.10.10.9 source 192.168.6.7** command from Device E.

```
user@C> ping 10.10.10.14 source 192.168.40.4
```

```
PING 10.10.10.14 (10.10.10.14): 56 data bytes
64 bytes from 10.10.10.14: icmp_seq=0 ttl=63 time=1.262 ms
64 bytes from 10.10.10.14: icmp_seq=1 ttl=63 time=1.202 ms
^C
```



```

--- 10.10.10.14 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.202/1.232/1.262/0.030 ms

```

```
user@E> ping 10.10.10.9 source 192.168.6.7
```

```

PING 10.10.10.9 (10.10.10.9): 56 data bytes
64 bytes from 10.10.10.9: icmp_seq=0 ttl=63 time=1.255 ms
64 bytes from 10.10.10.9: icmp_seq=1 ttl=63 time=1.158 ms
^C
--- 10.10.10.9 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.158/1.206/1.255/0.049 ms

```

Meaning The static routes are working if the pings work.

Verifying That BGP Sessions Are Established

Purpose Verify that the BGP sessions are up.

Action From operational mode, enter the `show bgp summary` command.

```
user@C> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0          2          0          0          0          0          0          0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.6.7      18      147      147        0        1    1:04:27
0/2/2/0          0/0/0/0

```

```
user@E> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0          2          0          0          0          0          0          0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.40.4     17      202      202        0        1    1:02:18
0/2/2/0          0/0/0/0

```

Meaning The output shows that both devices have one peer each. No peers are down.

Viewing Advertised Routes

Purpose Check to make sure that routes are being advertised by BGP.

Action From operational mode, enter the `show route advertising-protocol bgp neighbor` command.

```
user@C> show route advertising-protocol bgp 192.168.6.7
```

```

inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED    Lclpref  AS path
* 10.10.10.14/32    Self              0
* 192.168.6.7/32    Self              0

```

```
user@E> show route advertising-protocol bgp 192.168.40.4
```

```
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED      Lclpref   AS path
* 10.10.10.8/30      Self              0
* 192.168.40.4/32    Self              0
```

Meaning The **send-static** routing policy is exporting the static routes from the routing table into BGP. BGP is advertising these routes between the peers because the BGP peer session is established.

Related Documentation

- [Examples: Configuring External BGP Peering on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring BGP Route Preference (Administrative Distance)

- [Understanding Route Preference Values on page 3326](#)
- [Example: Configuring the Preference Value for BGP Routes on page 3327](#)

Understanding Route Preference Values

The Junos OS routing protocol process assigns a default preference value (also known as an *administrative distance*) to each route that the routing table receives. The default value depends on the source of the route. The preference value is a value from 0 through 4,294,967,295 ($2^{32} - 1$), with a lower value indicating a more preferred route.

[Table 285 on page 3326](#) lists the default preference values.

Table 285: Default Route Preference Values

How Route Is Learned	Default Preference	Statement to Modify Default Preference
Directly connected network	0	—
System routes	4	—
Static and Static LSPs	5	<i>static</i>
RSVP-signaled LSPs	7	RSVP preference as described in the <i>Junos OS MPLS Applications Configuration Guide</i>
LDP-signaled LSPs	9	LDP preference , as described in the <i>Junos OS MPLS Applications Configuration Guide</i>
OSPF internal route	10	OSPF preference
IS-IS Level 1 internal route	15	IS-IS preference
IS-IS Level 2 internal route	18	IS-IS preference

Table 285: Default Route Preference Values (*continued*)

How Route Is Learned	Default Preference	Statement to Modify Default Preference
Redirects	30	–
Kernel	40	–
SNMP	50	–
Router discovery	55	–
RIP	100	RIP preference
RIPng	100	RIPng preference
PIM	105	<i>Multicast Protocols Configuration Guide</i>
DVMRP	110	<i>Multicast Protocols Configuration Guide</i>
Aggregate	130	aggregate
OSPF AS external routes	150	OSPF external-preference
IS-IS Level 1 external route	160	IS-IS external-preference
IS-IS Level 2 external route	165	IS-IS external-preference
BGP	170	BGP preference , export , import
MSDP	175	<i>Multicast Protocols Configuration Guide</i>

In general, the narrower the scope of the statement, the higher precedence its preference value is given, but the smaller the set of routes it affects. To modify the default preference value for routes learned by routing protocols, you generally apply routing policy when configuring the individual routing protocols. You also can modify some preferences with other configuration statements, which are indicated in the table.

Example: Configuring the Preference Value for BGP Routes

This example shows how to specify the preference for routes learned from BGP. Routing information can be learned from multiple sources. To break ties among equally specific routes learned from multiple sources, each source has a preference value. Routes that are learned through explicit administrative action, such as static routes, are preferred over routes learned from a routing protocol, such as BGP or OSPF. This concept is called *administrative distance* by some vendors.

- [Requirements on page 3328](#)
- [Overview on page 3328](#)

- [Configuration on page 3329](#)
- [Verification on page 3331](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Routing information can be learned from multiple sources, such as through static configuration, BGP, or an interior gateway protocol (IGP). When Junos OS determines a route's preference to become the active route, it selects the route with the lowest preference as the active route and installs this route into the forwarding table. By default, the routing software assigns a preference of 170 to routes that originated from BGP. Of all the routing protocols, BGP has the highest default preference value, which means that routes learned by BGP are the least likely to become the active route.

Some vendors have a preference (distance) of 20 for external BGP (EBGP) and a distance of 200 for internal BGP (IBGP). Junos OS uses the same value (170) for both EBGP and IBGP. However, this difference between vendors has no operational impact because Junos OS always prefers EBGP routes over IBGP routes.

Another area in which vendors differ is in regard to IGP distance compared to BGP distance. For example, some vendors assign a distance of 110 to OSPF routes. This is higher than the EBGP distance of 20, and results in the selection of an EBGP route over an equivalent OSPF route. In the same scenario, Junos OS chooses the OSPF route, because of the default preference 10 for an internal OSPF route and 150 for an external OSPF route, which are both lower than the 170 preference assigned to all BGP routes.

In a multivendor environment, you might want to change the preference value for BGP routes so that Junos OS chooses an EBGP route instead of an OSPF route. To accomplish this goal, one option is to include the [preference](#) statement in the EBGP configuration. To modify the default BGP preference value, include the **preference** statement, specifying a value from 0 through 4,294,967,295 ($2^{32} - 1$).



TIP: Another way to achieve multivendor compatibility is to include the [advertise-inactive](#) statement in the EBGP configuration. This causes the routing table to export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The `advertise-inactive` statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the `advertise-inactive` statement, the Junos OS device uses the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.

Topology

In the sample network, Device R1 and Device R2 have EBGP routes to each other and also OSPF routes to each other.

This example shows the routing tables in the following cases:

- Accept the default preference values of 170 for BGP and 10 for OSPF.
- Change the BGP preference to 8.

Figure 90 on page 3329 shows the sample network.

Figure 90: BGP Preference Value Topology

ERROR: Unresolved graphic fileref="g041157.svg" not found in
 "//cmsxml/default/main/supplemental/STAGING/images/".

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1  set interfaces fe-1/2/0 unit 4 family inet address 1.12.0.1/30
            set interfaces lo0 unit 2 family inet address 10.255.71.24/32
            set protocols bgp export send-direct
            set protocols bgp group ext type external
            set protocols bgp group ext preference 8
            set protocols bgp group ext peer-as 65000
            set protocols bgp group ext neighbor 1.12.0.2
            set protocols ospf area 0.0.0.0 interface fe-1/2/0.4
            set protocols ospf area 0.0.0.0 interface 10.255.71.24
            set policy-options policy-statement send-direct term 1 from protocol direct
            set policy-options policy-statement send-direct term 1 then accept
            set routing-options autonomous-system 65500

Device R2  set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
            set interfaces lo0 unit 3 family inet address 10.255.14.177/32
            set protocols bgp export send-direct
            set protocols bgp group ext type external
            set protocols bgp group ext peer-as 65500
            set protocols bgp group ext neighbor 1.12.0.1
            set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
            set protocols ospf area 0.0.0.0 interface 10.255.14.177
            set policy-options policy-statement send-direct term 1 from protocol direct
            set policy-options policy-statement send-direct term 1 then accept
            set routing-options autonomous-system 65000
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 4 family inet address 1.12.0.1/30
user@R1# set lo0 unit 2 family inet address 10.255.71.24/32
```
2. Configure the local autonomous system.

```
[edit routing-options]
user@R1# set autonomous-system 65500
```
3. Configure the external peering with Device R2.

```
[edit protocols bgp]
user@R1# set export send-direct
user@R1# set group ext type external
user@R1# set group ext preference 8
user@R1# set group ext peer-as 65000
user@R1# set group ext neighbor 1.12.0.2
```
4. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.4
user@R1# set interface 10.255.71.24
```
5. Configure the routing policy.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 4 {
    family inet {
      address 1.12.0.1/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.255.71.24/32;
    }
  }
}
```

```

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
protocols {
  bgp {
    export send-direct;
    group ext {
      type external;
      preference 8;
      peer-as 65000;
      neighbor 1.12.0.2;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface fe-1/2/0.4;
      interface 10.255.71.24;
    }
  }
}

user@R1# show routing-options
autonomous-system 65500;

```

If you are done configuring the device, enter **commit** from configuration mode.
Repeat these steps on Device R2.

Verification

Confirm that the configuration is working properly.

Verifying the Preference

Purpose Make sure that the routing tables on Device R1 and Device R2 reflect the fact that Device R1 is using the configured EBGp preference of 8, and Device R2 is using the default EBGp preference of 170.

Action From operational mode, enter the **show route** command.

```

user@R1> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.12.0.0/30          *[Direct/0] 3d 07:03:01
                    > via fe-1/2/0.4
                    [BGP/8] 01:04:49, localpref 100
                    AS path: 65000 I
                    > to 1.12.0.2 via fe-1/2/0.4
1.12.0.1/32         *[Local/0] 3d 07:03:01
                    Local via fe-1/2/0.4
10.255.14.177/32    *[BGP/8] 01:04:49, localpref 100
                    AS path: 65000 I

```

```
> to 1.12.0.2 via fe-1/2/0.4
[OSPF/10] 3d 07:02:16, metric 1
> to 1.12.0.2 via fe-1/2/0.4
10.255.71.24/32 * [Direct/0] 3d 07:03:01
> via lo0.2
224.0.0.5/32 * [OSPF/10] 5d 03:42:16, metric 1
MultiRecv

user@R2> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.12.0.0/30 * [Direct/0] 3d 07:03:30
> via fe-1/2/0.6
[BGP/170] 00:45:36, localpref 100
AS path: 65500 I
> to 1.12.0.1 via fe-1/2/0.6
1.12.0.2/32 * [Local/0] 3d 07:03:30
Local via fe-1/2/0.6
10.255.14.177/32 * [Direct/0] 3d 07:03:30
> via lo0.3
10.255.71.24/32 * [OSPF/10] 3d 07:02:45, metric 1
> to 1.12.0.1 via fe-1/2/0.6
[BGP/170] 00:45:36, localpref 100
AS path: 65500 I
> to 1.12.0.1 via fe-1/2/0.6
224.0.0.5/32 * [OSPF/10] 5d 03:42:45, metric 1
MultiRecv
```

Meaning The output shows that on Device R1, the active path to Device R2's loopback interface (10.255.14.177/32) is a BGP route. The output also shows that on Device R2, the active path to Device R1's loopback interface (10.255.71.24/32) is an OSPF route.

Related Documentation

- [Route Preferences Overview](#)
- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring BGP Path Selection

- [Understanding BGP Path Selection on page 3332](#)
- [Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3336](#)

Understanding BGP Path Selection

For each prefix in the routing table, the routing protocol process selects a single best path. After the best path is selected, the route is installed in the routing table. The best path becomes the active route if the same prefix is not learned by a protocol with a lower (more preferred) global preference value, also known as the administrative distance. The algorithm for determining the active route is as follows:

1. Verify that the next hop can be resolved.
2. Choose the path with the lowest preference value (routing protocol process preference).

Routes that are not eligible to be used for forwarding (for example, because they were rejected by routing policy or because a next hop is inaccessible) have a preference of -1 and are never chosen.

3. Prefer the path with higher local preference.

For non-BGP paths, choose the path with the lowest **preference2** value.

4. If the accumulated interior gateway protocol (AIGP) attribute is enabled, prefer the path with the lower AIGP attribute.
5. Prefer the path with the shortest autonomous system (AS) path value (skipped if the **as-path-ignore** statement is configured).

A confederation segment (sequence or set) has a path length of 0. An AS set has a path length of 1.

6. Prefer the route with the lower origin code.

Routes learned from an IGP have a lower origin code than those learned from an exterior gateway protocol (EGP), and both have lower origin codes than incomplete routes (routes whose origin is unknown).

7. Prefer the path with the lowest multiple exit discriminator (MED) metric.

Depending on whether nondeterministic routing table path selection behavior is configured, there are two possible cases:

- If nondeterministic routing table path selection behavior is not configured (that is, if the **path-selection cisco-nondeterministic** statement is not included in the BGP configuration), for paths with the same neighboring AS numbers at the front of the AS path, prefer the path with the lowest MED metric. To always compare MEDs whether or not the peer ASs of the compared routes are the same, include the **path-selection always-compare-med** statement.
- If nondeterministic routing table path selection behavior is configured (that is, the **path-selection cisco-nondeterministic** statement is included in the BGP configuration), prefer the path with the lowest MED metric.

Confederations are not considered when determining neighboring ASs. A missing MED metric is treated as if a MED were present but zero.



NOTE: MED comparison works for single path selection within an AS (when the route does not include an AS path), though this usage is uncommon.

By default, only the MEDs of routes that have the same peer autonomous systems (ASs) are compared. You can configure routing table path selection options to obtain different behaviors.

8. Prefer strictly internal paths, which include IGP routes and locally generated routes (static, direct, local, and so forth).
9. Prefer strictly external BGP (EBGP) paths over external paths learned through internal BGP (IBGP) sessions.

10. Prefer the path whose next hop is resolved through the IGP route with the lowest metric.



NOTE: A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed after the previous step. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.

BGP multipath does not apply to paths that share the same MED-plus-IGP cost yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

11. If both paths are external, prefer the currently active path to minimize route-flapping. This rule is not used if any one of the following conditions is true:
- **path-selection external-router-id** is configured.
 - Both peers have the same router ID.
 - Either peer is a confederation peer.
 - Neither path is the current active path.
12. Prefer a primary route over a secondary route. A primary route is one that belongs to the routing table. A secondary route is one that is added to the routing table through an export policy.
13. Prefer the path from the peer with the lowest router ID. For any path with an originator ID attribute, substitute the originator ID for the router ID during router ID comparison.
14. Prefer the path with the shortest cluster list length. The length is 0 for no list.
15. Prefer the path from the peer with the lowest peer IP address.

Routing Table Path Selection

The shortest AS path step of the algorithm, by default, evaluates the length of the AS path and determines the active path. You can configure an option that enables Junos OS to skip this step of the algorithm by including the **as-path-ignore** option.



NOTE: The **as-path-ignore** option is not supported for routing instances.

To configure routing table path selection behavior, include the **path-selection** statement:

```
path-selection {
  (always-compare-med | cisco-non-deterministic | external-router-id);
  as-path-ignore;
  med-plus-igp {
    igp-multiplier number;
    med-multiplier number;
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Routing table path selection can be configured in one of the following ways:

- Emulate the Cisco IOS default behavior (**cisco-non-deterministic**). This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With **cisco-non-deterministic** mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGp; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGp peer. This allows the routing device to install path 1 as the active path for the route.



NOTE: We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

- Always comparing MEDs whether or not the peer ASs of the compared routes are the same (**always-compare-med**).
- Override the rule that If both paths are external, the currently active path is preferred (**external-router-id**). Continue with the next step (Step 12) in the path-selection process.
- Adding the IGP cost to the next-hop destination to the MED value before comparing MED values for path selection (**med-plus-igp**).

BGP multipath does not apply to paths that share the same MED-plus-IGP cost, yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

Effects of Advertising Multiple Paths to a Destination

BGP advertises only the active path, unless you configure BGP to advertise multiple paths to a destination.

Suppose a routing device has in its routing table four paths to a destination and is configured to advertise up to three paths (**add-path send path-count 3**). The three paths are chosen based on path selection criteria. That is, the three best paths are chosen in

path-selection order. The best path is the active path. This path is removed from consideration and a new best path is chosen. This process is repeated until the specified number of paths is reached.

Example: Ignoring the AS Path Attribute When Selecting the Best Path

If multiple BGP routes to the same destination exist, BGP selects the best path based on the route attributes of the paths. One of the route attributes that affects the best-path decision is the length of the AS paths of each route. Routes with shorter AS paths are preferred over those with longer AS paths. Although not typically practical, some scenarios might require that the AS path length be ignored in the route selection process. This example shows how to configure a routing device to ignore the AS path attribute.

- [Requirements on page 3336](#)
- [Overview on page 3336](#)
- [Configuration on page 3337](#)
- [Verification on page 3341](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

On externally connected routing devices, the purpose of skipping the AS path comparison might be to force an external BGP (EBGP) versus internal BGP (IBGP) decision to remove traffic from your network as soon as possible. On internally connected routing devices, you might want your IBGP-only routers to default to the local externally connected gateway. The local IBGP-only (internal) routers skip the AS path comparison and move down the decision tree to use the closest interior gateway protocol (IGP) gateway (lowest IGP metric). Doing this might be an effective way to force these routers to use a LAN connection instead of their WAN connection.



CAUTION: When you include the `as-path-ignore` statement on a routing device in your network, you might need to include it on all other BGP-enabled devices in your network to prevent routing loops and convergence issues. This is especially true for IBGP path comparisons.

In this example, Device R2 is learning about the loopback interface address on Device R4 (4.4.4.4/32) from Device R1 and Device R3. Device R1 is advertising 4.4.4.4/32 with an AS-path of 1 5 4, and Device R3 is advertising 4.4.4.4/32 with an AS-path of 3 4. Device R2 selects the path for 4.4.4.4/32 from Device R3 as the best path because the AS path is shorter than the AS path from Device R1.

This example modifies the BGP configuration on Device R2 so that the AS-path length is not used in the best-path selection.

Device R1 has a lower router ID (1.1.1.1) than Device R3 (1.1.1.1). If all other path selection criteria are equal (or, as in this case, ignored), the route learned from Device R1 is used.

Because the AS-path attribute is being ignored, the best path is toward Device R1 because of its lower router ID value.

Figure 91 on page 3337 shows the sample topology.

Figure 91: Topology for Ignoring the AS-Path Length

ERROR: Unresolved graphic fileref="g041166.svg" not found in
 "/cmsxml/default/main/supplemental/STAGING/images/".

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```

Device R1
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces fe-1/2/1 unit 10 family inet address 192.168.50.2/24
set interfaces lo0 unit 1 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.2 peer-as 2
set protocols bgp group ext neighbor 192.168.50.1 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.50.1
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1

Device R2
set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.2/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.2/24
set interfaces lo0 unit 2 family inet address 2.2.2.2/32
set protocols bgp path-selection as-path-ignore
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.1 peer-as 1
set protocols bgp group ext neighbor 192.168.20.1 peer-as 3
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.50.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.40.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.30.0/24 next-hop 192.168.20.1

```

```
set routing-options router-id 2.2.2.2
set routing-options autonomous-system 2
```

Device R3

```
set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces fe-1/2/1 unit 5 family inet address 192.168.30.1/24
set interfaces lo0 unit 3 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.20.2 peer-as 2
set protocols bgp group ext neighbor 192.168.30.2 peer-as 4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.20.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.2
set routing-options router-id 3.3.3.3
set routing-options autonomous-system 3
```

Device R4

```
set interfaces fe-1/2/0 unit 6 family inet address 192.168.30.2/24
set interfaces fe-1/2/1 unit 7 family inet address 192.168.40.1/24
set interfaces lo0 unit 4 family inet address 4.4.4.4/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.30.1 peer-as 3
set protocols bgp group ext neighbor 192.168.40.2 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.1
set routing-options router-id 4.4.4.4
set routing-options autonomous-system 4
```

Device R5

```
set interfaces fe-1/2/0 unit 8 family inet address 192.168.40.2/24
set interfaces fe-1/2/1 unit 9 family inet address 192.168.50.1/24
set interfaces lo0 unit 5 family inet address 5.5.5.5/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.40.1 peer-as 4
set protocols bgp group ext neighbor 192.168.50.2 peer-as 1
set policy-options policy-statement send-direct term 1 from protocol direct
```

```

set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.20.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.40.1
set routing-options router-id 5.5.5.5
set routing-options autonomous-system 5

```

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 192.168.10.2/24
user@R2# set fe-1/2/1 unit 3 family inet address 192.168.20.2/24
user@R2# set lo0 unit 2 family inet address 2.2.2.2/32

```
2. Configure EBGp.

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set export send-local
user@R2# set neighbor 192.168.10.1 peer-as 1
user@R2# set neighbor 192.168.20.1 peer-as 3

```
3. Configure the autonomous system (AS) path attribute to be ignored in the Junos OS path selection algorithm.

```

[edit protocols bgp]
user@R2# set path-selection as-path-ignore

```
4. Configure the routing policy.

```

[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-local term 1 from protocol local
user@R2# set policy-statement send-local term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept

```
5. Configure some static routes.

```

[edit routing-options static]
user@R2# set route 192.168.50.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.40.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.30.0/24 next-hop 192.168.20.1

```

6. Configure the autonomous system (AS) number and the router ID.

```
[edit routing-options]
user@R2# set router-id 2.2.2.2
user@R2# set autonomous-system 2
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 192.168.10.2/24;
    }
  }
}
fe-1/2/1 {
  unit 3 {
    family inet {
      address 192.168.20.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 2.2.2.2/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
policy-statement send-local {
  term 1 {
    from protocol local;
    then accept;
  }
}
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@R2# show protocols
bgp {
```



```

path-selection as-path-ignore;
group ext {
  type external;
  export [ send-direct send-static send-local ];
  neighbor 192.168.10.1 {
    peer-as 1;
  }
  neighbor 192.168.20.1 {
    peer-as 3;
  }
}
}

user@R21# show routing-options
static {
  route 192.168.50.0/24 next-hop 192.168.10.1;
  route 192.168.40.0/24 next-hop 192.168.10.1;
  route 192.168.30.0/24 next-hop 192.168.20.1;
}
router-id 2.2.2.2;
autonomous-system 2;

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on the other devices in the network, changing the interface names and IP addresses, as needed.

Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3341](#)

Checking the Neighbor Status

Purpose Make sure that from Device R2, the active path to get to AS 4 is through AS 1 and AS 5, not through AS 3.



NOTE: To verify the functionality of the `as-path-ignore` statement, you might need to run the `restart routing` command to force reevaluation of the active path. This is because for BGP, if both paths are external, the Junos OS behavior is to prefer the currently active path. This behavior helps to minimize route-flapping. Use caution when restarting the routing protocol process in a production network.

Action From operational mode, enter the `restart routing` command.

```

user@R2> restart routing
Routing protocols process started, pid 49396

```

From operational mode, enter the `show route 4.4.4.4 protocol bgp` command.

```

user@R2> show route 4.4.4.4 protocol bgp
inet.0: 12 destinations, 25 routes (12 active, 0 holdown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

```

```
4.4.4.4/32      *[BGP/170] 00:00:12, localpref 100
                  AS path: 154 I
                  > to 192.168.10.1 via fe-1/2/0.2
                  [BGP/170] 00:00:08, localpref 100
                  AS path: 34 I
                  > to 192.168.20.1 via fe-1/2/1.3
```

Meaning The asterisk (*) is next to the path learned from R1, meaning that this is the active path. The AS path for the active path is 1 5 4, which is longer than the AS path (3 4) for the nonactive path learned from Router R3.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Example: Removing Private AS Numbers

- [Understanding Private AS Number Removal from AS Paths on page 3342](#)
- [Example: Removing Private AS Numbers from AS Paths on page 3343](#)

Understanding Private AS Number Removal from AS Paths

By default, when BGP advertises AS paths to remote systems, it includes all AS numbers, including private AS numbers. You can configure the software so that it removes private AS numbers from AS paths. Doing this is useful when any of the following circumstances are true:

- A remote AS for which you provide connectivity is multihomed, but only to the local AS.
- The remote AS does not have an officially allocated AS number.
- It is not appropriate to make the remote AS a confederation member AS of the local AS.

Most companies acquire their own AS number. Some companies also use private AS numbers to connect to their public AS network. These companies might use a different private AS number for each region in which their company does business. In any implementation, announcing a private AS number to the Internet must be avoided. Service providers can use the **remove-private** statement to prevent advertising private AS numbers to the Internet.

In an enterprise scenario, suppose that you have multiple AS numbers in your company, some of which are private AS numbers, and one with a public AS number. The one with a public AS number has a direct connection to the service provider. In the AS that connects directly to the service provider, you can use the **remove-private** statement to filter out any private AS numbers in the advertisements that are sent to the service provider.



CAUTION: Changing configuration statements that affect BGP peers, such as enabling or disabling **remove-private** or renaming a BGP group, resets the

BGP sessions. Changes that affect BGP peers should only be made when resetting a BGP session is acceptable.

The AS numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and later, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The software is preconfigured with knowledge of the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document. The set of AS numbers reserved as private are in the range from 64,512 through 65,534, inclusive.

Example: Removing Private AS Numbers from AS Paths

This example demonstrates the removal of a private AS number from the advertised AS path to avoid announcing the private AS number to the Internet.

- [Requirements on page 3343](#)
- [Overview on page 3343](#)
- [Configuration on page 3344](#)
- [Verification on page 3346](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Service providers and enterprise networks use the `remove-private` statement to prevent advertising private AS numbers to the Internet. The `remove-private` statement works in the outbound direction. You configure the `remove-private` statement on a device that has a public AS number and that is connected to one or more devices that have private AS numbers. Generally, you would not configure this statement on a device that has a private AS number.

[Figure 92 on page 3344](#) shows the sample topology.

Figure 92: Topology for Removing a Private AS from the Advertised AS Path

ERROR: Unresolved graphic fileref="g041165.svg" not found in
"//cmsxml/default/main/supplemental/STAGING/images/".

In this example, Device R1 is connected to its service provider using private AS number 65535. The example shows the **remove-private** statement configured on Device ISP to prevent Device R1's private AS number from being announced to Device R2. Device R2 sees only the AS number of the service provider.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces lo0 unit 1 family inet address 10.10.10.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.10.10
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.10
set routing-options autonomous-system 65535
```

Device ISP

```
set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.10/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.20/24
set interfaces lo0 unit 2 family inet address 10.10.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext neighbor 192.168.10.1 peer-as 65535
set protocols bgp group ext neighbor 192.168.20.1 remove-private
set protocols bgp group ext neighbor 192.168.20.1 peer-as 200
set routing-options autonomous-system 100
```

Device R2

```
set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces lo0 unit 3 family inet address 10.10.20.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.20.20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.20
set routing-options autonomous-system 200
```

Device ISP

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device ISP:

1. Configure the interfaces.

```
[edit interfaces]
user@ISP# set fe-1/2/0 unit 2 family inet address 192.168.10.10/24
user@ISP# set fe-1/2/1 unit 3 family inet address 192.168.20.20/24
user@ISP# set lo0 unit 2 family inet address 10.10.0.1/32
```
2. Configure EBGP.

```
[edit protocols bgp group ext]
user@ISP# set type external
user@ISP# set neighbor 192.168.10.1 peer-as 65535
user@ISP# set neighbor 192.168.20.1 peer-as 200
```
3. For the neighbor in autonomous system (AS) 200 (Device R2), remove private AS numbers from the advertised AS paths.

```
[edit protocols bgp group ext]
user@ISP# set neighbor 192.168.20.1 remove-private
```
4. Configure the AS number.

```
[edit routing-options]
user@ISP# set autonomous-system 100
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@ISP# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 192.168.10.10/24;
    }
  }
}
fe-1/2/1 {
  unit 3 {
    family inet {
      address 192.168.20.20/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.10.0.1/32;
    }
  }
}
```

```
}  
}  
  
user@ISP# show protocols  
bgp {  
  group ext {  
    type external;  
    neighbor 192.168.10.1 {  
      peer-as 65535;  
    }  
    neighbor 192.168.20.1 {  
      remove-private;  
      peer-as 200;  
    }  
  }  
}  
  
user@ISP# show routing-options  
autonomous-system 100;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R1 and Device R2, changing the interface names and IP address, as needed, and adding the routing policy configuration.

Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3346](#)
- [Checking the Routing Tables on page 3347](#)
- [Checking the AS Path When the remove-private Statement Is Deactivated on page 3348](#)

Checking the Neighbor Status

Purpose Make sure that Device ISP has the **remove-private** setting enabled in its neighbor session with Device R2.

Action From operational mode, enter the **show bgp neighbor 192.168.20.1** command.

```
user@ISP> show bgp neighbor 192.168.20.1  
Peer: 192.168.20.1+179 AS 200 Local: 192.168.20.20+60216 AS 100  
Type: External State: Established Flags: <ImportEval Sync>  
Last State: OpenConfirm Last Event: RecvKeepAlive  
Last Error: None  
Options: <Preference RemovePrivateAS PeerAS Refresh>  
Holdtime: 90 Preference: 170  
Number of flaps: 0  
Peer ID: 10.10.20.1 Local ID: 10.10.0.1 Active Holdtime: 90  
Keepalive Interval: 30 Peer index: 0  
BFD: disabled, down  
Local Interface: fe-1/2/1.3  
NLRI for restart configured on peer: inet-unicast  
NLRI advertised by peer: inet-unicast  
NLRI for this session: inet-unicast  
Peer supports Refresh capability (2)  
Stale routes from peer are kept for: 300  
Peer does not support Restarter functionality  
NLRI that restart is negotiated for: inet-unicast
```

```

NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 200)
Peer does not support Addpath
Table inet.0 Bit: 10001
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        3
  Accepted prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 10   Sent 16   Checked 55
Input messages: Total 54   Updates 3   Refreshes 0   Octets 1091
Output messages: Total 54   Updates 1   Refreshes 0   Octets 1118
Output Queue[0]: 0

```

Meaning The `RemovePrivateAS` option shows that Device ISP has the expected setting.

Checking the Routing Tables

Purpose Make sure that the devices have the expected routes and AS paths.

Action From operational mode, enter the `show route protocol bgp` command.

```

user@R1> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

10.10.20.1/32      *[BGP/170] 00:28:57, localpref 100
                   AS path: 100 200 I
                   > to 192.168.10.10 via fe-1/2/0.1

```

```

user@ISP> show route protocol bgp

```

```

inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

10.10.10.1/32      *[BGP/170] 00:29:40, localpref 100
                   AS path: 65535 I
                   > to 192.168.10.1 via fe-1/2/0.2
10.10.20.1/32      *[BGP/170] 00:29:36, localpref 100
                   AS path: 200 I
                   > to 192.168.20.1 via fe-1/2/1.3
192.168.10.0/24    [BGP/170] 00:29:40, localpref 100
                   AS path: 65535 I
                   > to 192.168.10.1 via fe-1/2/0.2
192.168.20.0/24    [BGP/170] 00:29:36, localpref 100
                   AS path: 200 I
                   > to 192.168.20.1 via fe-1/2/1.3

```

```

user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

10.10.10.1/32      *[BGP/170] 00:29:53, localpref 100
                   AS path: 100 I
                   > to 192.168.20.20 via fe-1/2/0.4

```

Meaning Device ISP has the private AS number 65535 in its AS path to Device R1. However, Device ISP does not advertise this private AS number to Device R2. This is shown in the routing table of Device R2. Device R2's path to Device R1 contains only the AS number for Device ISP.

Checking the AS Path When the remove-private Statement Is Deactivated

Purpose Verify that without the **remove-private** statement, the private AS number appears in Device R2's routing table.

Action From configuration mode on Device ISP, enter the **deactivate remove-private** command and then recheck the routing table on Device R2.

```
[protocols bgp group ext neighbor 192.168.20.1]
user@ISP# deactivate remove-private
user@ISP# commit
```

```
user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32      *[BGP/170] 00:00:54, localpref 100
                  AS path: 100 65535 I
                  > to 192.168.20.20 via fe-1/2/0.4
```

Meaning Private AS number 65535 appears in Device R2's AS path to Device R1.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

BGP BFD Configuration

- [Example: Configuring BFD for BGP on page 3348](#)
- [Example: Configuring BFD Authentication for BGP on page 3357](#)

Example: Configuring BFD for BGP

- [Understanding BFD for BGP on page 3348](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3349](#)

Understanding BFD for BGP

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than default failure detection mechanisms for BGP, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions. In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only. In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.

Example: Configuring BFD on Internal BGP Peer Sessions

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 3349](#)
- [Overview on page 3349](#)
- [Configuration on page 3350](#)
- [Verification on page 3354](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

The minimum configuration to enable BFD on IBGP sessions is to include the **bfd-liveness-detection minimum-interval** statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see [bfd-liveness-detection](#).



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and less than 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 93 on page 3350 shows a typical network with internal peer sessions.

Figure 93: Typical Network with IBGP Sessions

ERROR: Unresolved graphic fileref="g040732.svg" not found in
 "/cmsxml/default/main/supplemental/STAGING/images/".

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```

set logical-systems A interfaces lt-1/2/0 unit 1 description to-B
set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
  minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
  
```

```

set logical-systems A policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17

```

Device B

```

set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C
set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
  minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17

```

Device C

```

set logical-systems C interfaces lt-1/2/0 unit 6 description to-B
set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers bfd-liveness-detection
  minimum-interval 1000
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

Configuring Device A

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.

```
user@host> set cli logical-system A
```

2. Configure the interfaces.

```
[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
user@host:A# set peer-unit 2
user@host:A# set family inet address 10.10.1/30
```

```
[edit interfaces lo0 unit 1]
user@host:A# set family inet address 192.168.6.5/32
```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@host:A# set type internal
user@host:A# set local-address 192.168.6.5
user@host:A# set export send-direct
user@host:A# set neighbor 192.163.6.4
user@host:A# set neighbor 192.168.40.4
```

4. Configure BFD.

```
[edit protocols bgp group internal-peers]
user@host:A# set bfd-liveness-detection minimum-interval 1000
```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```
[edit protocols bgp group internal-peers]
user@host:A# set traceoptions file bgp-bfd
user@host:A# set traceoptions flag bfd detail
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@host:A# set interface lo0.1 passive
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
```

```
user@host:A# set from protocol direct
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:A# set router-id 192.168.6.5
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:A# show interfaces
lt-1/2/0 {
  unit 1 {
    description to-B;
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@host:A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@host:A# show protocols
bgp {
  group internal-peers {
    type internal;
    traceoptions {
      file bgp-bfd;
      flag bfd detail;
    }
    local-address 192.168.6.5;
    export send-direct;
    bfd-liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

```
        neighbor 192.163.6.4;
        neighbor 192.168.40.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.1 {
            passive;
        }
        interface lt-1/2/0.1;
    }
}

user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;
```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 3354](#)
- [Verifying That BFD Sessions Are Up on page 3354](#)
- [Viewing Detailed BFD Events on page 3355](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 3356](#)

Verifying That BFD Is Enabled

Purpose Verify that BFD is enabled between the IBGP peers.

Action From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```
user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10
```

Meaning The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled, down**. The output also shows that BFD-related events are being written to a log file because trace operations are configured.

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@host:A> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.163.6.4	Up		3.000	1.000	3

Client BGP, TX interval 1.000, RX interval 1.000
 Session up time 00:54:40
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 12, routing table index 25
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 10, remote discriminator 9
 Echo mode disabled/inactive
 Multi-hop route table 25, local-address 192.168.6.5

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.40.4	Up		3.000	1.000	3

Client BGP, TX interval 1.000, RX interval 1.000
 Session up time 00:48:03
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 12, routing table index 25
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 14, remote discriminator 13
 Echo mode disabled/inactive
 Multi-hop route table 25, local-address 192.168.6.5

2 sessions, 2 clients

Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

Meaning The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
```

```

Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor
192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up

```

Meaning Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface

Purpose Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

Action 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal
AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS
17): No route to host

```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.


```
user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit
```

4. From operational mode, enter the `file show /var/log/A/bgp-bfd` command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capabilty to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up
```

**Related
Documentation**

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Example: Configuring BFD Authentication for BGP

- [Understanding BFD Authentication for BGP on page 3357](#)
- [Example: Configuring BFD Authentication for BGP on page 3359](#)

Understanding BFD Authentication for BGP

Bidirectional Forwarding Detection protocol (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over BGP. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3358](#)
- [Security Authentication Keychains on page 3358](#)
- [Strict Versus Loose Authentication on page 3359](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for BGP

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over BGP. Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the BGP protocol.
2. Associate the authentication keychain with the BGP protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on BGP:

- [Configuring BFD Authentication Parameters on page 3359](#)
- [Viewing Authentication Information for BFD Sessions on page 3361](#)

Configuring BFD Authentication Parameters

BFD authentication can be configured for the entire BGP protocol, or a specific BGP group, neighbor, or routing instance.

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.

[edit]

```
user@host# set protocols bgp bfd-liveness-detection authentication algorithm
keyed-sha-1
```

```
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
algorithm keyed-sha-1
```

```
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication algorithm keyed-sha-1
```



NOTE: Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on BGP with the unique security authentication keychain attributes.

The keychain name you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication keychain bfd-bgp
```



NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
 - The matching keychain name as specified in Step 2.
 - At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
 - The secret data used to allow access to the session.
 - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-bgp key 53 secret
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication loose-check
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
loose-check
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat these steps to configure the other end of the BFD session.



NOTE: BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **bgp-gr1** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-bgp**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols bgp]
group bgp-gr1 {
  bfd-liveness-detection {
    authentication {
      algorithm keyed-sha-1;
      key-chain bfd-bgp;
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-bgp {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9L.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3
Client BGP, TX interval 0.300, RX interval 0.300, Authenticate					

```
Session up time 3d 00:34
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
```

show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3

```
Client BGP, TX interval 0.300, RX interval 0.300, Authenticate
keychain bfd-bgp, algo keyed-sha-1, mode strict
Session up time 00:04:42
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-bgp, algo keyed-sha-1, mode strict
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3149](#)
 - [BGP Configuration Overview](#)

BGP Load Balancing Configuration

- [Examples: Configuring BGP Multipath on page 3362](#)
- [Example: Advertising Multiple BGP Paths to a Destination on page 3379](#)

Examples: Configuring BGP Multipath

- [Understanding BGP Multipath on page 3362](#)
- [Example: Load Balancing BGP Traffic on page 3363](#)
- [Example: Configuring Single-Hop EBGPeers to Accept Remote Next Hops on page 3367](#)

Understanding BGP Multipath

The Junos OS BGP multipath feature supports the following applications:

- Load balancing across multiple links between two routing devices belonging to different autonomous systems (ASs)
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to the same peer AS
- Load balancing across multiple links between two routing devices belonging to different external confederation peers
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to external confederation peers

In a common scenario for load balancing, a customer is multihomed to multiple routers in a point of presence (POP). The default behavior is to send all traffic across only one of the available links. Load balancing causes traffic to use two or more of the links.

BGP multipath does not apply to paths that share the same MED-plus-IGP cost, yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

Example: Load Balancing BGP Traffic

This example shows how to configure BGP to select multiple equal-cost external BGP (EBGP) or internal BGP (IBGP) paths as active paths.

- [Requirements on page 3363](#)
- [Overview on page 3363](#)
- [Configuration on page 3364](#)
- [Verification on page 3366](#)

Requirements

Before you begin:

- Configure the device interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure a routing policy that exports routes (such as direct routes or IGP routes) from the routing table into BGP.

Overview

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
  from {
    match-conditions;
    route-filter destination-prefix match-type <actions>;
    prefix-list name;
  }
  then {
    load-balance per-packet;
  }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {
  export policy-name;
}
```

You cannot apply the export policy to VRF routing instances.


3. Specify all next-hops of that route, if more than one exists, when allocating a label corresponding to a route that is being advertised.
4. Configure the forwarding-options hash key for MPLS to include the IP payload.

In this example, Device R1 is in AS 65000 and is connected to both Device R2 and Device R3, which are in AS 65001. This example shows the configuration on Device R1.

Topology

Figure 94 on page 3364 shows the topology used in this example.

Figure 94: BGP Load Balancing



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set protocols bgp group external type external
set protocols bgp group external peer-as 65001
set protocols bgp group external multipath
set protocols bgp group external neighbor 10.0.1.1
set protocols bgp group external neighbor 10.0.0.2
set policy-options policy-statement loadbal from route-filter 10.0.0.0/16 orlonger
set policy-options policy-statement loadbal then load-balance per-packet
set routing-options forwarding-table export loadbal
set routing-options autonomous-system 65000

```


Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the BGP group.


```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set peer-as 65001
user@R1# set neighbor 10.0.1.1
user@R1# set neighbor 10.0.0.2
```
2. Enable the BGP group to use multiple paths.



NOTE: To disable the default check requiring that paths accepted by BGP multipath must have the same neighboring autonomous system (AS), include the `multiple-as` option.

```
[edit protocols bgp group external]
user@R1# set multipath
```

3. Configure the load-balancing policy.


```
[edit policy-options policy-statement loadbal]
user@R1# set from route-filter 10.0.0.0/16 orlonger
user@R1# set then load-balance per-packet
```
4. Apply the load-balancing policy.


```
[edit routing-options]
user@R1# set forwarding-table export loadbal
```
5. Configure the local autonomous system (AS) number.


```
[edit routing-options]
user@R1# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show protocols
bgp {
  group external {
    type external;
    peer-as 65001;
    multipath;
    neighbor 10.0.1.1;
    neighbor 10.0.0.2;
  }
}
```

```
[edit]
user@R1# show policy-options
policy-statement loadbal {
  from {
    route-filter 10.0.0.0/16 orlonger;
  }
  then {
    load-balance per-packet;
  }
}

[edit]
user@R1# show routing-options
autonomous-system 65000;
forwarding-table {
  export loadbal;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly:

- [Verifying Routes on page 3366](#)
- [Verifying Forwarding on page 3367](#)

Verifying Routes

Purpose Verify that routes are learned from both routers in the neighboring AS.

Action From operational mode, run the **show route** command.

```
user@R1> show route 10.0.2.0
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.2.0/30          *[BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    to 10.0.1.1 via ge-1/2/0.0
                    > to 10.0.0.2 via ge-1/2/1.0
                    [BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    > to 10.0.1.1 via ge-1/2/0.0

user@R1> show route 10.0.2.0 detail
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
              Next hop type: Router, Next hop index: 262142
              Next-hop reference count: 3
              Source: 10.0.0.2
              Next hop: 10.0.1.1 via ge-1/2/0.0
              Next hop: 10.0.0.2 via ge-1/2/1.0, selected
              State: <Active Ext>
              Local AS: 65000 Peer AS: 65001
              Age: 3:18:30
              Task: BGP_65001.10.0.0.2+55402
```

```

BGP
  Announcement bits (1): 2-KRT
  AS path: 65001 I
  Accepted Multipath
  Localpref: 100
  Router ID: 192.168.2.1
  Preference: 170/-101
  Next hop type: Router, Next hop index: 602
  Next-hop reference count: 5
  Source: 10.0.1.1
  Next hop: 10.0.1.1 via ge-1/2/0.0, selected
  State: <NotBest Ext>
  Inactive reason: Not Best in its group - Active preferred
  Local AS: 65000 Peer AS: 65001
  Age: 3:18:30
  Task: BGP_65001.10.0.1.1+53135
  AS path: 65001 I
  Accepted
  Localpref: 100
  Router ID: 192.168.3.1

```

Meaning The active path, denoted with an asterisk (*), has two next hops: 10.0.1.1 and 10.0.0.2 to the 10.0.2.0 destination. The 10.0.1.1 next hop is copied from the inactive path to the active path.

Verifying Forwarding

Purpose Verify that both next hops are installed in the forwarding table.

Action From operational mode, run the **show route forwarding-table** command.

```

user@R1> show route forwarding-table destination 10.0.2.0
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
10.0.2.0/30      user   0          10.0.1.1      ucst  602   5 ge-1/2/0.0
                  10.0.0.2      ucst  522   6 ge-1/2/1.0

```

Example: Configuring Single-Hop EBGPeers to Accept Remote Next Hops

This example shows how to configure a single-hop external BGP (EBGP) peer to accept a remote next hop with which it does not share a common subnet.

- [Requirements on page 3367](#)
- [Overview on page 3368](#)
- [Configuration on page 3369](#)
- [Verification on page 3377](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

In some situations, it is necessary to configure a single-hop EBGP peer to accept a remote next hop with which it does not share a common subnet. The default behavior is for any next-hop address received from a single-hop EBGP peer that is not recognized as sharing a common subnet to be discarded. The ability to have a single-hop EBGP peer accept a remote next hop to which it is not directly connected also prevents you from having to configure the single-hop EBGP neighbor as a multihop session. When you configure a multihop session in this situation, all next-hop routes learned through this EBGP peer are labeled indirect even when they do share a common subnet. This situation breaks multipath functionality for routes that are recursively resolved over routes that include these next-hop addresses. Configuring the `accept-remote-nexthop` statement allows a single-hop EBGP peer to accept a remote next hop, which restores multipath functionality for routes that are resolved over these next-hop addresses. You can configure this statement at the global, group, and neighbor hierarchy levels for BGP. The statement is also supported on logical systems and the VPN routing and forwarding (VRF) routing instance type. Both the remote next-hop and the EBGP peer must support BGP route refresh as defined in RFC 2918, *Route Refresh Capability in BGP-4*. If the remote peer does not support BGP route refresh, the session is reset.




NOTE: You cannot configure both the `multihop` and `accept-remote-nexthop` statements for the same EBGP peer.

When you enable a single-hop EBGP peer to accept a remote next hop, you must also configure an import routing policy on the EBGP peer that specifies the remote next-hop address.

This example includes an import routing policy, `agg_route`, that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network. At the `[edit protocols bgp]` hierarchy level, the example includes the `import agg_route` statement to apply the policy to the external BGP peer and includes the `accept-remote-nexthop` statement to enable the single-hop EBGP peer to accept the remote next hop.

Figure 95 on page 3369 shows the sample topology.

Figure 95: Topology for Accepting a Remote Next Hop

**Configuration**

- [Device R0 on page 3370](#)
- [Configuring Device R1 on page 3373](#)
- [Configuring Device R2 on page 3375](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R0

```
set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
set interfaces fe-1/2/1 unit 2 family inet address 1.1.1.1/30
set interfaces lo0 unit 1 family inet address 10.255.14.179/32
set protocols bgp group ext type external
set protocols bgp group ext export test_route
set protocols bgp group ext export agg_route
set protocols bgp group ext peer-as 65000
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.2
set protocols bgp group ext neighbor 1.1.1.2
set policy-options policy-statement agg_route term 1 from protocol static
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then accept
set policy-options policy-statement test_route term 1 from protocol static
set policy-options policy-statement test_route term 1 from route-filter 1.1.10.10/32 exact
set policy-options policy-statement test_route term 1 then accept
set routing-options static route 1.1.10.10/32 reject
```

```
set routing-options static route 1.1.230.0/23 reject
set routing-options autonomous-system 65500
```

Device R1

```
set interfaces fe-1/2/0 unit 3 family inet address 1.1.0.2/30
set interfaces fe-1/2/1 unit 4 family inet address 1.12.0.1/30
set interfaces fe-1/2/2 unit 5 family inet address 1.1.1.2/30
set interfaces lo0 unit 2 family inet address 10.255.71.24/32
set protocols bgp accept-remote-nexthop
set protocols bgp group ext type external
set protocols bgp group ext import agg_route
set protocols bgp group ext peer-as 65500
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.1
set protocols bgp group ext neighbor 1.1.1.1
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.71.24
set protocols bgp group int neighbor 10.255.14.177
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface 10.255.71.24
set policy-options policy-statement agg_route term 1 from protocol bgp
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then next-hop 1.1.10.10
set policy-options policy-statement agg_route term 1 then accept
set routing-options autonomous-system 65000
```

Device R2

```
set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
set interfaces lo0 unit 3 family inet address 10.255.14.177/32
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.177
set protocols bgp group int neighbor 10.255.71.24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface 10.255.14.177
set routing-options autonomous-system 65000
```

Device R0

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R0:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 1]
user@R0# set family inet address 1.1.0.1/30

[edit interfaces fe-1/2/1 unit 2]
user@R0# set family inet address 1.1.1.1/30

[edit interfaces lo0 unit 1]
user@R0# set family inet address 10.255.14.179/32
```
2. Configure EBGp.

```
[edit protocols bgp group ext]
```

```

user@R0# set type external
user@R0# set peer-as 65000
user@R0# set neighbor 1.1.0.2
user@R0# set neighbor 1.1.1.2

```

3. Enable multipath BGP between Device R0 and Device R1.

```

[edit protocols bgp group ext]
user@R0# set multipath

```

4. Configure static routes to remote networks.
These routes are not part of the topology. The purpose of these routes is to demonstrate the functionality in this example.

```

[edit routing-options]
user@R0# set static route 1.1.10.10/32 reject
user@R0# set static route 1.1.230.0/23 reject

```

5. Configure routing policies that accept the static routes.

```

[edit policy-options policy-statement agg_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.230.0/23 exact
user@R0# set then accept

```

```

[edit policy-options policy-statement test_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.10.10/32 exact
user@R0# set then accept

```

6. Export the **agg_route** and **test_route** policies from the routing table into BGP.

```

[edit protocols bgp group ext]
user@R0# set export test_route
user@R0# set export agg_route

```

7. Configure the autonomous system (AS) number.

```

[edit routing-options]
user@R0# set autonomous-system 65500

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 1.1.1.1/30;
    }
  }
}

```

```
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.255.14.179/32;
    }
  }
}

user@R0# show policy-options
policy-statement agg_route {
  term 1 {
    from {
      protocol static;
      route-filter 1.1.230.0/23 exact;
    }
    then accept;
  }
}
policy-statement test_route {
  term 1 {
    from {
      protocol static;
      route-filter 1.1.10.10/32 exact;
    }
    then accept;
  }
}

user@R0# show protocols
bgp {
  group ext {
    type external;
    export [ test_route agg_route ];
    peer-as 65000;
    multipath;
    neighbor 1.1.0.2;
    neighbor 1.1.1.2;
  }
}

user@R0# show routing-options
static {
  route 1.1.10.10/32 reject;
  route 1.1.230.0/23 reject;
}
autonomous-system 65500;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.


```
[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 1.1.0.2/30

[edit interfaces fe-1/2/1 unit 4]
user@R1# set family inet address 1.12.0.1/30

[edit interfaces fe-1/2/2 unit 5]
user@R1# set family inet address 1.1.1.2/30

[edit interfaces lo0 unit 2]
user@R1# set family inet address 10.255.71.24/32
```
2. Configure OSPF.


```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/1.4
user@R1# set interface 10.255.71.24
```
3. Enable Device R1 to accept the remote next hop.


```
[edit protocols bgp]
user@R1# set accept-remote-nexthop
```
4. Configure IBGP.


```
[edit protocols bgp group int]
user@R1# set type internal
user@R1# set local-address 10.255.71.24
user@R1# set neighbor 10.255.14.177
```
5. Configure EBGP.


```
[edit protocols bgp group ext]
user@R1# set type external
user@R1# set peer-as 65500
user@R1# set neighbor 1.1.0.1
user@R1# set neighbor 1.1.1.1
```
6. Enable multipath BGP between Device R0 and Device R1.


```
[edit protocols bgp group ext]
user@R1# set multipath
```
7. Configure a routing policy that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network.


```
[edit policy-options policy-statement agg_route term 1]
user@R1# set from protocol bgp
user@R1# set from route-filter 1.1.230.0/23 exact
```

```
user@R1# set then next-hop 1.1.10.10
user@R1# set then accept
```

8. Import the **agg_route** policy into the routing table on Device R1.

```
[edit protocols bgp group ext]
user@R1# set import agg_route
```

9. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 1.1.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 1.12.0.1/30;
    }
  }
}
fe-1/2/2 {
  unit 5 {
    family inet {
      address 1.1.1.2/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.255.71.24/32;
    }
  }
}

user@R1# show policy-options
policy-statement agg_route {
  term 1 {
    from {
      protocol bgp;
      route-filter 1.1.230.0/23 exact;
    }
    then {
      next-hop 1.1.10.10;
    }
  }
}
```

```

        accept;
    }
}
}

user@R1# show protocols
bgp {
    accept-remote-nexthop;
    group ext {
        type external;
        import agg_route;
        peer-as 65500;
        multipath;
        neighbor 1.1.0.1;
        neighbor 1.1.1.1;
    }
    group int {
        type internal;
        local-address 10.255.71.24;
        neighbor 10.255.14.177;
    }
}
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/1.4;
        interface 10.255.71.24;
    }
}
}

user@R1# show routing-options
autonomous-system 65000;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device R2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.


```

[edit interfaces fe-1/2/0 unit 6]
user@R2# set family inet address 1.12.0.2/30

[edit interfaces lo0 unit 3]
user@R2# set family inet address 10.255.14.177/32

```
2. Configure OSPF.


```

[edit protocols ospf area 0.0.0.0]
user@R2# set interface fe-1/2/0.6
user@R2# set interface 10.255.14.177

```
3. Configure IBGP.

```
[edit protocols bgp group int]
user@R2# set type internal
user@R2# set local-address 10.255.14.177
user@R2# set neighbor 10.255.71.24
```

4. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 6 {
    family inet {
      address 1.12.0.2/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 10.255.14.177/32;
    }
  }
}

user@R2# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.177;
    neighbor 10.255.71.24;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.6;
    interface 10.255.14.177;
  }
}

user@R2# show routing-options
autonomous-system 65000;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table on page 3377](#)
- [Deactivating and Reactivating the accept-remote-nexthop Statement on page 3378](#)

Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table

Purpose Verify that Device R1 has a route to the 1.1.230.0/23 network.

Action From operational mode, enter the **show route 1.1.230.0 extensive** command.

```

user@R1> show route 1.1.230.0 extensive
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
1.1.230.0/23 (2 entries, 1 announced)
TSI:
KRT in-kernel 1.1.230.0/23 -> {indirect(262142)}
Page 0 idx 1 Type 1 val 9168f6c
  Nexthop: 1.1.10.10
  Localpref: 100
  AS path: [65000] 65500 I
  Communities:
Path 1.1.230.0 from 1.1.0.1 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x90c44d8
    Next-hop reference count: 4
    Source: 1.1.0.1
    Next hop type: Router, Next hop index: 262143
    Next hop: 1.1.0.1 via fe-1/2/0.3, selected
    Next hop: 1.1.1.1 via fe-1/2/2.5
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    State: <Active Ext>
    Local AS: 65000 Peer AS: 65500
    Age: 2:55:31 Metric2: 0
    Task: BGP_65500.1.1.0.1+64631
    Announcement bits (3): 2-KRT 3-BGP_RT_Background 4-Resolve tree
1
    AS path: 65500 I
    Accepted Multipath
    Localpref: 100
    Router ID: 10.255.14.179
    Indirect next hops: 1
      Protocol next hop: 1.1.10.10
      Indirect next hop: 91c0000 262142
      Indirect path forwarding next hops: 2
        Next hop type: Router
        Next hop: 1.1.0.1 via fe-1/2/0.3
        Next hop: 1.1.1.1 via fe-1/2/2.5
      1.1.10.10/32 Originating RIB: inet.0
      Node path count: 1
      Forwarding nexthops: 2
        Nexthop: 1.1.0.1 via fe-1/2/0.3
        Nexthop: 1.1.1.1 via fe-1/2/2.5
  BGP Preference: 170/-101

```

```

Next hop type: Indirect
Address: 0x90c44d8
Next-hop reference count: 4
Source: 1.1.1.1
Next hop type: Router, Next hop index: 262143
Next hop: 1.1.0.1 via fe-1/2/0.3, selected
Next hop: 1.1.1.1 via fe-1/2/2.5
Protocol next hop: 1.1.10.10
Indirect next hop: 91c0000 262142
State: <NotBest Ext>
Inactive reason: Not Best in its group - Update source
Local AS: 65500 Peer AS: 65500
Age: 2:55:27 Metric2: 0
Task: BGP_65500.1.1.1.1+53260
AS path: 65500 I
Accepted
Localpref: 100
Router ID: 10.255.14.179
Indirect next hops: 1
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    Indirect path forwarding next hops: 2
        Next hop type: Router
        Next hop: 1.1.0.1 via fe-1/2/0.3
        Next hop: 1.1.1.1 via fe-1/2/2.5
    1.1.10.10/32 Originating RIB: inet.0
    Node path count: 1
    Forwarding nexthops: 2
        Nexthop: 1.1.0.1 via fe-1/2/0.3
        Nexthop: 1.1.1.1 via fe-1/2/2.5

```

Meaning The output shows that Device R1 has a route to the 1.1.230.0 network with the multipath feature enabled (**Accepted Multipath**). The output also shows that the route has an indirect next hop of 1.1.10.10.

Deactivating and Reactivating the accept-remote-nexthop Statement

Purpose Make sure that the multipath route with the indirect next hop is removed from the routing table when you deactivate the **accept-remote-nexthop** statement.

Action 1. From configuration mode, enter the **deactivate protocols bgp accept-remote-nexthop** command.

```

user@R1# deactivate protocols bgp accept-remote-nexthop
user@R1# commit

```

2. From operational mode, enter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

3. From configuration mode, reactivate the statement by entering the **activate protocols bgp accept-remote-nexthop** command.

```

user@R1# activate protocols bgp accept-remote-nexthop
user@R1# commit

```

4. From operational mode, reenter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

```

inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)

```

Restart Complete

+ = Active Route, - = Last Active, * = Both

```
1.1.230.0/23      *[BGP/170] 03:13:19, localpref 100
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5
                  [BGP/170] 03:13:15, localpref 100, from 1.1.1.1
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5
```

Meaning When the **accept-remote-nexthop** statement is deactivated, the multipath route to the 1.1.230.0 network is removed from the routing table .

Related Documentation • *Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches*

Example: Advertising Multiple BGP Paths to a Destination

- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3379](#)
- [Example: Advertising Multiple Paths in BGP on page 3380](#)

Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

BGP peers advertise routes to each other in update messages. BGP stores its routes in the Junos OS routing table (**inet.0**). For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

Instead of advertising only the active path to a destination, you can configure BGP to advertise multiple paths to the destination. Within an autonomous system (AS), the availability of multiple exit points to reach a destination provides the following benefits:

- **Fault tolerance**—Path diversity leads to reduction in restoration time after failure. For instance, a border router after receiving multiple paths to the same destination can precompute a backup path and have it ready so that when the primary path becomes invalid, the border router can use the backup to quickly restore connectivity. Without a backup path, the restoration time depends on BGP reconvergence, which includes withdraw and advertisement messages in the network before a new best path can be learned.
- **Load balancing**—The availability of multiple paths to reach the same destination enables load balancing of traffic, if the routing within the AS meets certain constraints.
- **Maintenance**—The availability of alternate exit points allows for graceful maintenance operation of routers.

The following limitations apply to advertising multiple routes in BGP:

- Address families supported:

- IPv4 unicast (**family inet unicast**)
- IPv6 unicast (**family inet6 unicast**)
- IPv4 labeled unicast (**family inet labeled-unicast**)
- IPv6 labeled unicast (**family inet6 labeled-unicast**)
- Internal BGP (IBGP) peers only. No support on external BGP (EBGP) peers.
- Master instance only. No support for routing instances.
- Graceful restart supported, but not nonstop active routing (NSR).
- No BGP Monitoring Protocol (BMP) support.
- No support for EBGP sessions between confederations.
- Prefix policies enable you to filter routes on a router that is configured to advertise multiple paths to a destination. Prefix policies can only match prefixes. They cannot match route attributes, and they cannot change the attributes of routes.

Example: Advertising Multiple Paths in BGP

In this example, BGP routers are configured to advertise multiple paths instead of advertising only the active path. Advertising multiple paths in BGP is specified in Internet draft draft-ietf-idr-add-paths-04.txt, *Advertisement of Multiple Paths in BGP*.

- [Requirements on page 3380](#)
- [Overview on page 3380](#)
- [Configuration on page 3381](#)
- [Verification on page 3399](#)

Requirements

This example uses the following hardware and software components:

- Eight BGP-speaking devices.
- Five of the BGP-enabled devices do not necessarily need to be routers. For example, they can be EX Series Ethernet Switches.
- Three of the BGP-enabled devices are configured to send multiple paths or receive multiple paths (or both send and receive multiple paths). These three BGP-enabled devices must be M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- The three routers must be running Junos OS Release 11.4 or later.

Overview

The following statements are used for configuring multiple paths to a destination:

```
[edit protocols bgp group group-name family family],  
add-path {  
    receive;  
    send {
```



```

    path-count number;
    prefix-policy [ policy-names ];
  }
}

```

In this example, Router R5, Router R6, and Router R7 redistribute static routes into BGP. Router R1 and Router R4 are route reflectors. Router R2 and Router R3 are clients to Route Reflector R1. Router R8 is a client to Route Reflector R4.

Route reflection is optional when multiple-path advertisement is enabled in BGP.

With the **add-path send path-count 6** configuration, Router R1 is configured to send up to six paths (per destination) to Router R4.

With the **add-path receive** configuration, Router R4 is configured to receive multiple paths from Router R1.

With the **add-path send path-count 6** configuration, Router R4 is also configured to send up to six paths to Router R8.


With the **add-path receive** configuration, Router R8 is configured to receive multiple paths from Router R4.

The **add-path send prefix-policy allow_199** policy configuration (along with the corresponding route filter) limits Router R4 to sending multiple paths for only the 199.1.1.1/32 route.

Topology Diagram

Figure 96 on page 3381 shows the topology used in this example.

Figure 96: Advertisement of Multiple Paths in BGP



Configuration

- [Configuring Router R1 on page 3384](#)
- [Configuring Router R2 on page 3387](#)
- [Configuring Router R3 on page 3389](#)
- [Configuring Router R4 on page 3391](#)
- [Configuring Router R5 on page 3393](#)

- [Configuring Router R6 on page 3395](#)
- [Configuring Router R7 on page 3396](#)
- [Configuring Router R8 on page 3398](#)
- [Results on page 3398](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set interfaces fe-0/0/0 unit 12 family inet address 10.0.12.1/24
set interfaces fe-0/0/1 unit 13 family inet address 10.0.13.1/24
set interfaces fe-1/0/0 unit 14 family inet address 10.0.14.1/24
set interfaces fe-1/2/0 unit 15 family inet address 10.0.15.1/24
set interfaces lo0 unit 10 family inet address 10.0.0.10/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.10
set protocols bgp group rr cluster 10.0.0.10
set protocols bgp group rr neighbor 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.30
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.2 local-address 10.0.15.1
set protocols bgp group e1 neighbor 10.0.15.2 peer-as 2
set protocols bgp group rr_rr type internal
set protocols bgp group rr_rr local-address 10.0.0.10
set protocols bgp group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
  path-count 6
set protocols ospf area 0.0.0.0 interface lo0.10 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.12
set protocols ospf area 0.0.0.0 interface fe-0/0/1.13
set protocols ospf area 0.0.0.0 interface fe-1/0/0.14
set protocols ospf area 0.0.0.0 interface fe-1/2/0.15
set routing-options router-id 10.0.0.10
set routing-options autonomous-system 1
```

Router R2

```
set interfaces fe-1/2/0 unit 21 family inet address 10.0.12.2/24
set interfaces fe-1/2/1 unit 26 family inet address 10.0.26.1/24
set interfaces lo0 unit 20 family inet address 10.0.0.20/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.20 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.21
set protocols ospf area 0.0.0.0 interface fe-1/2/1.28
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1
```

Router R3

```
set interfaces fe-1/0/1 unit 31 family inet address 10.0.13.2/24
set interfaces fe-1/0/2 unit 37 family inet address 10.0.37.1/24
set interfaces lo0 unit 30 family inet address 10.0.0.30/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.30
```

```

set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.30 passive
set protocols ospf area 0.0.0.0 interface fe-1/0/1.31
set protocols ospf area 0.0.0.0 interface fe-1/0/2.37
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1

```

Router R4

```

set interfaces fe-1/2/0 unit 41 family inet address 10.0.14.2/24
set interfaces fe-1/2/1 unit 48 family inet address 10.0.48.1/24
set interfaces lo0 unit 40 family inet address 10.0.0.40/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.40
set protocols bgp group rr family inet unicast add-path receive
set protocols bgp group rr neighbor 10.0.0.10
set protocols bgp group rr_client type internal
set protocols bgp group rr_client local-address 10.0.0.40
set protocols bgp group rr_client cluster 10.0.0.40
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  path-count 6
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  prefix-policy allow_199
set protocols ospf area 0.0.0.0 interface fe-1/2/0.41
set protocols ospf area 0.0.0.0 interface lo0.40 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.48
set routing-options autonomous-system 1
set policy-options policy-statement allow_199 from route-filter 199.1.1.1/32 exact
set policy-options policy-statement allow_199 then accept

```

Router R5

```

set interfaces fe-1/2/0 unit 51 family inet address 10.0.15.2/24
set interfaces lo0 unit 50 family inet address 10.0.0.50/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.1 export s2b
set protocols bgp group e1 neighbor 10.0.15.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then as-path-expand 2
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject

```

Router R6

```

set interfaces fe-1/2/0 unit 62 family inet address 10.0.26.2/24
set interfaces lo0 unit 60 family inet address 10.0.0.60/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.1 export s2b
set protocols bgp group e1 neighbor 10.0.26.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject

```

Router R7

```
set interfaces fe-1/2/0 unit 73 family inet address 10.0.37.2/24
set interfaces lo0 unit 70 family inet address 10.0.0.70/32
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.1 export s2b
set protocols bgp group e1 neighbor 10.0.37.1 peer-as 1
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
```

Router R8

```
set interfaces fe-1/2/0 unit 84 family inet address 10.0.48.2/24
set interfaces lo0 unit 80 family inet address 10.0.0.80/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.80
set protocols bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
set protocols ospf area 0.0.0.0 interface lo0.80 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.84
set routing-options autonomous-system 1
```

Configuring Router R1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Configure the interfaces to Router R2, Router R3, Router R4, and Router R5, and configure the loopback (lo0) interface.

```
[edit interfaces]
user@R1# set fe-0/0/0 unit 12 family inet address 10.0.12.1/24

user@R1# set fe-0/0/1 unit 13 family inet address 10.0.13.1/24

user@R1# set fe-1/0/0 unit 14 family inet address 10.0.14.1/24

user@R1# set fe-1/2/0 unit 15 family inet address 10.0.15.1/24

user@R1# set lo0 unit 10 family inet address 10.0.0.10/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
user@R1# set group rr type internal
user@R1# set group rr local-address 10.0.0.10
user@R1# set group rr cluster 10.0.0.10
user@R1# set group rr neighbor 10.0.0.20
user@R1# set group rr neighbor 10.0.0.30

user@R1# set group rr_rr type internal
user@R1# set group rr_rr local-address 10.0.0.10
```

```

user@R1# set group e1 type external
user@R1# set group e1 neighbor 10.0.15.2 local-address 10.0.15.1
user@R1# set group e1 neighbor 10.0.15.2 peer-as 2

```

3. Configure Router R1 to send up to six paths to its neighbor, Router R4.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```

[edit protocols bgp]
user@R1# set group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
path-count 6

```

4. Configure OSPF on the interfaces.

```

[edit protocols ospf]
user@R1# set area 0.0.0.0 interface lo0.10 passive
user@R1# set area 0.0.0.0 interface fe-0/0/0.12
user@R1# set area 0.0.0.0 interface fe-0/0/1.13
user@R1# set area 0.0.0.0 interface fe-1/0/0.14
user@R1# set area 0.0.0.0 interface fe-1/2/0.15

```

5. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@R1# set router-id 10.0.0.10
user@R1# set autonomous-system 1

```

6. If you are done configuring the device, commit the configuration.

```

user@R1# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-0/0/0 {
  unit 12 {
    family inet {
      address 10.0.12.1/24;
    }
  }
}
fe-0/0/1 {
  unit 13 {
    family inet {
      address 10.0.13.1/24;
    }
  }
}
fe-1/0/0 {
  unit 14 {
    family inet {
      address 10.0.14.1/24;
    }
  }
}

```

```
}
fe-1/2/0 {
  unit 15 {
    family inet {
      address 10.0.15.1/24;
    }
  }
}
lo0 {
  unit 10 {
    family inet {
      address 10.0.0.10/32;
    }
  }
}

user@R1# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.10;
    cluster 10.0.0.10;
    neighbor 10.0.0.20;
    neighbor 10.0.0.30;
  }
  group e1 {
    type external;
    neighbor 10.0.15.2 {
      local-address 10.0.15.1;
      peer-as 2;
    }
  }
  group rr_rr {
    type internal;
    local-address 10.0.0.10;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            send {
              path-count 6;
            }
          }
        }
      }
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface lo0.10 {
      passive;
    }
    interface fe-0/0/0.12;
    interface fe-0/0/1.13;
    interface fe-1/0/0.14;
```

```

        interface fe-1/2/0.15;
    }
}

user@R1# show routing-options
router-id 10.0.0.10;
autonomous-system 1;

```

Configuring Router R2

Step-by-Step Procedure

To configure Router R2:

1. Configure the loopback (lo0) interface and the interfaces to Router R6 and Router R1.

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 21 family inet address 10.0.12.2/24

user@R2# set fe-1/2/1 unit 26 family inet address 10.0.26.1/24

user@R2# set lo0 unit 20 family inet address 10.0.0.20/32

```

2. Configure BGP and OSPF on Router R2's interfaces.

```

[edit protocols]
user@R2# set bgp group rr type internal
user@R2# set bgp group rr local-address 10.0.0.20

user@R2# set bgp group e1 type external
user@R2# set bgp group e1 neighbor 10.0.26.2 peer-as 2

user@R2# set ospf area 0.0.0.0 interface lo0.20 passive
user@R2# set ospf area 0.0.0.0 interface fe-1/2/0.21
user@R2# set ospf area 0.0.0.0 interface fe-1/2/1.28

```

3. For routes sent from Router R2 to Router R1, advertise Router R2 as the next hop, because Router R1 does not have a route to Router R6's address on the 10.0.26.0/24 network.

```

[edit]
user@R2# set policy-options policy-statement set_nh_self then next-hop self
user@R2# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self

```

4. Configure the autonomous system number.

```

[edit]
user@R2# set routing-options autonomous-system 1

```

5. If you are done configuring the device, commit the configuration.

```

user@R2# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 21 {
    family inet {
      address 10.0.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 26 {
    family inet {
      address 10.0.26.1/24;
    }
  }
}
lo0 {
  unit 20 {
    family inet {
      address 10.0.0.20/32;
    }
  }
}

user@R2# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R2# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.20;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.26.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.20 {
      passive;
    }
    interface fe-1/2/0.21;
    interface fe-1/2/1.28;
  }
}

user@R2# show routing-options
```



```
autonomous-system 1;
```

Configuring Router R3

Step-by-Step Procedure

To configure Router R3:

1. Configure the loopback (lo0) interface and the interfaces to Router R7 and Router R1.

```
[edit interfaces]
```

```
user@R3# set fe-1/0/1 unit 31 family inet address 10.0.13.2/24
```

```
user@R3# set fe-1/0/2 unit 37 family inet address 10.0.37.1/24
```

```
user@R3# set lo0 unit 30 family inet address 10.0.0.30/32
```

2. Configure BGP and OSPF on Router R3's interfaces.

```
[edit protocols]
```

```
user@R3# set bgp group rr type internal
```

```
user@R3# set bgp group rr local-address 10.0.0.30
```

```
user@R3# set bgp group e1 type external
```

```
user@R3# set bgp group e1 neighbor 10.0.37.2 peer-as 2
```

```
user@R3# set ospf area 0.0.0.0 interface lo0.30 passive
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/1.31
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/2.37
```

3. For routes sent from Router R3 to Router R1, advertise Router R3 as the next hop, because Router R1 does not have a route to Router R7's address on the 10.0.37.0/24 network.

```
[edit]
```

```
user@R3# set policy-options policy-statement set_nh_self then next-hop self
```

```
user@R3# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R3# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R3# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
```

```
fe-1/0/1 {
```

```
unit 31 {
```

```
family inet {
```

```
address 10.0.13.2/24;
```

```
    }
  }
}
fe-1/0/2 {
  unit 37 {
    family inet {
      address 10.0.37.1/24;
    }
  }
}
lo0 {
  unit 30 {
    family inet {
      address 10.0.0.30/32;
    }
  }
}

user@R3# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R3# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.30;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.37.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.30 {
      passive;
    }
    interface fe-1/0/1.31;
    interface fe-1/0/2.37;
  }
}

user@R3# show routing-options
autonomous-system 1;
```

Configuring Router R4**Step-by-Step Procedure**

To configure Router R4:

1. Configure the interfaces to Router R1 and Router R8, and configure the loopback (lo0) interface.

```
[edit interfaces]
```

```
user@R4# set fe-1/2/0 unit 41 family inet address 10.0.14.2/24
```

```
user@R4# set fe-1/2/1 unit 48 family inet address 10.0.48.1/24
```

```
user@R4# set lo0 unit 40 family inet address 10.0.0.40/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
```

```
user@R4# set group rr type internal
```

```
user@R4# set group rr local-address 10.0.0.40
```

```
user@R4# set group rr neighbor 10.0.0.10
```

```
user@R4# set group rr_client type internal
```

```
user@R4# set group rr_client local-address 10.0.0.40
```

```
user@R4# set group rr_client cluster 10.0.0.40
```

3. Configure Router R4 to send up to six paths to its neighbor, Router R8.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R4# set group rr_client neighbor 10.0.0.80 family inet unicast add-path send path-count 6
```

4. Configure Router R4 to receive multiple paths from its neighbor, Router R1.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R4# set group rr family inet unicast add-path receive
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf]
```

```
user@R4# set area 0.0.0.0 interface fe-1/2/0.41
```

```
user@R4# set area 0.0.0.0 interface lo0.40 passive
```

```
user@R4# set area 0.0.0.0 interface fe-1/2/1.48
```

6. Configure a policy that allows Router R4 to send Router R8 multiple paths to the 199.1.1.1/32 route.

Router R4 receives multiple paths for the 198.1.1.1/32 route and the 199.1.1.1/32 route. However, because of this policy, Router R4 only sends multiple paths for the 199.1.1.1/32 route.

```
[edit]
```

```
user@R4# set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast
add-path send prefix-policy allow_199
user@R4# set policy-options policy-statement allow_199 from route-filter 199.1.1.1/32
exact
user@R4# set policy-options policy-statement allow_199 then accept
```

7. Configure the autonomous system number.

```
[edit routing-options]
user@R4# set autonomous-system 1
```

8. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 41 {
    family inet {
      address 10.0.14.2/24;
    }
  }
}
fe-1/2/1 {
  unit 48 {
    family inet {
      address 10.0.48.1/24;
    }
  }
}
lo0 {
  unit 40 {
    family inet {
      address 10.0.0.40/32;
    }
  }
}

user@R4# show policy-options
policy-statement allow_199 {
  from {
    route-filter 199.1.1.1/32 exact;
  }
  then accept;
}

user@R4# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.40;
    family inet {
```

```

        unicast {
            add-path {
                receive;
            }
        }
    }
    neighbor 10.0.0.10;
}
group rr_client {
    type internal;
    local-address 10.0.0.40;
    cluster 10.0.0.40;
    neighbor 10.0.0.80 {
        family inet {
            unicast {
                add-path {
                    send {
                        path-count 6;
                        prefix-policy allow_199;
                    }
                }
            }
        }
    }
}
}
}
}
}
}
}
}
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.40 {
            passive;
        }
        interface fe-1/2/0.41;
        interface fe-1/2/1.48;
    }
}
}

user@R4# show routing-options
autonomous-system 1;

```

Configuring Router R5

Step-by-Step Procedure

To configure Router R5:

1. Configure the loopback (lo0) interface and the interface to Router R1.

[edit interfaces]

```
user@R5# set fe-1/2/0 unit 51 family inet address 10.0.15.2/24
```

```
user@R5# set lo0 unit 50 family inet address 10.0.0.50/32
```

2. Configure BGP on Router R5's interface.

[edit protocols]

```
user@R5# set bgp group e1 type external
```

```
user@R5# set bgp group e1 neighbor 10.0.15.1 peer-as 1
```

3. Create static routes for redistribution into BGP.

```
[edit]
user@R5# set routing-options static route 199.1.1.1/32 reject
user@R5# set routing-options static route 198.1.1.1/32 reject
```

4. Redistribute static and direct routes into BGP.

```
[edit]
user@R5# set protocols bgp group e1 neighbor 10.0.15.1 export s2b
user@R5# set policy-options policy-statement s2b from protocol static
user@R5# set policy-options policy-statement s2b from protocol direct
user@R5# set policy-options policy-statement s2b then as-path-expand 2
user@R5# set policy-options policy-statement s2b then accept
```

5. Configure the autonomous system number.

```
[edit]
user@R5# set routing-options autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R5# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R5# show interfaces
fe-1/2/0 {
  unit 51 {
    family inet {
      address 10.0.15.2/24;
    }
  }
}
lo0 {
  unit 50 {
    family inet {
      address 10.0.0.50/32;
    }
  }
}

user@R5# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then {
    as-path-expand 2;
    accept;
  }
}

user@R5# show protocols
bgp {
  group e1 {
```

```

    type external;
    neighbor 10.0.15.1 {
        export s2b;
        peer-as 1;
    }
}

user@R5# show routing-options
static {
    route 198.1.1.1/32 reject;
    route 199.1.1.1/32 reject;
}
autonomous-system 2;

```

Configuring Router R6

Step-by-Step Procedure

To configure Router R6:

1. Configure the loopback (lo0) interface and the interface to Router R2.


```

[edit interfaces]
user@R6# set fe-1/2/0 unit 62 family inet address 10.0.26.2/24

user@R6# set lo0 unit 60 family inet address 10.0.0.60/32

```
2. Configure BGP on Router R6's interface.


```

[edit protocols]
user@R6# set bgp group e1 type external
user@R6# set bgp group e1 neighbor 10.0.26.1 peer-as 1

```
3. Create static routes for redistribution into BGP.


```

[edit]
user@R6# set routing-options static route 199.1.1.1/32 reject
user@R6# set routing-options static route 198.1.1.1/32 reject

```
4. Redistribute static and direct routes from Router R6's routing table into BGP.


```

[edit]
user@R6# set protocols bgp group e1 neighbor 10.0.26.1 export s2b
user@R6# set policy-options policy-statement s2b from protocol static
user@R6# set policy-options policy-statement s2b from protocol direct
user@R6# set policy-options policy-statement s2b then accept

```
5. Configure the autonomous system number.


```

[edit]
user@R6# set routing-options autonomous-system 2

```
6. If you are done configuring the device, commit the configuration.


```

user@R6# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R6# show interfaces
fe-1/2/0 {
  unit 62 {
    family inet {
      address 10.0.26.2/24;
    }
  }
}
lo0 {
  unit 60 {
    family inet {
      address 10.0.0.60/32;
    }
  }
}

user@R6# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R6# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.26.1 {
      export s2b;
      peer-as 1;
    }
  }
}

user@R6# show routing-options
static {
  route 198.1.1.1/32 reject;
  route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

Configuring Router R7

Step-by-Step Procedure

To configure Router R7:

1. Configure the loopback (**lo0**) interface and the interface to Router R3.

[edit interfaces]

```
user@R7# set fe-1/2/0 unit 73 family inet address 10.0.37.2/24
```

```
user@R7# set lo0 unit 70 family inet address 10.0.0.70/32
```

2. Configure BGP on Router R7's interface.

[edit protocols]

```
user@R7# set bgp group e1 type external
```



```
user@R7# set bgp group e1 neighbor 10.0.37.1 peer-as 1
```

3. Create a static route for redistribution into BGP.

```
[edit]
user@R7# set routing-options static route 199.1.1.1/32 reject
```

4. Redistribute static and direct routes from Router R7's routing table into BGP.

```
[edit]
user@R7# set protocols bgp group e1 neighbor 10.0.37.1 export s2b
user@R7# set policy-options policy-statement s2b from protocol static
user@R7# set policy-options policy-statement s2b from protocol direct
user@R7# set policy-options policy-statement s2b then accept
```

5. Configure the autonomous system number.

```
[edit]
user@R7# set routing-options autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R7# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R7# show interfaces
fe-1/2/0 {
  unit 73 {
    family inet {
      address 10.0.37.2/24;
    }
  }
}
lo0 {
  unit 70 {
    family inet {
      address 10.0.0.70/32;
    }
  }
}

user@R7# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R7# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.37.1 {
      export s2b;
      peer-as 1;
    }
  }
}
```

```
    }  
  }  
}  
  
user@R7# show routing-options  
static {  
    route 199.1.1.1/32 reject;  
}  
autonomous-system 2;
```

Configuring Router R8

Step-by-Step Procedure

To configure Router R8:

1. Configure the loopback (lo0) interface and the interface to Router R4.

[edit interfaces]
user@R8# set fe-1/2/0 unit 84 family inet address 10.0.48.2/24

user@R8# set lo0 unit 80 family inet address 10.0.0.80/32
2. Configure BGP and OSPF on Router R8's interface.

[edit protocols]
user@R8# set bgp group rr type internal
user@R8# set bgp group rr local-address 10.0.0.80

user@R8# set ospf area 0.0.0.0 interface lo0.80 passive
user@R8# set ospf area 0.0.0.0 interface fe-1/2/0.84
3. Configure Router R8 to receive multiple paths from its neighbor, Router R4.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

[edit protocols]
user@R8# set bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
4. Configure the autonomous system number.

[edit]
user@R8# set routing-options autonomous-system 1
5. If you are done configuring the device, commit the configuration.

user@R8# commit

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R8# show interfaces  
fe-1/2/0 {  
    unit 84 {  
        family inet {  
            address 10.0.48.2/24;
```

```

    }
  }
}
lo0 {
  unit 80 {
    family inet {
      address 10.0.0.80/32;
    }
  }
}

user@R8# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.80;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            receive;
          }
        }
      }
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface lo0.80 {
      passive;
    }
    interface fe-1/2/0.84;
  }
}

user@R8# show routing-options
autonomous-system 1;

```

Verification

Confirm that the configuration is working properly.

- [Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths on page 3400](#)
- [Verifying That Router R1 Is Advertising Multiple Paths on page 3400](#)
- [Verifying That Router R4 Is Receiving and Advertising Multiple Paths on page 3401](#)
- [Verifying That Router R8 Is Receiving Multiple Paths on page 3402](#)
- [Checking the Path ID on page 3402](#)

Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths

Purpose Make sure that one or both of the following strings appear in the output of the **show bgp neighbor** command:

- NLRI's for which peer can receive multiple paths: inet-unicast
- NLRI's for which peer can send multiple paths: inet-unicast

Action

```
user@R1> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.10+65237 AS 1
  Type: Internal  State: Established  Flags: <Sync>
... NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.10
Peer: 10.0.0.10+65237 AS 1    Local: 10.0.0.40+179 AS 1
  Type: Internal  State: Established  Flags: <Sync>
...
  NLRI's for which peer can send multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.80
Peer: 10.0.0.80+55416 AS 1    Local: 10.0.0.40+179 AS 1
  Type: Internal  State: Established (route reflector client)Flags: <Sync>
'''
  NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R8> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.80+55416 AS 1
  Type: Internal  State: Established  Flags: <Sync>
...
  NLRI's for which peer can send multiple paths: inet-unicast
...
```

Verifying That Router R1 Is Advertising Multiple Paths

Purpose Make sure that multiple paths to the 198.1.1.1/32 destination and multiple paths to the 199.1.1.1/32 destination are advertised to Router R4.

Action user@R1> show route advertising-protocol bgp 10.0.0.40
 inet.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
	10.0.15.2		100	2 2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

Meaning When you see one prefix and more than one next hop, it means that multiple paths are advertised to Router R4.

Verifying That Router R4 Is Receiving and Advertising Multiple Paths

Purpose Make sure that multiple paths to the 199.1.1.1/32 destination are received from Router R1 and advertised to Router R8. Make sure that multiple paths to the 198.1.1.1/32 destination are received from Router R1, but only one path to this destination is advertised to Router R8.

Action user@R4> show route receive-protocol bgp 10.0.0.10
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
	10.0.15.2		100	2 2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

user@R4> show route advertising-protocol bgp 10.0.0.80
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

Meaning The **show route receive-protocol** command shows that Router R4 receives two paths to the 198.1.1.1/32 destination and three paths to the 199.1.1.1/32 destination. The **show route advertising-protocol** command shows that Router R4 advertises only one path to the 198.1.1.1/32 destination and advertises all three paths to the 199.1.1.1/32 destination.

Because of the prefix policy that is applied to Router R4, Router R4 does not advertise multiple paths to the 198.1.1.1/32 destination. Router R4 advertises only one path to the 198.1.1.1/32 destination even though it receives multiple paths to this destination.

Verifying That Router R8 Is Receiving Multiple Paths

Purpose Make sure that Router R8 receives multiple paths to the 199.1.1.1/32 destination through Router R4. Make sure that Router R8 receives only one path to the 198.1.1.1/32 destination through Router R4.

Action user@R8> show route receive-protocol bgp 10.0.0.40
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)

Prefix	Nexthop	MED	Lclpref	AS path
* 10.0.0.50/32	10.0.15.2		100	2 2 I
* 10.0.0.60/32	10.0.0.20		100	2 I
* 10.0.0.70/32	10.0.0.30		100	2 I
* 198.1.1.1/32	10.0.0.20		100	2 I
* 199.1.1.1/32	10.0.0.20		100	2 I
	10.0.0.30		100	2 I
	10.0.15.2		100	2 2 I
* 200.1.1.0/30	10.0.0.20		100	2 I

Checking the Path ID

Purpose On the downstream devices, Router R4 and Router R8, verify that a path ID uniquely identifies the path. Look for the **Addpath Path ID:** string.

Action user@R4> show route 199.1.1.1/32 detail

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 3 announced)
  *BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 9
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.0.20
    Indirect next hop: 92041c8 262146
    State: <Active Int Ext>
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (3): 2-KRT 3-BGP RT Background 4-Resolve tree

  1
    AS path: 2 I (Originator) Cluster list: 10.0.0.10
    AS path: Originator ID: 10.0.0.20
    Accepted
    Localpref: 100
    Router ID: 10.0.0.10
    Addpath Path ID: 1
  BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 4
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.0.30
    Indirect next hop: 92042ac 262151
    State: <NotBest Int Ext>
    Inactive reason: Not Best in its group - Router ID
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (1): 3-BGP RT Background
    AS path: 2 I (Originator) Cluster list: 10.0.0.10
    AS path: Originator ID: 10.0.0.30
    Accepted
    Localpref: 100
    Router ID: 10.0.0.10
    Addpath Path ID: 2
  BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 4
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.15.2
    Indirect next hop: 92040e4 262150
    State: <Int Ext>
    Inactive reason: AS path
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (1): 3-BGP RT Background
    AS path: 2 2 I
    Accepted
```

```

Localpref: 100
Router ID: 10.0.0.10
Addpath Path ID: 3

```

```
user@R8> show route 199.1.1.1/32 detail
```

```

inet.0: 17 destinations, 19 routes (17 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 1 announced)
*BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 9
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.20
      Indirect next hop: 91fc0e4 262148
      State: <Active Int Ext>
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      Announcement bits (2): 2-KRT 4-Resolve tree 1
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.20
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 1
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.30
      Indirect next hop: 91fc1c8 262152
      State: <NotBest Int Ext>
      Inactive reason: Not Best in its group - Router ID
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.30
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 2
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.15.2
      Indirect next hop: 91fc2ac 262153
      State: <Int Ext>
      Inactive reason: AS path
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 2 I (Originator) Cluster list: 10.0.0.40
      AS path: Originator ID: 10.0.0.10

```



```
Accepted
Localpref: 100
Router ID: 10.0.0.40
Addpath Path ID: 3
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3149](#)
 - [BGP Configuration Overview](#)

IBGP Scaling Configuration

- [Example: Configuring BGP Route Reflectors on page 3405](#)
- [Example: Configuring BGP Confederations on page 3422](#)

Example: Configuring BGP Route Reflectors

- [Understanding BGP Route Reflectors on page 3405](#)
- [Example: Configuring a Route Reflector on page 3407](#)

Understanding BGP Route Reflectors

Because of the internal BGP (IBGP) full-mesh requirement, most networks use route reflectors to simplify configuration. The formula to compute the number of sessions required for a full mesh is $v * (v - 1) / 2$, where v is the number of BGP-enabled devices. The full-mesh model does not scale well. Using a route reflector, you group routers into clusters, which are identified by numeric identifiers unique to the autonomous system (AS). Within the cluster, you must configure a BGP session from a single router (the route reflector) to each internal peer. With this configuration, the IBGP full-mesh requirement is met.

To use route reflection in an AS, you designate one or more routers as a route reflector—typically, one per point of presence (POP). Route reflectors have the special BGP ability to readvertise routes learned from an internal peer to other internal peers. So rather than requiring all internal peers to be fully meshed with each other, route reflection requires only that the route reflector be fully meshed with all internal peers. The route reflector and all of its internal peers form a cluster, as shown in [Figure 97 on page 3406](#).



NOTE: For some Juniper Networks devices, you must have an Advanced BGP Feature license installed on each device that uses a route reflector. For license details, see the *Junos OS Initial Configuration Guide for Security Devices*.


Figure 97: Simple Route Reflector Topology (One Cluster)

Figure 97 on page 3406 shows Router RR configured as the route reflector for Cluster 127. The other routers are designated internal peers within the cluster. BGP routes are advertised to Router RR by any of the internal peers. RR then readvertises those routes to all other peers within the cluster.

You can configure multiple clusters and link them by configuring a full mesh of route reflectors (see Figure 98 on page 3406).


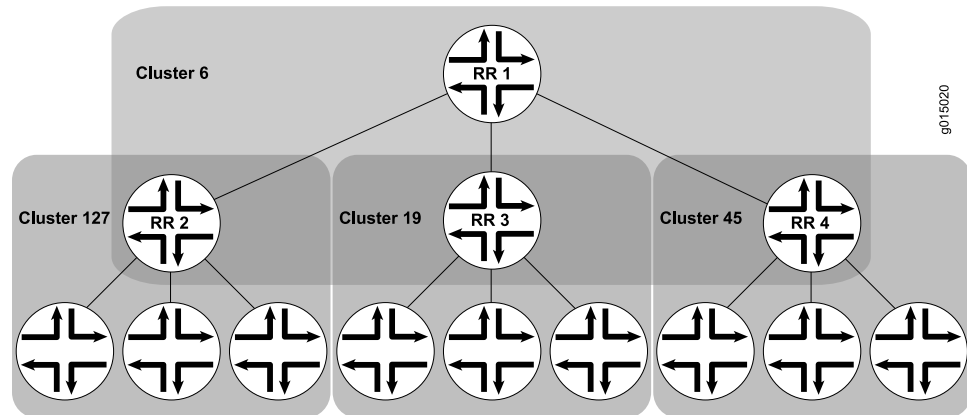
Figure 98: Basic Route Reflection (Multiple Clusters)

Figure 98 on page 3406 shows Route Reflectors RR 1, RR 2, RR 3, and RR 4 as fully meshed internal peers. When a router advertises a route to RR 1, RR 1 readvertises the route to the other route reflectors, which, in turn, readvertise the route to the remaining routers within the AS. Route reflection allows the route to be propagated throughout the AS without the scaling problems created by the full mesh requirement.

However, as clusters become large, a full mesh with a route reflector becomes difficult to scale, as does a full mesh between route reflectors. To help offset this problem, you can group clusters of routers together into clusters of clusters for hierarchical route reflection (see [Figure 99 on page 3407](#)).

Figure 99: Hierarchical Route Reflection (Clusters of Clusters)



[Figure 99 on page 3407](#) shows RR 2, RR 3, and RR 4 as the route reflectors for Clusters 127, 19, and 45, respectively. Rather than fully mesh those route reflectors, the network administrator has configured them as part of another cluster (Cluster 6) for which RR 1 is the route reflector. When a router advertises a route to RR 2, RR 2 readvertises the route to all the routers within its own cluster, and then readvertises the route to RR 1. RR 1 readvertises the route to the routers in its cluster, and those routers propagate the route down through their clusters.

Example: Configuring a Route Reflector

This example shows how to configure a route reflector.

- [Requirements on page 3407](#)
- [Overview on page 3407](#)
- [Configuration on page 3409](#)
- [Verification on page 3417](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Generally, internal BGP (IBGP)-enabled devices need to be fully meshed, because IBGP does not readvertise updates to other IBGP-enabled devices. The full mesh is a logical mesh achieved through configuration of multiple **neighbor** statements on each IBGP-enabled device. The full mesh is not necessarily a physical full mesh. Maintaining a full mesh (logical or physical) does not scale well in large deployments.

Figure 100 on page 3409 shows an IBGP network with Device A acting as a route reflector. Device B and Device C are clients of the route reflector. Device D and Device E are outside the cluster, so they are nonclients of the route reflector.

On Device A (the route reflector), you must form peer relationships with all of the IBGP-enabled devices by including the **neighbor** statement for the clients (Device B and Device C) and the nonclients (Device D and Device E). You must also include the **cluster** statement and a cluster identifier. The cluster identifier can be any 32-bit value. This example uses the loopback interface IP address of the route reflector.


On Device B and Device C, the route reflector clients, you only need one **neighbor** statement that forms a peer relationship with the route reflector, Device A.

On Device D and Device E, the nonclients, you need a **neighbor** statement for each nonclient device (D-to-E and E-to-D). You also need a **neighbor** statement for the route reflector (D-to-A and E-to-A). Device D and Device E do not need **neighbor** statements for the client devices (Device B and Device C).



TIP: Device D and Device E are considered to be nonclients because they have explicitly configured peer relationships with each other. To make them RRroute reflector clients, remove the **neighbor 192.168.5.5** statement from the configuration on Device D, and remove the **neighbor 192.168.0.1** statement from the configuration on Device E.

Figure 100: IBGP Network Using a Route Reflector

**Configuration**

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```

set interfaces fe-0/0/0 unit 1 description to-B
set interfaces fe-0/0/0 unit 1 family inet address 10.10.10.1/30
set interfaces fe-0/0/1 unit 3 description to-D
set interfaces fe-0/0/1 unit 3 family inet address 10.10.10.9/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers cluster 192.168.6.5
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.1
set protocols ospf area 0.0.0.0 interface fe-0/0/1.3
set policy-options policy-statement send-ospf term 2 from protocol ospf

```

```
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17
```

Device B

```
set interfaces fe-0/0/0 unit 2 description to-A
set interfaces fe-0/0/0 unit 2 family inet address 10.10.10.2/30
set interfaces fe-0/0/1 unit 5 description to-C
set interfaces fe-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.2
set protocols ospf area 0.0.0.0 interface fe-0/0/1.5
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

Device C

```
set interfaces fe-0/0/0 unit 6 description to-B
set interfaces fe-0/0/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.6
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

Device D

```
set interfaces fe-0/0/0 unit 4 description to-A
set interfaces fe-0/0/0 unit 4 family inet address 10.10.10.10/30
set interfaces fe-0/0/1 unit 7 description to-E
set interfaces fe-0/0/1 unit 7 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.0.1/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.0.1
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.4
set protocols ospf area 0.0.0.0 interface fe-0/0/1.7
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 17
```

Device E

```
set interfaces fe-0/0/0 unit 8 description to-D
set interfaces fe-0/0/0 unit 8 family inet address 10.10.10.14/30
```

```

set interfaces lo0 unit 5 family inet address 192.168.5.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.5.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.8
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.5.5
set routing-options autonomous-system 17

```

Configuring the Route Reflector

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IBGP in the network using Juniper Networks Device A as a route reflector:

1. Configure the interfaces.

```

[edit interfaces]
user@A# set fe-0/0/0 unit 1 description to-B
user@A# set fe-0/0/0 unit 1 family inet address 10.10.1/30
user@A# set fe-0/0/1 unit 3 description to-D
user@A# set fe-0/0/1 unit 3 family inet address 10.10.9/30
user@A# set lo0 unit 1 family inet address 192.168.6.5/32

```
2. Configure BGP, including the cluster identifier and neighbor relationships with all IBGP-enabled devices in the autonomous system (AS).

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set local-address 192.168.6.5
user@A# set export send-ospf
user@A# set cluster 192.168.6.5
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
user@A# set neighbor 192.168.0.1
user@A# set neighbor 192.168.5.5

```

3. Configure static routing or an interior gateway protocol (IGP).

This example uses OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@A# set interface lo0.1 passive
user@A# set interface fe-0/0/0.1
user@A# set interface fe-0/0/1.3

```

4. Configure the policy that redistributes OSPF routes into BGP.

```

[edit policy-options policy-statement send-ospf term 2]
user@A# set from protocol ospf
user@A# set then accept

```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show interfaces
fe-0/0/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
fe-0/0/1 {
  unit 3 {
    description to-D;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.6.5;
    export send-ospf;
    cluster 192.168.6.5;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
    neighbor 192.168.0.1;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-0/0/0.1;
    interface fe-0/0/1.3;
```



```

    }
  }

user@A# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each nonclient BGP peer within the cluster that you are configuring, if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Configuring Client Peers

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure client peers:

1. Configure the interfaces.


```

[edit interfaces]
user@B# set fe-0/0/0 unit 2 description to-A
user@B# set fe-0/0/0 unit 2 family inet address 10.10.10.2/30
user@B# set fe-0/0/1 unit 5 description to-C
user@B# set fe-0/0/1 unit 5 family inet address 10.10.10.5/30
user@B# set lo0 unit 2 family inet address 192.163.6.4/32
      
```

2. Configure the BGP neighbor relationship with the route reflector.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@B# set type internal
user@B# set local-address 192.163.6.4
user@B# set export send-ospf
user@B# set neighbor 192.168.6.5

```

3. Configure OSPF.


```

[edit protocols ospf area 0.0.0.0]
user@B# set interface lo0.2 passive
user@B# set interface fe-0/0/0.2
user@B# set interface fe-0/0/1.5
      
```
4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@B# set from protocol ospf
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/0/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
fe-0/0/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}

user@B# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.163.6.4;
    export send-ospf;
    neighbor 192.168.6.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-0/0/0.2;
    interface fe-0/0/1.5;
```

```

    }
  }

user@B# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each client BGP peer within the cluster that you are configuring if the other client devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Configuring Nonclient Peers

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure nonclient peers:

1. Configure the interfaces.


```

[edit interfaces]
user@D# set fe-0/0/0 unit 4 description to-A
user@D# set fe-0/0/0 unit 4 family inet address 10.10.10.10/30
user@D# set fe-0/0/1 unit 7 description to-E
user@D# set fe-0/0/1 unit 7 family inet address 10.10.10.13/30
user@D# set lo0 unit 4 family inet address 192.168.0.1/32
      
```
2. Configure the BGP neighbor relationships with the RRroute reflector and with the other nonclient peers.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@D# set type internal
user@D# set local-address 192.168.0.1
user@D# set export send-ospf
user@D# set neighbor 192.168.6.5
user@D# set neighbor 192.168.5.5

```

3. Configure OSPF.


```

[edit protocols ospf area 0.0.0.0]
user@D# set interface lo0.4 passive
user@D# set interface fe-0/0/0.4
      
```

```
user@D# set interface fe-0/0/1.7
```

4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@D# set from protocol ospf
user@D# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@D# set router-id 192.168.0.1
user@D# set autonomous-system 17
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
fe-0/0/0 {
  unit 4 {
    description to-A;
    family inet {
      address 10.10.10.10/30;
    }
  }
}
fe-0/0/1 {
  unit 7 {
    description to-E;
    family inet {
      address 10.10.10.13/30;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}
```

```
user@D# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.0.1;
    export send-ospf;
    neighbor 192.168.6.5;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.4 {
```

```

        passive;
    }
    interface fe-0/0/0.4;
    interface fe-0/0/1.7;
}
}

user@D# show policy-options
policy-statement send-ospf {
    term 2 {
        from protocol ospf;
        then accept;
    }
}

user@D# show routing-options
router-id 192.168.0.1;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat these steps for each nonclient BGP peer within the cluster that you are configuring if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3417](#)
- [Verifying BGP Groups on page 3420](#)
- [Verifying BGP Summary Information on page 3420](#)
- [Verifying Routing Table Information on page 3420](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is established for each neighbor address.

Action From operational mode, enter the **show bgp neighbor** command.

```

user@A> show bgp neighbor
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+62857 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  BFD: disabled, down

```

```

NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5    Sent 3    Checked 19
Input messages: Total 2961    Updates 7    Refreshes 0    Octets 56480
Output messages: Total 2945    Updates 6    Refreshes 0    Octets 56235
Output Queue[0]: 0

Peer: 192.168.0.1+179 AS 17    Local: 192.168.6.5+60068 AS 17
Type: Internal    State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ send-ospf ]
Options: <Preference LocalAddress Cluster Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.0.1    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 3
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 18    Sent 20    Checked 12
Input messages: Total 15    Updates 5    Refreshes 0    Octets 447
Output messages: Total 554    Updates 4    Refreshes 0    Octets 32307

```

Output Queue[0]: 0

```

Peer: 192.168.5.5+57458 AS 17 Local: 192.168.6.5+179 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.5.5      Local ID: 192.168.6.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 2
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 17)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        7
    Accepted prefixes:        7
    Suppressed due to damping: 0
    Advertised prefixes:      6
  Last traffic (seconds): Received 17 Sent 3 Checked 9
  Input messages: Total 2967 Updates 7 Refreshes 0 Octets 56629
  Output messages: Total 2943 Updates 6 Refreshes 0 Octets 56197
  Output Queue[0]: 0

```

```

Peer: 192.168.40.4+53990 AS 17 Local: 192.168.6.5+179 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.40.4     Local ID: 192.168.6.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 1
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast

```

```

Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        7
  Accepted prefixes:        7
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5   Sent 23   Checked 52
Input messages: Total 2960   Updates 7   Refreshes 0   Octets 56496
Output messages: Total 2943   Updates 6   Refreshes 0   Octets 56197
Output Queue[0]: 0

```

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal   AS: 17                      Local AS: 17
Name: internal-peers  Index: 0                     Flags: <>
Export: [ send-ospf ]
Options: <Cluster>
Holdtime: 0
Total peers: 4         Established: 4
192.163.6.4+179
192.168.40.4+53990
192.168.0.1+179
192.168.5.5+57458
inet.0: 0/26/16/0

Groups: 1 Peers: 4 External: 0 Internal: 4 Down peers: 0 Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State Pending
inet.0      26         0         0         0         0      0      0

```

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary

Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State Pending
inet.0      26         0         0         0         0      0      0
Peer      AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4 17      2981     2965      0        0    22:19:15 0/6/1/0      0/0/0/0
192.168.0.1 17       36       575      0         0    13:43 0/6/1/0      0/0/0/0
192.168.5.5 17      2988     2964      0         0    22:19:10 0/7/7/0      0/0/0/0
192.168.40.4 17      2980     2964      0         0    22:19:14 0/7/7/0      0/0/0/0

```

Verifying Routing Table Information

Purpose Verify that the routing table contains the IBGP routes.

Action From operational mode, enter the **show route** command.

```

user@A> show route
inet.0: 12 destinations, 38 routes (12 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      * [Direct/0] 22:22:03
                  > via fe-0/0/0.1
                  [BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
10.10.10.1/32     * [Local/0] 22:22:03
                  Local via fe-0/0/0.1
10.10.10.4/30     * [OSPF/10] 22:21:13, metric 2
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
10.10.10.8/30     * [Direct/0] 22:22:03
                  > via fe-0/0/1.3
                  [BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
10.10.10.9/32     * [Local/0] 22:22:03
                  Local via fe-0/0/1.3
10.10.10.12/30    * [OSPF/10] 22:21:08, metric 2
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.163.6.4/32    * [OSPF/10] 22:21:13, metric 1
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:55, MED 1, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
                  [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
192.168.0.1/32    * [OSPF/10] 22:21:08, metric 1
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:51, MED 1, localpref 100, from 192.168.5.5
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.168.5.5/32    * [OSPF/10] 22:21:08, metric 2
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 00:15:24, MED 1, localpref 100, from 192.168.0.1
                  AS path: I
                  > to 10.10.10.10 via fe-0/0/1.3
                  [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via fe-0/0/0.1
192.168.6.5/32    * [Direct/0] 22:22:04

```

```
> via lo0.1
[BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
AS path: I
> to 10.10.10.10 via fe-0/0/1.3
[BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via fe-0/0/0.1
192.168.40.4/32 * [OSPF/10] 22:21:13, metric 2
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:55, MED 1, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
AS path: I
> to 10.10.10.10 via fe-0/0/1.3
224.0.0.5/32 * [OSPF/10] 22:22:07, metric 1
MultiRecv
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3149](#)
 - [BGP Configuration Overview](#)

Example: Configuring BGP Confederations

- [Understanding BGP Confederations on page 3422](#)
- [Example: Configuring BGP Confederations on page 3423](#)

Understanding BGP Confederations

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large autonomous system (AS) into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64,512 and 65,535.

Within a sub-AS, the same internal BGP (IBGP) full mesh requirement exists. Connections to other confederations are made with standard external BGP (EBGP), and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS.

[Figure 101 on page 3423](#) shows an AS divided into four confederations.

Figure 101: BGP Confederations




Figure 101 on page 3423 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

Example: Configuring BGP Confederations

This example shows how to configure BGP confederations.

- [Requirements on page 3423](#)
- [Overview on page 3423](#)
- [Configuration on page 3424](#)
- [Verification on page 3426](#)

Requirements

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3150](#).
- Configure interior gateway protocol (IGP) sessions between peers.
- Configure a routing policy to advertise the BGP routes.

Overview


Within a BGP confederation, the links between the confederation member autonomous systems (ASs) must be external BGP (EBGP) links, not internal BGP (IBGP) links.

Similar to route reflectors, BGP confederations reduce the number of peer sessions and TCP sessions to maintain connections between IBGP routing devices. BGP confederation is one method used to solve the scaling problems created by the IBGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous

systems. Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535. Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

Figure 102 on page 3424 shows a sample network in which AS 17 has two separate confederations: sub-AS 64512 and sub-AS 64513, each of which has multiple routers. Within a sub-AS, an IGP is used to establish network connectivity with internal peers. Between sub-ASs, an EBGP peer session is established.

Figure 102: Typical Network Using BGP Confederations



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

All Devices in Sub-AS 64512	<pre> set routing-options autonomous-system 64512 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64512 type internal set protocols bgp group sub-AS-64512 local-address 192.168.5.1 set protocols bgp group sub-AS-64512 neighbor 192.168.8.1 set protocols bgp group sub-AS-64512 neighbor 192.168.15.1 </pre>
Border Device in Sub-AS 64512	<pre> set protocols bgp group to-sub-AS-64513 type external set protocols bgp group to-sub-AS-64513 peer-as 64513 set protocols bgp group to-sub-AS-64513 neighbor 192.168.5.2 </pre>
All Devices in Sub-AS 64513	<pre> set routing-options autonomous-system 64513 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64513 type internal set protocols bgp group sub-AS-64513 local-address 192.168.5.2 set protocols bgp group sub-AS-64513 neighbor 192.168.9.1 </pre>

	<pre>set protocols bgp group sub-AS-64513 neighbor 192.168.16.1</pre>
Border Device in Sub-AS 64513	<pre>set protocols bgp group to-sub-AS-64512 type external set protocols bgp group to-sub-AS-64512 peer-as 64512 set protocols bgp group to-sub-AS-64512 neighbor 192.168.5.1</pre>
Step-by-Step Procedure	<p>This procedure shows the steps for the devices that are in sub-AS 64512.</p> <p>The autonomous-system statement sets the sub-AS number of the device.</p> <p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure BGP confederations:</p> <ol style="list-style-type: none"> Set the sub-AS number for the device. <pre>[edit routing-options] user@host# set autonomous-system 64512</pre> In the confederation, include all sub-ASs in the main AS. <p>The number 17 represents the main AS. The members statement lists all the sub-ASs in the main AS.</p> <pre>[edit routing-options confederation] user@host# set 17 members 64512 user@host# set 17 members 64513</pre> On the border device in sub-AS 64512, configure an EBGP connection to the border device in AS 64513. <pre>[edit protocols bgp group to-sub-AS-64513] user@host# set type external user@host# set neighbor 192.168.5.2 user@host# set peer-as 64513</pre> Configure an IBGP group for peering with the devices within sub-AS 64512. <pre>[edit protocols bgp group sub-AS-64512] user@host# set type internal user@host# set local-address 192.168.5.1 user@host# neighbor 192.168.8.1 user@host# neighbor 192.168.15.1</pre>
Results	<p>From configuration mode, confirm your configuration by entering the show routing-options and show protocols commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.</p> <pre>user@host# show routing-options autonomous-system 64512; confederation 17 members [64512 64513]; user@host# show protocols bgp { group to-sub-AS-64513 { # On the border devices only type external;</pre>

```
    peer-as 64513;
    neighbor 192.168.5.2;
  }
  group sub-AS-64512 {
    type internal;
    local-address 192.168.5.1;
    neighbor 192.168.8.1;
    neighbor 192.168.15.1;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.
Repeat these steps for sSub-AS 64513.

Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3426](#)
- [Verifying BGP Groups on page 3427](#)
- [Verifying BGP Summary Information on page 3428](#)

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the CLI, enter the **show bgp neighbor** command.

Sample Output

```
user@host> show bgp neighbor
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: Sync
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh

  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
```

```

Active prefixes: 4
Received prefixes: 6
Suppressed due to damping: 0
Table inet6.0 Bit: 20000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

Meaning The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For **State**, each BGP session is **Established**.
- For **Type**, each peer is configured as the correct type (either internal or external).
- For **AS**, the AS number of the BGP neighbor is correct.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the CLI, enter the **show bgp group** command.

Sample Output

```

user@host> show bgp group
Group Type: Internal  AS: 10045      Local AS: 10045
Name: pe-to-asbr2
Export: [ match-all ]
Total peers: 1      Established: 1
10.0.0.4+179
bgp.l3vpn.0: 1/1/0
vpn-green.inet.0: 1/1/0

Groups: 1  Peers: 1  External: 0  Internal: 1  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
bgp.l3vpn.0      1          1          0          0          0          0

```

Meaning The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For **AS**, each group's remote AS is configured correctly.
- For **Local AS**, each group's local AS is configured correctly.
- For **Group Type**, each group has the correct type (either internal or external).

- For **Total peers**, the expected number of peers within the group is shown.
- For **Established**, the expected number of peers within the group have BGP sessions in the **Established** state.
- The IP addresses of all the peers within the group are present.

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the CLI, enter the **show bgp summary** command.

Sample Output

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table      Tot Paths  Act Paths  Suppressed    History Damp State    Pending
inet.0          6          4          0          0      0      0      0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2    65002    88675    88652      0        2      42:38 2/4/0
           0/0/0
10.0.0.3    65002    54528    54532      0        1     2w4d22h 0/0/0
           0/0/0
10.0.0.4    65002    51597    51584      0        0     2w3d22h 2/2/0
           0/0/0
```

Meaning The output shows a summary of BGP session information. Verify the following information:

- For **Groups**, the total number of configured groups is shown.
- For **Peers**, the total number of BGP peers is shown.
- For **Down Peers**, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under **Peer**, the IP address for each configured peer is shown.
- Under **AS**, the peer AS for each configured peer is correct.
- Under **Up/Dwn State**, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

BGP Security Configuration

- [Example: Configuring BGP Route Authentication on page 3429](#)
- [Examples: Configuring TCP and BGP Security on page 3435](#)

Example: Configuring BGP Route Authentication

- [Understanding Route Authentication on page 3429](#)
- [Example: Configuring Route Authentication for BGP on page 3430](#)

Understanding Route Authentication

The use of router and route authentication and route integrity greatly mitigates the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. In this kind of attack, the attacked router can be tricked into creating a routing loop, or the attacked router's routing table can be greatly increased thus impacting performance, or routing information can be redirected to a place in the network for the attacker to analyze it. Bogus route advertisements can be sent out on a segment. These updates can be accepted into the routing tables of neighbor routers unless an authentication mechanism is in place to verify the source of the routes.

Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, it accepts the route. By using a hashing algorithm, the key is not sent over the wire in plain text. Instead, a hash is calculated using the configured key. The routing update is used as the input text, along with the key, into the hashing function. This hash is sent along with the route update to the receiving router. The receiving router compares the received hash with a hash it generates on the route update using the preshared key configured on it. If the two hashes are the same, the route is assumed to be from a trusted source. The key is known only to the sending and receiving routers.

To further strengthen security, you can configure a series of authentication keys (a *keychain*). Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as *hitless* because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.

The sending peer uses the following rules to identify the active authentication key:

- The start time is less than or equal to the current time (in other words, not in the future).
- The start time is greater than that of all other keys in the chain whose start time is less than the current time (in other words, closest to the current time).

The receiving peer determines the key with which it authenticates based on the incoming key identifier.

The sending peer identifies the current authentication key based on a configured start time and then generates a hash value using the current key. The sending peer then inserts a TCP-enhanced authentication option object into the BGP update message. The object contains an object ID (assigned by IANA), the object length, the current key, and a hash value.

The receiving peer examines the incoming TCP-enhanced authentication option, looks up the received authentication key, and determines whether the key is acceptable based on the start time, the system time, and the tolerance parameter. If the key is accepted, the receiving peer calculates a hash and authenticates the update message.

Initial application of a keychain to a TCP session causes the session to reset. However, once the keychain is applied, the addition or removal of a password from the keychain does not cause the TCP session to reset. Also, the TCP session does not reset when the keychain changes from one authentication algorithm to another.

Example: Configuring Route Authentication for BGP

All BGP protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in autonomous system (AS) routing updates. By default, authentication is disabled.

- [Requirements on page 3430](#)
- [Overview on page 3430](#)
- [Configuration on page 3431](#)
- [Verification on page 3433](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

Overview

When you configure authentication, the algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

This example includes the following statements for configuring and applying the keychain:

- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.

The key can be up to 126 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
- **tolerance**—(Optional) For each keychain, you can configure a clock-skew tolerance value in seconds. The clock-skew tolerance is applicable to the receiver accepting keys for BGP updates. The configurable range is 0 through 999,999,999 seconds. During the tolerance period, either the current or previous password is acceptable.
- **key-chain**—For each keychain, you must specify a name. This example defines one keychain: **bgp-auth**. You can have multiple keychains on a routing device. For example, you can have a keychain for BGP, a keychain for OSPF, and a keychain for LDP.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the keychain.
- **authentication-key-chain**—Enables you to apply a keychain at the global BGP level for all peers, for a group, or for a neighbor. This example applies the keychain to the peers defined in the external BGP (EBGP) group called **ext**.
- **authentication-algorithm**—For each keychain, you can specify a hashing algorithm. The algorithm can be AES-128, MD5, or SHA-1.


You associate a keychain and an authentication algorithm with a BGP neighboring session.

This example configures a keychain named **bgp-auth**. Key 0 will be sent and accepted starting at 2011-6-23.20:19:33 -0700, and will stop being sent and accepted when the next key in the keychain (key 1) becomes active. Key 1 becomes active one year later at 2012-6-23.20:19:33 -0700, and will not stop being sent and accepted unless another key is configured with a start time that is later than the start time of key 1. A clock-skew tolerance of 30 seconds applies to the receiver accepting the keys. During the tolerance period, either the current or previous key is acceptable. The keys are shared-secret passwords. This means that the neighbors receiving the authenticated routing updates must have the same authentication keychain configuration, including the same keys (passwords). So Router R0 and Router R1 must have the same authentication-key-chain configuration if they are configured as peers. This example shows the configuration on only one of the routing devices.

Topology Diagram

Figure 103 on page 3431 shows the topology used in this example.

Figure 103: Authentication for BGP



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65530
set protocols bgp group ext neighbor 172.16.2.1
set routing-options autonomous-system 65533
```

```
set protocols bgp group ext authentication-key-chain bgp-auth
set protocols bgp group ext authentication-algorithm md5
set security authentication-key-chains key-chain bgp-auth tolerance 30
set security authentication-key-chains key-chain bgp-auth key 0 secret
  this-is-the-secret-password
set security authentication-key-chains key-chain bgp-auth key 0 start-time
  2011-6-23.20:19:33-0700
set security authentication-key-chains key-chain bgp-auth key 1 secret
  this-is-another-secret-password
set security authentication-key-chains key-chain bgp-auth key 1 start-time
  2012-6-23.20:19:33-0700
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.

```
[edit routing-options]
user@R1# set autonomous-system 65533
```

2. Configure one or more BGP groups.

```
[edit protocols bgp group ext]
user@R1# set type external
user@R1# set peer-as 65530
user@R1# set neighbor 172.16.2.1
```

3. Configure authentication with multiple keys.

```
[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set key 0 secret this-is-the-secret-password
user@R1# set key 0 start-time 2011-6-23.20:19:33-0700
user@R1# set key 1 secret this-is-another-secret-password
user@R1# set key 1 start-time 2012-6-23.20:19:33-0700
```

The start time of each key must be unique within the keychain.

4. Apply the authentication keychain to BGP, and set the hashing algorithm.

```
[edit protocols bgp group ext]
user@R1# set authentication-key-chain bgp-auth
user@R1# set authentication-algorithm md5
```

5. (Optional) Apply a clock-skew tolerance value in seconds.

```
[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set tolerance 30
```

Results From configuration mode, confirm your configuration by entering the **show protocols**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
bgp {
```

```

group ext {
  type external;
  peer-as 65530;
  neighbor 172.16.2.1;
  authentication-key-chain bgp-auth;
  authentication-algorithm md5;
}
}

user@R1# show routing-options
autonomous-system 65533;

user@R1# show security
authentication-key-chains {
  key-chain bgp-auth {
    tolerance 30;
    key 0 {
      secret
        "$9$ST6AREyK8RhXNdwaJn/Ct0IcykWWx9AylMWdVgoJDjqP5FCA0z3IEhcMWLxNbgJDiF6A";
      ## SECRET-DATA
      start-time "2011-6-23.20:19:33 -0700";
    }
    key 1 {
      secret "$9$UyD.59Cu0Ih9AylKW-dqmfT369CuRhSP5hrvMN-JGDiqfu0lleWpuh.";
      ## SECRET-DATA
      start-time "2012-6-23.20:19:33 -0700";
    }
  }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

Verification

Confirm that the configuration is working properly.

- [Verifying Authentication for the Neighbor on page 3433](#)
- [Verifying That Authorization Messages Are Sent on page 3434](#)
- [Checking Authentication Errors on page 3435](#)
- [Verifying the Operation of the Keychain on page 3435](#)

Verifying Authentication for the Neighbor

Purpose Make sure that the **AuthKeyChain** option appears in the output of the **show bgp neighbor** command.

Action From operational mode, enter the **show bgp neighbor** command.

```

user@R1> show bgp neighbor
Peer: 172.16.2.1+179 AS 65530  Local: 172.16.2.2+1222 AS 65533
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None

```

```

Export: [ direct-lo0 ]
Options: <Preference PeerAS Refresh>
Options: <AuthKeyChain>
Authentication key is configured
Authentication key chain: jni
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 172.16.2.1      Local ID: 10.255.124.35   Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
Local Interface: fe-0/0/1.0
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 2    Sent 2    Checked 2
Input messages: Total 21    Updates 2    Refreshes 0    Octets 477
Output messages: Total 22    Updates 1    Refreshes 0    Octets 471
Output Queue[0]: 0

```

Verifying That Authorization Messages Are Sent

Purpose Confirm that BGP has the enhanced authorization option.

Action From operational mode, enter the **monitor traffic interface fe-0/0/1** command.

```

user@R1> monitor traffic interface fe-0/0/1
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on fe-0/0/1, capture size 96 bytes

13:08:00.618402 In arp who-has 172.16.2.66 tell 172.16.2.69
13:08:02.408249 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P
1889289217:1889289235(18) ack 2215740969 win 58486 <nop,nop,timestamp 167557
1465469,nop,Enhanced Auth keyid 0 diglen 12 digest: fe3366001f45767165f17037>:
13:08:02.418396 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 1:19(18) ack 18 win
57100 <nop,nop,timestamp 1466460 167557,nop,Enhanced Auth keyid 0 diglen 12
digest: a18c31eda1b14b2900921675>:
13:08:02.518146 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 19 win 58468
<nop,nop,timestamp 167568 1466460,nop,Enhanced Auth keyid 0 diglen 12 digest:
c3b6422eb6bd3fd9cf79742b>
13:08:28.199557 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: P
286842489:286842508(19) ack 931203976 win 57200 <nop,Enhanced Auth keyid 0
diglen 12 digest: fc0e42900a73736bcc07c1a4>: BGP, length: 19
13:08:28.209661 In IP 172.16.2.1.bgp > 172.16.2.2.nerv: P 1:20(19) ack 19 win
56835 <nop,Enhanced Auth keyid 0 diglen 12 digest: 0fc8578c489fabce63aeb2c3>:
BGP, length: 19
13:08:28.309525 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: . ack 20 win 57181
<nop,Enhanced Auth keyid 0 diglen 12 digest: ef03f282fb2ece0039491df8>
13:08:32.439708 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P 54:72(18) ack 55 win
58432 <nop,nop,timestamp 170560 1468472,nop,Enhanced Auth keyid 0 diglen 12
digest: 76e0cf926f348b726c631944>:
13:08:32.449795 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 55:73(18) ack 72 win
57046 <nop,nop,timestamp 1469463 170560,nop,Enhanced Auth keyid 0 diglen 12
digest: dae3eec390d18a114431f4d8>:
13:08:32.549726 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 73 win 58414
<nop,nop,timestamp 170571 1469463,nop,Enhanced Auth keyid 0 diglen 12 digest:

```

```

851df771aee2ea7a43a0c46c>
13:08:33.719880 In arp who-has 172.16.2.66 tell 172.16.2.69
^C
35 packets received by filter
0 packets dropped by kernel

```

Checking Authentication Errors

Purpose Check the number of packets dropped by TCP because of authentication errors.

Action From operational mode, enter the **show system statistics tcp | match auth** command.

```

user@R1> show system statistics tcp | match auth
0 send packets dropped by TCP due to auth errors
58 rcv packets dropped by TCP due to auth errors

```

Verifying the Operation of the Keychain

Purpose Check the number of packets dropped by TCP because of authentication errors.

Action From operational mode, enter the **show security keychain detail** command.

```

user@R1> show security keychain detail
keychain          Active-ID      Next-ID      Transition  Tolerance
                  Send Receive   Send Receive
bgp-auth          3      3      1      1      1d 23:58    30
Id 3, Algorithm hmac-md5, State send-receive, Option basic
Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
Id 1, Algorithm hmac-md5, State inactive, Option basic
Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Examples: Configuring TCP and BGP Security

- [Understanding Security Options for BGP with TCP on page 3435](#)
- [Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers on page 3436](#)
- [Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List on page 3441](#)
- [Example: Limiting TCP Segment Size for BGP on page 3444](#)

Understanding Security Options for BGP with TCP

Among routing protocols, BGP is unique in using TCP as its transport protocol. BGP peers are established by manual configuration between routing devices to create a TCP session on port 179. A BGP-enabled device periodically sends keepalive messages to maintain the connection.

Over time, BGP has become the dominant interdomain routing protocol on the Internet. However, it has limited guarantees of stability and security. Configuring security options for BGP must balance suitable security measures with acceptable costs. No one method

has emerged as superior to other methods. Each network administrator must configure security measures that meet the needs of the network being used.

For detailed information about the security issues associated with BGP's use of TCP as a transport protocol, see RFC 4272, *BGP Security Vulnerabilities Analysis*.

Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers

This example shows how to configure a standard stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except from specified BGP peers.

- [Requirements on page 3436](#)
- [Overview on page 3436](#)
- [Configuration on page 3436](#)
- [Verification on page 3439](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.


Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except the specified BGP peers.

The stateless firewall filter **filter_bgp179** matches all packets from the directly connected interfaces on Device A and Device B to the destination port number 179.

Figure 104 on page 3436 shows the topology used in this example. Device C attempts to make a TCP connection to Device E. Device E blocks the connection attempt. This example shows the configuration on Device E.

Figure 104: Typical Network with BGP Peer Sessions



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.


```

Device C    set interfaces ge-1/2/0 unit 10 description to-E
            set interfaces ge-1/2/0 unit 10 family inet address 10.10.10.10/30
            set protocols bgp group external-peers type external
            set protocols bgp group external-peers peer-as 17
            set protocols bgp group external-peers neighbor 10.10.10.9
            set routing-options autonomous-system 22

Device E    set interfaces ge-1/2/0 unit 0 description to-A
            set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
            set interfaces ge-1/2/1 unit 5 description to-B
            set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
            set interfaces ge-1/0/0 unit 9 description to-C
            set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
            set interfaces lo0 unit 2 family inet filter input filter_bgp179
            set interfaces lo0 unit 2 family inet address 192.168.0.1/32
            set protocols bgp group external-peers type external
            set protocols bgp group external-peers peer-as 22
            set protocols bgp group external-peers neighbor 10.10.10.2
            set protocols bgp group external-peers neighbor 10.10.10.6
            set protocols bgp group external-peers neighbor 10.10.10.10
            set routing-options autonomous-system 17
            set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.2/32
            set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.6/32
            set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
            set firewall family inet filter filter_bgp179 term 1 then accept
            set firewall family inet filter filter_bgp179 term 2 then reject

```

Configuring Device E

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E with a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requestors except specified BGP peers:

1. Configure the interfaces.


```

user@E# set interfaces ge-1/2/0 unit 0 description to-A
user@E# set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30

user@E# set interfaces ge-1/2/1 unit 5 description to-B
user@E# set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30

user@E# set interfaces ge-1/0/0 unit 9 description to-C
user@E# set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30

```
2. Configure BGP.


```

[edit protocols bgp group external-peers]
user@E# set type external
user@E# set peer-as 22
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10

```

3. Configure the autonomous system number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

4. Define the filter term that accepts TCP connection attempts to port 179 from the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 1 from source-address 10.10.10.2/32
user@E# set term 1 from source-address 10.10.10.6/32
user@E# set term 1 from destination-port bgp
user@E# set term 1 then accept
```

5. Define the other filter term to reject packets from other sources.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 2 then reject
```

6. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 2 family inet]
user@E# set filter input filter_bgp179
user@E# set address 192.168.0.1/32
```

Results From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@E# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          10.10.10.2/32;
          10.10.10.6/32;
        }
        destination-port bgp;
      }
      then accept;
    }
    term 2 {
      then {
        reject;
      }
    }
  }
}

user@E# show interfaces
lo0 {
  unit 2 {
    family inet {
      filter {
        input filter_bgp179;
```

```

    }
    address 192.168.0.1/32;
  }
}
ge-1/2/0 {
  unit 0 {
    description to-A;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
ge-1/2/1 {
  unit 5 {
    description to-B;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
ge-1/0/0 {
  unit 9 {
    description to-C;
    family inet {
      address 10.10.10.9/30;
    }
  }
}

user@E# show protocols
bgp {
  group external-peers {
    type external;
    peer-as 22;
    neighbor 10.10.10.2;
    neighbor 10.10.10.6;
    neighbor 10.10.10.10;
  }
}

user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Filter Is Configured on page 3440](#)
- [Verifying the TCP Connections on page 3440](#)
- [Monitoring Traffic on the Interfaces on page 3440](#)

Verifying That the Filter Is Configured

Purpose Make sure that the filter is listed in output of the **show firewall filter** command.

Action user@E> show firewall filter filter_bgp179
Filter: filter_bgp179

Verifying the TCP Connections

Purpose Verify the TCP connections.

Action From operational mode, run the **show system connections extensive** command on Device C and Device E.

The output on Device C shows the attempt to establish a TCP connection. The output on Device E shows that connections are established with Device A and Device B only.

user@C> show system connections extensive | match 10.10.10

tcp4	0	0	10.10.10.9.51872	10.10.10.10.179	SYN_SENT
------	---	---	------------------	-----------------	----------

user@E> show system connections extensive | match 10.10.10

tcp4	0	0	10.10.10.5.179	10.10.10.6.62096	ESTABLISHED
tcp4	0	0	10.10.10.6.62096	10.10.10.5.179	ESTABLISHED
tcp4	0	0	10.10.10.1.179	10.10.10.2.61506	ESTABLISHED
tcp4	0	0	10.10.10.2.61506	10.10.10.1.179	ESTABLISHED

Monitoring Traffic on the Interfaces

Purpose Use the **monitor traffic** command to compare the traffic on an interface that establishes a TCP connection with the traffic on an interface that does not establish a TCP connection.

Action From operational mode, run the **monitor traffic** command on the Device E interface to Device B and on the Device E interface to Device C. The following sample output verifies that in the first example, acknowledgment (**ack**) messages are received. In the second example, **ack** messages are not received.

```
user@E> monitor traffic size 1500 interface ge-1/2/1.5
19:02:49.700912 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P
3330573561:3330573580(19) ack 915601686 win 16384 <nop,nop,timestamp 1869518816
1869504850>: BGP, length: 19
19:02:49.801244 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 19 win 16384
<nop,nop,timestamp 1869518916 1869518816>
19:03:03.323018 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: P 1:20(19) ack 19 win
16384 <nop,nop,timestamp 1869532439 1869518816>: BGP, length: 19
19:03:03.422418 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: . ack 20 win 16384
<nop,nop,timestamp 1869532539 1869532439>
19:03:17.220162 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P 19:38(19) ack 20 win
16384 <nop,nop,timestamp 1869546338 1869532439>: BGP, length: 19
19:03:17.320501 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 38 win 16384
<nop,nop,timestamp 1869546438 1869546338>
```

user@E> monitor traffic size 1500 interface ge-1/0/0.9

```
18:54:20.175471 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
```

```

win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869009240 0,sackOK,eol>
18:54:23.174422 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869012240 0,sackOK,eol>
18:54:26.374118 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869015440 0,sackOK,eol>
18:54:29.573799 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:32.773493 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:35.973185 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>

```

Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

This example shows how to configure a standard stateless firewall filter that limits certain TCP and Internet Control Message Protocol (ICMP) traffic destined for the Routing Engine by specifying a list of prefix sources that contain allowed BGP peers.

- [Requirements on page 3441](#)
- [Overview on page 3441](#)
- [Configuration on page 3441](#)
- [Verification on page 3443](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except BGP peers that have a specified prefix.

A source prefix list, **plist_bgp179**, is created that specifies the list of source prefixes that contain allowed BGP peers.

The stateless firewall filter **filter_bgp179** matches all packets from the source prefix list **plist_bgp179** to the destination port number 179.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

- [Configure the Filter on page 3442](#)
- [Results on page 3442](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set policy-options prefix-list plist_bgp179 apply-path "protocols bgp group <*> neighbor
<*>"
set firewall family inet filter filter_bgp179 term 1 from source-address 0.0.0.0/0

```

```

set firewall family inet filter filter_bgp179 term 1 from source-prefix-list plist_bgp179 except
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then reject
set firewall family inet filter filter_bgp179 term 2 then accept
set interfaces lo0 unit 0 family inet filter input filter_bgp179
set interfaces lo0 unit 0 family inet address 127.0.0.1/32

```

Configure the Filter

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the filter:

1. Expand the prefix list **bgp179** to include all prefixes pointed to by the BGP peer group defined by **protocols bgp group <*> neighbor <*>**.

```

[edit policy-options prefix-list plist_bgp179]
user@host# set apply-path "protocols bgp group <*> neighbor <*>"

```

2. Define the filter term that rejects TCP connection attempts to port 179 from all requesters except the specified BGP peers.

```

[edit firewall family inet filter filter_bgp179]
user@host# set term term1 from source-address 0.0.0.0/0
user@host# set term term1 from source-prefix-list bgp179 except
user@host# set term term1 from destination-port bgp
user@host# set term term1 then reject

```

3. Define the other filter term to accept all packets.

```

[edit firewall family inet filter filter_bgp179]
user@host# set term term2 then accept

```

4. Apply the firewall filter to the loopback interface.

```

[edit interfaces lo0 unit 0 family inet]
user@host# set filter input filter_bgp179
user@host# set address 127.0.0.1/32

```

Results

From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          0.0.0.0/0;
        }
        source-prefix-list {
          plist_bgp179 except;

```

```

    }
    destination-port bgp;
  }
  then {
    reject;
  }
}
term 2 {
  then {
    accept;
  }
}
}
}

user@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      filter {
        input filter_bgp179;
      }
      address 127.0.0.1/32;
    }
  }
}

user@host# show policy-options
prefix-list plist_bgp179 {
  apply-path "protocols bgp group <*> neighbor <*>";
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure, where appropriate, for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

Verification

Confirm that the configuration is working properly.

Displaying the Firewall Filter Applied to the Loopback Interface

Purpose Verify that the firewall filter **filter_bgp179** is applied to the IPv4 input traffic at logical interface **lo0.0**.

Action Use the **show interfaces statistics** operational mode command for logical interface **lo0.0**, and include the **detail** option. Under the **Protocol inet** section of the command output section, the **Input Filters** field displays the name of the stateless firewall filter applied to the logical interface in the input direction:

```

[edit]
user@host> show interfaces statistics lo0.0 detail
Logical interface lo0.0 (Index 321) (SNMP ifIndex 16) (Generation 130)
Flags: SNMP-Traps Encapsulation: Unspecified
Traffic statistics:
Input bytes : 0

```

```
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: Unlimited, Generation: 145, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Input Filters: filter_bgp179
Addresses, Flags: Primary
Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 138
```

Example: Limiting TCP Segment Size for BGP

This example shows how to avoid Internet Control Message Protocol (ICMP) vulnerability issues by limiting TCP segment size when you are using maximum transmission unit (MTU) discovery. Using MTU discovery on TCP paths is one method of avoiding BGP packet fragmentation.

- [Requirements on page 3444](#)
- [Overview on page 3444](#)
- [Configuration on page 3445](#)
- [Verification on page 3447](#)
- [Troubleshooting on page 3447](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

TCP negotiates a maximum segment size (MSS) value during session connection establishment between two peers. The MSS value negotiated is primarily based on the maximum transmission unit (MTU) of the interfaces to which the communicating peers are directly connected. However, due to variations in link MTU on the path taken by the TCP packets, some packets in the network that are well within the MSS value might be fragmented when the packet size exceeds the link's MTU.

To configure the TCP MSS value, include the **tcp-mss** statement with a segment size from 1 through 4096.

If the router receives a TCP packet with the SYN bit and the MSS option set, and the MSS option specified in the packet is larger than the MSS value specified by the **tcp-mss** statement, the router replaces the MSS value in the packet with the lower value specified by the **tcp-mss** statement.


The configured MSS value is used as the maximum segment size for the sender. The assumption is that the TCP MSS value used by the sender to communicate with the BGP neighbor is the same as the TCP MSS value that the sender can accept from the BGP neighbor. If the MSS value from the BGP neighbor is less than the MSS value configured, the MSS value from the BGP neighbor is used as the maximum segment size for the sender.

This feature is supported with TCP over IPv4 and TCP over IPv6.

Topology Diagram

Figure 105 on page 3445 shows the topology used in this example.

Figure 105: TCP Maximum Segment Size for BGP



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
R0
set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
set interfaces lo0 unit 1 family inet address 10.255.14.179/32
set protocols bgp group-int tcp-mss 2020
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.179
set protocols bgp group int mtu-discovery
set protocols bgp group int neighbor 10.255.71.24 tcp-mss 2000
set protocols bgp group int neighbor 10.255.14.177
set protocols bgp group int neighbor 10.0.14.4 tcp-mss 4000
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface 10.255.14.179
set routing-options autonomous-system 65000
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R0:

1. Configure the interfaces.

```
[edit interfaces]
user@R0# set fe-1/2/0 unit 1 family inet address 1.1.0.1/30
user@R0# set lo0 unit 1 family inet address 10.255.14.179/32
```
2. Configure an interior gateway protocol (IGP), OSPF in this example.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R0# set interface fe-1/2/0.1
user@R0# set interface 10.255.14.179
```

3. Configure one or more BGP groups.

```
[edit protocols bgp group int]
user@R0# set type internal
user@R0# set local-address 10.255.14.179
```

4. Configure MTU discovery to prevent packet fragmentation.

```
[edit protocols bgp group int]
user@R0# set mtu-discovery
```

5. Configure the BGP neighbors, with the TCP MSS set globally for the group or specifically for the various neighbors.

```
[edit protocols bgo group int]
user@R0# set tcp-mss 2020
user@R0# set neighbor 10.255.14.177
user@R0# set neighbor 10.255.71.24 tcp-mss 2000
user@R0# set neighbor 10.0.14.4 tcp-mss 4000
```



NOTE: The TCP MSS neighbor setting overrides the group setting.

6. Configure the local autonomous system.

```
[edit routing-options]
user@R0# set autonomous-system 65000
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.255.14.179/32;
    }
  }
}

user@R0# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.179;
```

```

mtu-discovery;
tcp-mss 2020;
neighbor 10.255.71.24 {
    tcp-mss 2000;
}
neighbor 10.255.14.177;
neighbor 10.0.14.4 {
    tcp-mss 4000;
}
}
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/0.1;
        interface 10.255.14.179;
    }
}
}

user@R0# show routing-options
autonomous-system 65000;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, run the following commands:

- **show system connections extensive | find <neighbor-address>**, to check the negotiated TCP MSS value.
- **monitor traffic interface**, to monitor BGP traffic and to make sure that the configured TCP MSS value is used as the MSS option in the TCP SYN packet.

Troubleshooting

- [MSS Calculation with MTU Discovery on page 3447](#)

MSS Calculation with MTU Discovery

Problem Consider an example in which two routing devices (R1 and R2) have an internal BGP (IBGP) connection. On both of the routers, the connected interfaces have 4034 as the IPv4 MTU.

```

user@R1# show protocols bgp | display set
[edit]
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 45.45.45.2
set protocols bgp group ibgp mtu-discovery
set protocols bgp group ibgp neighbor 45.45.45.1

```

```

user@R1# run show interfaces xe-0/0/3 extensive | match mtu

```

```

Link-level type: Ethernet, MTU: 4048, LAN-PHY mode, Speed: 10Gbps,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Protocol inet, MTU: 4034, Generation: 180, Route table: 0
Protocol multiservice, MTU: Unlimited, Generation: 181, Route table: 0

```

In the following packet capture on Device R1, the negotiated MSS is 3994. In the **show system connections extensive** information for MSS, it is set to 2048.

```
05:50:01.575218 Out
  Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
  -----original packet-----
  00:21:59:e1:e8:03 > 00:19:e2:20:79:01, ethertype IPv4 (0x0800), length
78: (tos 0xc0, ttl 64, id 53193, offset 0, flags [DF], proto: TCP (6), length:
64) 45.45.45.2.62840 > 45.45.45.1.bgp: S 2939345813:2939345813(0) win 16384 **mss
3994,nop,wscale 0,nop,nop,timestamp 70559970 0,sackOK,eol>
05:50:01.575875 In
  Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
  -----original packet-----
  PFE proto 2 (ipv4): (tos 0xc0, ttl 255, id 37709, offset 0, flags [DF], proto:
TCP (6), length: 64) 45.45.45.1.bgp > 45.45.45.2.62840: S 2634967984:2634967984(0)
ack 2939345814 win 16384 **mss 3994,nop,wscale 0,nop,nop,timestamp 174167273
70559970,sackOK,eol>
```

user@R1# run show system connections extensive | find 45.45

```
tcp4      0      0 45.45.45.2.62840          45.45.45.1.179
ESTABLISHED
  sndsbcc:      0 sndsbmbcnt:      0 sndsbmbmax:      131072
  sndsblowat:    2048 sndsbhiwat:    16384
  rcvsbcc:      0 rcvsbmbcnt:      0 rcvsbmbmax:      131072
  rcvsblowat:    1 rcvsbhiwat:    16384
  proc id:      19725 proc name:      rpd
    iss: 2939345813      sndup: 2939345972
  snduna: 2939345991      sndnxt: 2939345991      sndwnd:      16384
  sndmax: 2939345991      sndcwnd:      10240 sndssthresh: 1073725440
  irs: 2634967984      rcvup: 2634968162
  rcvnxt: 2634968162      rcvadv: 2634984546      rcvwnd:      16384
  rtt:      0      srtt:      1538      rttv:      1040
  rxtcur:      1200 rxtshift:      0      rtseq: 2939345972
  rttmin:      1000 mss:      2048
```

Solution This is expected behavior with Junos OS. The MSS value is equal to the MTU value minus the IP or IPv6 and TCP headers. This means that the MSS value is generally 40 bytes less than the MTU (for IPv4) and 60 bytes less than the MTU (for IPv6). This value is negotiated between the peers. In this example, it is $4034 - 40 = 3994$. Junos OS then rounds this value to a multiple of 2 KB. The value is $3994 / 2048 * 2048 = 2048$. So it is not necessary to see same MSS value with in the **show system connections** output.

$3994 / 2048 = 1.95$

1.95 is rounded to 1.

$1 * 2048 = 2048$

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3149](#)
 - [BGP Configuration Overview](#)

BGP Flap Configuration

- [Example: Preventing BGP Session Resets on page 3449](#)
- [Examples: Configuring BGP Flap Damping on page 3456](#)

Example: Preventing BGP Session Resets

- [Understanding BGP Session Resets on page 3449](#)
- [Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3449](#)

Understanding BGP Session Resets

Certain configuration actions and events cause BGP sessions to be reset (dropped and then reestablished).

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same autonomous system (AS) number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an internal BGP (IBGP) group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.
- Changing configuration statements that affect BGP peers, such as renaming a BGP group, resets the BGP sessions.
- If you change the address family specified in the **[edit protocols bgp family]** hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Example: Preventing BGP Session Flaps When VPN Families Are Configured

This example shows a workaround for a known issue in which BGP sessions sometimes go down and then come back up (in other words, flap) when virtual private network (VPN) families are configured. If any VPN family (for example, **inet-vpn**, **inet6-vpn**, **inet-mpvn**, **inet-mdt**, **inet6-mpvn**, **l2vpn**, **iso-vpn**, and so on) is configured on a BGP master instance, a flap of either a route reflector (RR) internal BGP (IBGP) session or an external BGP (EBGP) session causes flaps of other BGP sessions configured with the same VPN family.

- [Requirements on page 3450](#)
- [Overview on page 3451](#)
- [Configuration on page 3452](#)
- [Verification on page 3455](#)

Requirements

Before you begin:

- Configure router interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure VPNs.

Overview

When a router or switch is configured as either a route reflector (RR) or an AS boundary router (an external BGP peer) and a VPN family (for example, the **family inet-vpn unicast** statement) is configured, a flap of either the RR IBGP session or the EBGp session causes flaps of all other BGP sessions that are configured with the **family inet-vpn unicast** statement. This example shows how to prevent these unnecessary session flaps.

The reason for the flapping behavior is related to BGP operation in Junos OS when originating VPN routes.

BGP has the following two modes of operation with respect to originating VPN routes:

- If BGP does not need to propagate VPN routes because the session has no EBGp peer and no RR clients, BGP exports VPN routes directly from the **instance.inet.0** routing table to other PE routers. This behavior is efficient in that it avoids the creation of two copies of many routes (one in the **instance.inet.0** table and one in the **bgp.l3vpn.0** table).
- If BGP does need to propagate VPN routes because the session has an EBGp peer or RR clients, BGP first exports the VPN routes from the **instance.inet.0** table to the **bgp.l3vpn.0** table. Then BGP exports the routes to other PE routers. In this scenario, two copies of the route are needed to enable best-route selection. A PE router might receive the same VPN route from a CE device and also from an RR client or EBGp peer.

When, because of a configuration change, BGP transitions from needing two copies of a route to not needing two copies of a route (or the reverse), all sessions over which VPN routes are exchanged go down and then come back up. Although this example focuses on the **family inet-vpn unicast** statement, the concept applies to all VPN network layer reachability information (NLRI) families. This issue impacts logical systems as well. All BGP sessions in the master instance related to the VPN NLRI family are brought down to implement the table advertisement change for the VPN NLRI family. Changing an RR to a non-RR or the reverse (by adding or removing the **cluster** statement) causes the table advertisement change. Also, configuring the first EBGp session or removing the EBGp session from the configuration in the master instance for a VPN NLRI family causes the table advertisement change.

The way to prevent these unnecessary session flaps is to configure an extra RR client or EBGp session as a passive session with a neighbor address that does not exist. This example focuses on the EBGp case, but the same workaround works for the RR case.

When a session is passive, the routing device does not send Open requests to a peer. Once you configure the routing device to be passive, the routing device does not originate the TCP connection. However, when the routing device receives a connection from the peer and an Open message, it replies with another BGP Open message. Each routing device declares its own capabilities.

[Figure 106 on page 3452](#) shows the topology for the EBGp case. Router R1 has an IBGP session with Routers R2 and R3 and an EBGp session with Router R4. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R4 EBGp session flaps, the R1-R2 and R1-R3 BGP sessions flap also.

Figure 106: Topology for the EBGP Case





Figure 107 on page 3452 shows the topology for the RR case. Router R1 is the RR, and Router R3 is the client. Router R1 has IBGP sessions with Routers R2 and R3. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R3 session flaps, the R1-R2 and R1-R4 sessions flap also.

Figure 107: Topology for the RR Case



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp family inet-vpn unicast
set protocols bgp family l2vpn signaling
set protocols bgp group R1-R4 type external
set protocols bgp group R1-R4 local-address 4.4.4.2
set protocols bgp group R1-R4 neighbor 4.4.4.1 peer-as 200
set protocols bgp group R1-R2-R3 type internal
set protocols bgp group R1-R2-R3 log-updown
set protocols bgp group R1-R2-R3 local-address 15.15.15.15
set protocols bgp group R1-R2-R3 neighbor 12.12.12.12
set protocols bgp group R1-R2-R3 neighbor 13.13.13.13
set protocols bgp group Fake type external
```



```
set protocols bgp group Fake passive
set protocols bgp group Fake neighbor 100.100.100.100 peer-as 500
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the EBGp scenario:

1. Configure one or more VPN families.

```
[edit protocols bgp]
user@R1# set family inet-vpn unicast
user@R1# set family l2vpn signaling
```

2. Configure the EBGp session.

```
[edit protocols bgp]
user@R1# set group R1-R4 type external
user@R1# set group R1-R4 local-address 4.4.4.2
user@R1# set group R1-R4 neighbor 4.4.4.1 peer-as 200
```

3. Configure the IBGP sessions.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 type internal
user@R1# set group R1-R2-R3 local-address 15.15.15.15
user@R1# set group R1-R2-R3 neighbor 12.12.12.12
user@R1# set group R1-R2-R3 neighbor 13.13.13.13
```

4. (Optional) Configure BGP so that it generates a **syslog** message whenever a BGP peer makes a state transition.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 log-updown
```

Enabling the **log-updown** statement causes BGP state transitions to be logged at **warning** level.

Step-by-Step Procedure To verify that unnecessary session flaps are occurring:

1. Run the **show bgp summary** command to verify that the sessions have been established.

```
user@R1> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0 0 0 0 0 0
bgp.12vpn.0 0 0 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1 200 6 5 0 0 1:08 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 3 7 0 0 1:18 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
13.13.13.13 100 3 6 0 0 1:14 Establ
```

```

bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0

```

2. Deactivate the EBGP session.

```

user@R1# deactivate group R1-R4
user@R1# commit

```

```

Mar 10 18:27:40 R1: rpd[1464]: bgp_peer_delete:6589: NOTIFICATION sent to 4.4.4.1 (External AS 200): code
 6 (Cease) subcode 3 (Peer Unconfigured), Reason: Peer Deletion
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 12.12.12.12 (Internal AS
100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table
advertise
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 13.13.13.13 (Internal AS
100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table
advertise

```

3. Run the **show bgp summary** command to view the session flaps.

```

user@R1> show bgp summary
Groups: 1 Peers: 2 Down peers: 2
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0          0          0          0          0          0
bgp.12vpn.0 0          0          0          0          0          0
inet.0      0          0          0          0          0          0
Peer      AS   InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 4      9      0      1      19 Active
13.13.13.13 100 4      8      0      1      19 Active

user@R1> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0          0          0          0          0          0
bgp.12vpn.0 0          0          0          0          0          0
inet.0      0          0          0          0          0          0
Peer      AS   InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 2      3      0      1      0 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
13.13.13.13 100 2      3      0      1      0 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To prevent unnecessary BGP session flaps:

1. Add a passive EBGP session with a neighbor address that does not exist in the peer autonomous system (AS).

```

[edit protocols bgp]
user@R1# set group Fake type external
user@R1# set group Fake passive
user@R1# set neighbor 100.100.100.100 peer-as 500

```

2. Run the **show bgp summary** command to verify that the real sessions have been established and the passive session is idle.

```

user@R1> show bgp summary
Groups: 3 Peers: 4 Down peers: 1
Table          Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0          0          0          0          0          0
bgp.12vpn.0 0          0          0          0          0          0
Peer           AS   InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1        200 9500  9439   0   0    2d   23:14:23 Estab1
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12    100 10309 10239   0   0    3d   5:17:49 Estab1
bgp.13vpn.0: 0/0/0/0
13.13.13.13    100 10306 10241   0   0    3d   5:18:25 Estab1
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0      0      0   0    2d   23:38:52 Idle

```

Verification

Confirm that the configuration is working properly.

- [Bringing Down the EBGp Session on page 3455](#)
- [Verifying That the IBGP Sessions Remain Up on page 3455](#)

Bringing Down the EBGp Session

Purpose Try to cause the flap issue after the workaround is configured.

Action user@R1# deactivate group R1-R4
user@R1# commit

Verifying That the IBGP Sessions Remain Up

Purpose Make sure that the IBGP sessions do not flap after the EBGp session is deactivated.

Action user@R1> show bgp summary

```

Groups: 2 Peers: 3 Down peers: 1
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0      0
Peer      AS  InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 10312 10242 0 0 3d 5:19:01 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10309 10244 0 0 3d 5:19:37 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 2d 23:40:04 Idle

```

user@R1> show bgp summary

```

Groups: 3 Peers: 4 Down peers: 1
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0
Peer      AS  InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1    200 5 4 0 0 28 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 10314 10244 0 0 3d 5:19:55 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10311 10246 0 0 3d 5:20:31 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 2d 23:40:58 Idle

```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3149](#)
 - [BGP Configuration Overview](#)

Examples: Configuring BGP Flap Damping

- [Understanding Damping Parameters on page 3456](#)
- [Example: Configuring Damping Parameters on page 3457](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3466](#)

Understanding Damping Parameters

BGP *route flapping* describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. BGP *flap damping* is a method of reducing the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do

not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level, which is supported in Junos OS Release 12.2 and later. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

By default, route flap damping is not enabled. Damping is applied to external peers and to peers at confederation boundaries.

When you enable damping, default parameters are applied, as summarized in [Table 286 on page 3457](#).

Table 286: Damping Parameters

Damping Parameter	Description	Default Value	Possible Values
half-life <i>minutes</i>	Decay half-life—Number of minutes after which an arbitrary value is halved if a route stays stable.	15 (minutes)	1 through 45
max-suppress <i>minutes</i>	Maximum hold-down time for a route, in minutes.	60 (minutes)	1 through 720
reuse	Reuse threshold—Arbitrary value below which a suppressed route can be used again.	750	1 through 20,000
suppress	Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements.	3000	1 through 20,000

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

Example: Configuring Damping Parameters

This example shows how to configure damping parameters.

- [Requirements on page 3457](#)
- [Overview on page 3457](#)
- [Configuration on page 3458](#)
- [Verification on page 3462](#)

Requirements

Before you begin, configure router interfaces and configure routing protocols, as explained in *Routing Policies Configuration Overview*.

Overview

This example has three routing devices. Device R2 has external BGP (EBGP) connections with Device R1 and Device R3.

Device R1 and Device R3 have some static routes configured for testing purposes, and these static routes are advertised through BGP to Device R2.


Device R2 damps routes received from Device R1 and Device R3 according to these criteria:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only the active routes are exported from the routing table.

Figure 108 on page 3458 shows the sample network.

Figure 108: BGP Flap Damping Topology



“CLI Quick Configuration” on page 3458 shows the configuration for all of the devices in Figure 108 on page 3458.

The section “Step-by-Step Procedure” on page 3459 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct-and-static
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.0.0.2
set policy-options policy-statement send-direct-and-static term 1 from protocol direct
set policy-options policy-statement send-direct-and-static term 1 from protocol static
set policy-options policy-statement send-direct-and-static term 1 then accept
set routing-options static route 172.16.0.0/16 reject
set routing-options static route 172.16.128.0/17 reject
set routing-options static route 172.16.192.0/20 reject
set routing-options static route 10.0.0.0/9 reject
set routing-options static route 224.0.0.0/7 reject
set routing-options static route 10.224.0.0/11 reject
set routing-options static route 0.0.0.0/0 reject
set routing-options autonomous-system 100

Device R2
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30

```

```

set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set protocols bgp damping
set protocols bgp group ext type external
set protocols bgp group ext import damp
set protocols bgp group ext export send-direct
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement damp term 1 from route-filter 10.128.0.0/9 exact
damping dry
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
prefix-length-range /0-/8 damping timid
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
prefix-length-range /17-/32 damping aggressive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options damping aggressive half-life 30
set policy-options damping aggressive suppress 2500
set policy-options damping timid half-life 5
set policy-options damping dry disable
set routing-options autonomous-system 200

```

Device R3

```

set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.2/30
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct-and-static
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct-and-static term 1 from protocol direct
set policy-options policy-statement send-direct-and-static term 1 from protocol static
set policy-options policy-statement send-direct-and-static term 1 then accept
set routing-options static route 10.128.0.0/9 reject
set routing-options autonomous-system 300

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure damping parameters:

1. Specify the routes to dampen and associate each group of routes with a group name.

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.2/30
user@R2# set fe-1/2/1 unit 0 family inet address 10.1.0.1/30
user@R2# set lo0 unit 0 family inet address 192.168.0.2/32

```

2. Configure the BGP neighbors.

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300

```

3. Create and configure the damping parameter groups.

```
[edit policy-options]
user@R2# set damping aggressive half-life 30
user@R2# set damping aggressive suppress 2500
user@R2# set damping timid half-life 5
user@R2# set damping dry disable
```

4. Configure the damping policy.

```
[edit policy-options policy-statement damp term 1]
user@R2# set from route-filter 10.128.0.0/9 exact damping dry
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping
    aggressive
```

5. Enable damping for BGP.

```
[edit protocols bgp]
user@R2# set damping
```

6. Apply the policy as an import policy for the BGP neighbor.

```
[edit protocols bgp group ext]
user@R2# set import damp
```



NOTE: You can refer to the same routing policy one or more times in the same or different import statements.

7. Configure an export policy.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

8. Apply the export policy.

```
[edit protocols bgp group ext]
user@R2# set export send-direct
```

9. Configure the autonomous system (AS) number.

```
user@R2# set routing-options autonomous-system 200
```

Results Confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interface
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
```



```

}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.1.0.1/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.2/32;
    }
  }
}

user@R2# show protocols
bgp {
  damping;
  group ext {
    type external;
    import damp;
    export send-direct;
    neighbor 10.0.0.1 {
      peer-as 100;
    }
    neighbor 10.1.0.2 {
      peer-as 300;
    }
  }
}

user@R2# show policy-options
policy-statement damp {
  term 1 {
    from {
      route-filter 10.128.0.0/9 exact damping dry;
      route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid;
      route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping aggressive;
    }
  }
}
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
damping aggressive {
  half-life 30;
  suppress 2500;
}
damping timid {
  half-life 5;
}
damping dry {
  disable;
}

```

```
}
}
```

```
user@R2# show routing-options
autonomous-system 200;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Causing Some Routes to Flap

Purpose To verify your route flap damping policy, some routes must flap. Having a live Internet feed almost guarantees that a certain number of route flaps will be present. If you have control over a remote system that is advertising the routes, you can modify the advertising router's policy to effect the advertisement and withdrawal of all routes or of a given prefix. In a test environment, you can cause routes to flap by clearing the BGP neighbors or by restarting the routing process on the BGP neighbors, as shown here.

Action From operational mode on Device R1 and Device R3, enter the **restart routing** command.



CAUTION: Use this command cautiously in a production network.

```
user@R1> restart routing
```

```
R1 started, pid 10474
```

```
user@R3> restart routing
```

```
R3 started, pid 10478
```

Meaning On Device R2, all of the routes from the neighbors are withdrawn and readvertised.

Checking the Route Flaps

Purpose View the number of neighbor flaps.

Action From operational mode, enter the **show bgp summary** command.

```
user@R2> show bgp summary
```

```
Groups: 1 Peers: 2 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	12	1	11	0	11	0	
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn	
State #Active/Received/Accepted/Damped...							
10.0.0.1	100	10	10	0	4	2:50	
0/9/0/9	0/0/0/0						
10.1.0.2	300	10	10	0	4	2:53	
1/3/1/2	0/0/0/0						

Meaning This output was captured after the routing process was restarted on Device R2's neighbors four times.

Verifying Route Flap Damping

Purpose Verify that routes are being hidden due to damping.

Action From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed
```

```
inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0.0.0.0/0      [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/9     [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/30    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.1.0.0/30    [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
10.224.0.0/11  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.0.0/16  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.128.0/17 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.192.0/20 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.1/32 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.3/32 [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
224.0.0.0/7    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
```

Meaning The output shows some routing instability. 11 routes are hidden due to damping.

Displaying the Details of a Damped Route

Purpose Display the details of damped routes provides useful information.

Action From operational mode, enter the **show route damping suppressed 172.16.192.0/20 detail** command.

```
user@R2> show route damping suppressed 172.16.192.0/20 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.192.0/20 (1 entry, 0 announced)
    BGP                    /-101
        Next hop type: Router, Next hop index: 758
        Address: 0x9414484
        Next-hop reference count: 9
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.0, selected
        Session Id: 0x100201
        State: <Hidden Ext>
        Local AS: 200 Peer AS: 100
        Age: 52
        Validation State: unverified
        Task: BGP_100.10.0.0.1+55922
        AS path: 100 I
        Localpref: 100
        Router ID: 192.168.0.1
        Merit (last update/now): 4278/4196
        damping-parameters: aggressive
        Last update: 00:00:52 First update: 01:01:55
        Flaps: 8
        Suppressed. Reusable in: 01:14:40
        Preference will be: 170
```

Meaning This output indicates that the displayed route has a mask length that is equal to or greater than /17, and confirms that it has been correctly mapped to the aggressive damping profile. You can also see the route's current (and last) figure of merit value, and when the route is expected to become active if it remains stable.

Verifying the Default Damping Parameters Are in Effect

Purpose Locating a damped route with a /16 mask confirms that the default parameters are in effect

Action From operational mode, enter the **show route damping suppressed detail | match 0/16** command.

```
user@R2> show route damping suppressed detail | match 0/16

172.16.0.0/16 (1 entry, 0 announced)

user@R2> show route damping suppressed 172.16.0.0/16 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.0.0/16 (1 entry, 0 announced)
    BGP                    /-101
        Next hop type: Router, Next hop index: 758
        Address: 0x9414484
        Next-hop reference count: 9
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.0, selected
        Session Id: 0x100201
        State: <Hidden Ext>
```

```

Local AS: 200 Peer AS: 100
Age: 1:58
Validation State: unverified
Task: BGP_100.10.0.0.1+55922
AS path: 100 I
Localpref: 100
Router ID: 192.168.0.1
Merit (last update/now): 3486/3202
Default damping parameters used
Last update: 00:01:58 First update: 01:03:01
Flaps: 8
Suppressed. Reusable in: 00:31:40
Preference will be: 170

```

Meaning /16 routes are not impacted by the custom damping rules. Therefore, the default damping rules are in effect.

To repeat, the custom rules are as follows:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

Filtering the Damping Information

Purpose Use OR groupings or cascaded piping to simplify the determination of what damping profile is being used for routes with a given mask length.

Action From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed detail | match "0 announced | damp"
```

```

0.0.0.0/0 (1 entry, 0 announced)
    damping-parameters: timid
10.0.0.0/9 (1 entry, 0 announced)
    Default damping parameters used
    damping-parameters: aggressive
    damping-parameters: aggressive
10.224.0.0/11 (1 entry, 0 announced)
    Default damping parameters used
172.16.0.0/16 (1 entry, 0 announced)
    Default damping parameters used
172.16.128.0/17 (1 entry, 0 announced)
    damping-parameters: aggressive
172.16.192.0/20 (1 entry, 0 announced)
    damping-parameters: aggressive
192.168.0.1/32 (1 entry, 0 announced)
    damping-parameters: aggressive
192.168.0.3/32 (1 entry, 0 announced)
    damping-parameters: aggressive
224.0.0.0/7 (1 entry, 0 announced)
    damping-parameters: timid

```

Meaning When you are satisfied that your EBGp routes are correctly associated with a damping profile, you can issue the **clear bgp damping** operational mode command to restore an active status to your damped routes, which will return your connectivity to normal operation.

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

This example shows how to configure an multiprotocol BGP multicast VPN (also called Next-Generation MVPN) with BGP route flap damping.

- [Requirements on page 3466](#)
- [Overview on page 3466](#)
- [Configuration on page 3467](#)
- [Verification on page 3474](#)

Requirements

This example uses Junos OS Release 12.2. BGP route flap damping support for MBGP MVPN, specifically, and on an address family basis, in general, is introduced in Junos OS Release 12.2.


Overview

BGP route flap damping helps to diminish route instability caused by routes being repeatedly withdrawn and readvertised when a link is intermittently failing.

This example uses the default damping parameters and demonstrates an MBGP MVPN scenario with three provider edge (PE) routing devices, three customer edge (CE) routing devices, and one provider (P) routing device.

[Figure 109 on page 3466](#) shows the topology used in this example.

Figure 109: MBGP MVPN with BGP Route Flap Damping



On PE Device R4, BGP route flap damping is configured for address family **inet-mvpn**. A routing policy called **dampPolicy** uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5. All other MVPN route types are not damped.

This example shows the full configuration on all devices in the [“CLI Quick Configuration” on page 3467](#) section. The [“Configuring Device R4” on page 3470](#) section shows the step-by-step configuration for PE Device R4.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 1 family mpls
set interfaces lo0 unit 1 family inet address 1.1.1.1/32
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.1

```

Device R2

```

set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30
set interfaces ge-1/2/0 unit 2 family mpls
set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 5 family mpls
set interfaces vt-1/2/0 unit 2 family inet
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set interfaces lo0 unit 102 family inet address 100.1.1.2/32
set protocols mpls interface ge-1/2/1.5
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
set protocols ldp interface ge-1/2/1.5
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-1/2/0.2
set routing-instances vpn-1 interface vt-1/2/0.2
set routing-instances vpn-1 interface lo0.102
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 1001

```

Device R3

```
set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 9 family mpls
set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9
set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 1.1.1.3
```

Device R4

```
set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set interfaces lo0 unit 104 family inet address 100.1.1.4/32
set protocols rsvp interface all aggregate
set protocols mpls interface all
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling damping
set protocols bgp group ibgp neighbor 1.1.1.2 import dampPolicy
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement dampPolicy term term1 from family inet-mvpn
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 3
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 4
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 5
set policy-options policy-statement dampPolicy term term1 then accept
set policy-options policy-statement dampPolicy then damping no-damp
set policy-options policy-statement dampPolicy then accept
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set policy-options damping no-damp disable
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
```



```

set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 1001

```

Device R5

```

set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 1.1.1.5/32
set interfaces lo0 unit 105 family inet address 100.1.1.5/32
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 1001

```

Device R6

```

set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls
set interfaces lo0 unit 6 family inet address 1.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 100.1.1.2

```

```
set protocols pim interface all
set routing-options router-id 1.1.1.6
```

Device R7

```
set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 1.1.1.7/32
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.7
```

Configuring Device R4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```
[edit interfaces]
user@R4# set ge-1/2/0 unit 10 family inet address 10.1.1.10/30
user@R4# set ge-1/2/0 unit 10 family mpls
user@R4# set ge-1/2/1 unit 17 family inet address 10.1.1.17/30
user@R4# set ge-1/2/1 unit 17 family mpls
user@R4# set vt-1/2/0 unit 4 family inet
user@R4# set lo0 unit 4 family inet address 1.1.1.4/32
user@R4# set lo0 unit 104 family inet address 100.1.1.4/32
```

2. Configure MPLS and the signaling protocols on the interfaces.

```
[edit protocols]
user@R4# set mpls interface all
user@R4# set mpls interface ge-1/2/0.10
user@R4# set rsvp interface all aggregate
user@R4# set ldp interface ge-1/2/0.10
user@R4# set ldp p2mp
```

3. Configure BGP.

The BGP configuration enables BGP route flap damping for the **inet-mvpn** address family. The BGP configuration also imports into the routing table the routing policy called **dampPolicy**. This policy is applied to neighbor PE Device R2.

```
[edit protocols bgp group ibgp]
user@R4# set type internal
user@R4# set local-address 1.1.1.4
user@R4# set family inet-vpn unicast
user@R4# set family inet-vpn any
user@R4# set family inet-mvpn signaling damping
user@R4# set neighbor 1.1.1.2 import dampPolicy
user@R4# set neighbor 1.1.1.5
```

4. Configure an interior gateway protocol.

```
[edit protocols ospf]
```

```

user@R4# set traffic-engineering
[edit protocols ospf area 0.0.0.0]
user@R4# set interface all
user@R4# set interface lo0.4 passive
user@R4# set interface ge-1/2/0.10

```

5. Configure a damping policy that uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5.

```

[edit policy-options policy-statement dampPolicy term term1]
user@R4# set from family inet-mvpn
user@R4# set from nlri-route-type 3
user@R4# set from nlri-route-type 4
user@R4# set from nlri-route-type 5
user@R4# set then accept

```

6. Configure the **damping** policy to disable BGP route flap damping.

The **no-damp** policy (**damping no-damp disable**) causes any damping state that is present in the routing table to be deleted. The **then damping no-damp** statement applies the **no-damp** policy as an action and has no **from** match conditions. Therefore, all routes that are not matched by **term1** are matched by this term, with the result that all other MVPN route types are not damped.

```

[edit policy-options policy-statement dampPolicy]
user@R4# set then damping no-damp
user@R4# set then accept
[edit policy-options]
user@R4# set damping no-damp disable

```

7. Configure the **parent_vpn_routes** to accept all other BGP routes that are not from the **inet-mvpn** address family.

This policy is applied as an OSPF export policy in the routing instance.

```

[edit policy-options policy-statement parent_vpn_routes]
user@R4# set from protocol bgp
user@R4# set then accept

```

8. Configure the VPN routing and forwarding (VRF) instance.

```

[edit routing-instances vpn-1]
user@R4# set instance-type vrf
user@R4# set interface vt-1/2/0.4
user@R4# set interface ge-1/2/1.17
user@R4# set interface lo0.104
user@R4# set route-distinguisher 100:100
user@R4# set vrf-target target:1:1
user@R4# set protocols ospf export parent_vpn_routes
user@R4# set protocols ospf area 0.0.0.0 interface lo0.104 passive
user@R4# set protocols ospf area 0.0.0.0 interface ge-1/2/1.17
user@R4# set protocols pim rp static address 100.1.1.2
user@R4# set protocols pim interface ge-1/2/1.17 mode sparse
user@R4# set protocols mvpn

```

9. Configure the router ID and the autonomous system (AS) number.

```

[edit routing-instances vpn-A]
user@R4# set routing-options router-id 1.1.1.4

```

```
user@R4# set routing-options autonomous-system 1001
```

10. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
ge-1/2/0 {
  unit 10 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 17 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 4 {
    family inet;
  }
}
lo0 {
  unit 4 {
    family inet {
      address 1.1.1.4/32;
    }
  }
  unit 104 {
    family inet {
      address 100.1.1.4/32;
    }
  }
}

user@R4# show policy-options
policy-statement dampPolicy {
  term term1 {
    from {
      family inet-mvpn;
      nlri-route-type [ 3 4 5 ];
    }
    then accept;
  }
}
```

```

    }
    then {
        damping no-damp;
        accept;
    }
}
policy-statement parent_vpn_routes {
    from protocol bgp;
    then accept;
}
damping no-damp {
    disable;
}

user@R4# show protocols
rsvp {
    interface all {
        aggregate;
    }
}
mpls {
    interface all;
    interface ge-1/2/0.10;
}
bgp {
    group ibgp {
        type internal;
        local-address 1.1.1.4;
        family inet-vpn {
            unicast;
            any;
        }
        family inet-mvpn {
            signaling {
                damping;
            }
        }
        neighbor 1.1.1.2 {
            import dampPolicy;
        }
        neighbor 1.1.1.5;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface lo0.4 {
            passive;
        }
        interface ge-1/2/0.10;
    }
}
ldp {
    interface ge-1/2/0.10;
    p2mp;
}

```

```
}  
user@R4# show routing-instances  
vpn-1 {  
  instance-type vrf;  
  interface vt-1/2/0.4;  
  interface ge-1/2/1.17;  
  interface lo0.104;  
  route-distinguisher 100:100;  
  vrf-target target:1:1;  
  protocols {  
    ospf {  
      export parent_vpn_routes;  
      area 0.0.0.0 {  
        interface lo0.104 {  
          passive;  
        }  
        interface ge-1/2/1.17;  
      }  
    }  
    pim {  
      rp {  
        static {  
          address 100.1.1.2;  
        }  
      }  
      interface ge-1/2/1.17 {  
        mode sparse;  
      }  
    }  
    mvpn;  
  }  
}  
  
user@R4# show routing-opts  
router-id 1.1.1.4;  
autonomous-system 1001;
```

Verification

Confirm that the configuration is working properly.

- [Verifying That Route Flap Damping Is Disabled on page 3474](#)
- [Verifying Route Flap Damping on page 3475](#)

Verifying That Route Flap Damping Is Disabled

Purpose Verify the presence of the **no-damp** policy, which disables damping for MVPN route types other than 3, 4, and 5.

Action From operational mode, enter the **show policy damping** command.

```
user@R4> show policy damping  
Default damping information:  
  Halflife: 15 minutes  
  Reuse merit: 750 Suppress/cutoff merit: 3000  
  Maximum suppress time: 60 minutes
```

```

Computed values:
  Merit ceiling: 12110
  Maximum decay: 6193
Damping information for "no-damp":
Damping disabled

```

Meaning The output shows that the default damping parameters are in effect and that the **no-damp** policy is also in effect for the specified route types.

Verifying Route Flap Damping

Purpose Check whether BGP routes have been damped.

Action From operational mode, enter the **show bgp summary** command.

```

user@R4> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0
          6          6          0          0          0          0
bgp.13vpn.2
          0          0          0          0          0          0
bgp.mvpn.0
          2          2          0          0          0          0
Peer      AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2    1001    3159     3155      0        0    23:43:47
Establ
  bgp.13vpn.0: 3/3/3/0
  bgp.13vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
1.1.1.5    1001    3157     3154      0        0    23:43:40
Establ
  bgp.13vpn.0: 3/3/3/0
  bgp.13vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0

```

Meaning When **Damp State** field shows that zero routes in the bgp.mvpn.0 routing table have been damped. Further down, the last number in the State field shows that zero routes have been damped for BGP peer 1.1.1.2.

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

BGP Monitoring Configuration

- [Example: Configuring BGP Trace Operations on page 3476](#)

Example: Configuring BGP Trace Operations

- [Understanding Trace Operations for BGP Protocol Traffic on page 3476](#)
- [Example: Viewing BGP Trace Files on Logical Systems on page 3477](#)

Understanding Trace Operations for BGP Protocol Traffic

You can trace various BGP protocol traffic to help you debug BGP protocol issues. To trace BGP protocol traffic, include the **traceoptions** statement at the **[edit protocols bgp]** hierarchy level. For routing instances, include the **traceoptions** statement at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

```
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

You can specify the following BGP protocol-specific trace options using the **flag** statement:

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages.
- **nsr-synchronization**—Nonstop active routing synchronization events.
- **open**—BGP open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—BGP update packets. These packets provide routing updates to BGP systems.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the BGP protocol using the **traceoptions flag** statement included at the **[edit protocols bgp]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions

- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information.
- **filter**—Filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Packets being received.
- **send**—Packets being transmitted.



NOTE: Use the **all** trace flag and the **detail** flag modifier with caution because these might cause the CPU to become very busy.



NOTE: If you only enable the **update** flag, received keepalive messages do not generate a trace message.

You can filter trace statements and display only the statement information that passes through the filter by specifying the **filter** flag modifier. The **filter** modifier is only supported for the **route** and **damping** tracing flags.

The **match-on** statement specifies filter matches based on prefixes. It is used to match on route filters.



NOTE: Per-neighbor trace filtering is not supported on a BGP per-neighbor level for **route** and **damping** flags. Trace option filtering support is on a peer group level.

Example: Viewing BGP Trace Files on Logical Systems

This example shows how to list and view files that are stored on a logical system.

- [Requirements on page 3477](#)
- [Overview on page 3478](#)
- [Configuration on page 3478](#)
- [Verification on page 3482](#)

Requirements

- You must have the **view** privilege for the logical system.
- Configure a network, such as the BGP network shown in “[Example: Configuring Internal BGP Peering Sessions on Logical Systems](#)” on page 3184.

Overview

Logical systems have their individual directory structure created in the `/var/logical-systems/logical-system-name` directory. It contains the following subdirectories:

- `/config`—Contains the active configuration specific to the logical system.
- `/log`—Contains system log and tracing files specific to the logical system.

To maintain backward compatibility for the log files with previous versions of Junos OS, a symbolic link (symlink) from the `/var/logs/logical-system-name` directory to the `/var/logical-systems/logical-system-name` directory is created when a logical system is configured.

- `/tmp`—Contains temporary files specific to the logical system.

The file system for each logical system enables logical system users to view trace logs and modify logical system files. Logical system administrators have full access to view and modify all files specific to the logical system.

Logical system users and administrators can save and load configuration files at the logical-system level using the **save** and **load** configuration mode commands. In addition, they can also issue the **show log**, **monitor**, and **file** operational mode commands at the logical-system level.

This example shows how to configure and view a BGP trace file on a logical system. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.



TIP: To view a list of hierarchy levels that support tracing operations, enter the **help apropos traceoptions** command in configuration mode.

Configuration

- [Configuring Trace Operations on page 3479](#)
- [Viewing the Trace File on page 3479](#)
- [Deactivating and Reactivating Trace Logging on page 3481](#)
- [Results on page 3482](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-log
set logical-systems A protocols bgp group internal-peers traceoptions file size 10k
set logical-systems A protocols bgp group internal-peers traceoptions file files 2
set logical-systems A protocols bgp group internal-peers traceoptions flag update detail
```

Configuring Trace Operations

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the trace operations:

1. Configure trace operations on the logical system.

```
[edit logical-systems A protocols bgp group internal-peers]
user@host# set traceoptions file bgp-log
user@host# set traceoptions file size 10k
user@host# set traceoptions file files 2
user@host# set traceoptions flag update detail
```

2. If you are done configuring the device, commit the configuration.

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

Viewing the Trace File

Step-by-Step Procedure To view the trace file:

1. In operational mode on the main router, list the directories on the logical system.

```
user@host> file list /var/logical-systems/A
/var/logical-systems/A:
config/
log/
tmp/
```

2. In operational mode on the main router, list the log files on the logical system.

```
user@host> file list /var/logical-systems/A/log/
/var/logical-systems/A/log:
bgp-log
```

3. View the contents of the **bgp-log** file.

```
user@host> file show /var/logical-systems/A/log/bgp-log
Aug 10 17:12:01 trace_on: Tracing to "/var/log/A/bgp-log" started
Aug 10 17:14:22.826182 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.826445 bgp_send: sending 21 bytes to 192.163.6.4 (Internal
AS 17)
Aug 10 17:14:22.826499
Aug 10 17:14:22.826499 BGP SEND 192.168.6.5+64965 -> 192.163.6.4+179
Aug 10 17:14:22.826559 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.826598 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
Aug 10 17:14:22.831756 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.168.40.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.831851 bgp_send: sending 21 bytes to 192.168.40.4 (Internal
AS 17)
Aug 10 17:14:22.831901
Aug 10 17:14:22.831901 BGP SEND 192.168.6.5+53889 -> 192.168.40.4+179
```

```
Aug 10 17:14:22.831959 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.831999 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
...
```

4. Filter the output of the log file.

```
user@host> file show /var/logical-systems/A/log/bgp-log | match "flags 0x40"
Aug 10 17:14:54.867460 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.867595 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.867650 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.867692 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.884529 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.884581 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.884628 BGP RECV flags 0x40 code NextHop(3): 192.168.6.4
Aug 10 17:14:54.884667 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.911377 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.911422 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.911466 BGP RECV flags 0x40 code NextHop(3): 192.168.40.4
Aug 10 17:14:54.911507 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.916008 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.916054 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.916100 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.916143 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.920304 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.920348 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.920393 BGP RECV flags 0x40 code NextHop(3): 10.0.0.10
Aug 10 17:14:54.920434 BGP RECV flags 0x40 code LocalPref(5): 100
```

5. View the tracing operations in real time.

```
user@host> clear bgp neighbor logical-system A
Cleared 2 connections
```



CAUTION: Clearing the BGP neighbor table is disruptive in a production environment.

6. Run the **monitor start** command with an optional **match** condition.

```
user@host> monitor start A/bgp-log | match 0.0.0.0/0
Aug 10 19:21:40.773467 BGP RECV          0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlr: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlr: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlr: 0.0.0.0/0 qualified bnp->ribact 0x0
12afcb 0x0
```

7. Pause the **monitor** command by pressing Esc-Q.

To unpause the output, press Esc-Q again.

8. Halt the **monitor** command by pressing Enter and typing **monitor stop**.

[Enter]

```
user@host> monitor stop
```

9. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
```

```
user@host:A# deactivate traceoptions
```

```
user@host:A# commit
```

When configuration is deactivated, it appears in the configuration with the **inactive** tag. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# show

type internal;
inactive: traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
    flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

10. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

Deactivating and Reactivating Trace Logging

Step-by-Step Procedure

To deactivate and reactivate the trace file:

1. When you are finished troubleshooting, consider deactivating trace logging to avoid an unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
user@host:A# deactivate traceoptions
user@host:A# commit
```

When configuration is deactivated, the statement appears in the configuration with the **inactive** tag.

```
[edit protocols bgp group internal-peers]
user@host:A# show

type internal;
inactive: traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
    flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

2. To reactivate logging, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

Results

From configuration mode, confirm your configuration by entering the **show logical-systems A protocols bgp group internal-peers** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems A protocols bgp group internal-peers
traceoptions {
  file bgp-log size 10k files 2;
  flag update detail;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Trace Log File Is Operating

Purpose Make sure that events are being written to the log file.

Action user@host:A> **show log bgp-log**
Aug 12 11:20:57 trace_on: Tracing to "/var/log/A/bgp-log" started

Related Documentation

- [Understanding External BGP Peering Sessions on page 3149](#)
- [BGP Configuration Overview](#)

Configuration Statements

- [accept-remote-nexthop on page 3485](#)
- [advertise-external on page 3486](#)
- [advertise-inactive on page 3487](#)
- [advertise-peer-as on page 3488](#)
- [algorithm \(BGP BFD Authentication\) on page 3489](#)
- [apply-groups on page 3491](#)
- [apply-groups-except on page 3491](#)
- [authentication \(BGP BFD Liveness Detection\) on page 3492](#)
- [authentication-algorithm on page 3494](#)
- [authentication-key \(Protocols BGP\) on page 3495](#)
- [authentication-key-chain \(Protocols BGP\) on page 3496](#)
- [bfd-liveness-detection \(Protocols BGP\) on page 3497](#)
- [bgp on page 3501](#)
- [bgp-orf-cisco-mode on page 3502](#)
- [cluster on page 3504](#)
- [damping \(Protocols BGP\) on page 3506](#)


- [description \(Protocols BGP\) on page 3508](#)
- [detection-time \(BFD for BGP\) on page 3509](#)
- [disable \(Protocols BGP\) on page 3510](#)
- [disable \(BGP Graceful Restart\) on page 3511](#)
- [export \(Protocols BGP\) on page 3512](#)
- [family \(Protocols BGP\) on page 3513](#)
- [graceful-restart \(Protocols BGP\) on page 3517](#)
- [group \(Protocols BGP\) on page 3518](#)
- [hold-down-interval \(BGP BFD Liveness Detection\) on page 3521](#)
- [hold-time \(Protocols BGP\) on page 3523](#)
- [import \(Protocols BGP\) on page 3525](#)
- [include-mp-next-hop on page 3526](#)
- [keep on page 3527](#)
- [key-chain \(BGP BFD Authentication\) on page 3529](#)
- [local-address \(Protocols BGP\) on page 3531](#)
- [local-as on page 3533](#)
- [local-preference on page 3536](#)
- [log-updown \(Protocols BGP\) on page 3537](#)
- [loops on page 3538](#)
- [loose-check \(BGP BFD Authentication\) on page 3540](#)
- [metric-out \(Protocols BGP\) on page 3542](#)
- [minimum-interval \(BGP BFD Liveness Detection\) on page 3544](#)
- [minimum-interval \(BGP BFD Transmit Interval\) on page 3545](#)
- [minimum-receive-interval \(BGP BFD Liveness Detection\) on page 3547](#)
- [mtu-discovery on page 3548](#)
- [multihop on page 3550](#)
- [multiplier \(BGP BFD Liveness Detection\) on page 3552](#)
- [neighbor \(Protocols BGP\) on page 3553](#)
- [no-adaptation \(BGP BFD Liveness Detection\) on page 3556](#)
- [no-advertise-peer-as on page 3557](#)
- [no-aggregator-id on page 3558](#)
- [no-client-reflect on page 3559](#)
- [out-delay on page 3560](#)
- [outbound-route-filter on page 3561](#)
- [passive \(Protocols BGP\) on page 3562](#)
- [path-selection on page 3563](#)
- [peer-as \(Protocols BGP\) on page 3565](#)

- [preference \(Protocols BGP\) on page 3567](#)
- [remove-private on page 3568](#)
- [restart-time \(BGP Graceful Restart\) on page 3570](#)
- [session-mode on page 3571](#)
- [stale-routes-time on page 3572](#)
- [tcp-mss \(Protocols BGP\) on page 3573](#)
- [threshold \(BGP BFD Detection Time\) on page 3574](#)
- [threshold \(BGP BFD Transmit Interval\) on page 3576](#)
- [traceoptions \(Protocols BGP\) on page 3577](#)
- [transmit-interval \(BGP BFD Liveness Detection\) on page 3580](#)
- [version \(BGP BFD Liveness Detection\) on page 3581](#)

accept-remote-nexthop

Syntax	accept-remote-nexthop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure multihop and accept-remote-nexthop statements for the same EBGP peer.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops on page 3367 • Understanding Route Advertisement on page 3309 • <i>multipath</i>

advertise-external

Syntax	<code>advertise-external {conditional};</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route. In order to configure the advertise-external statement on a route reflector, you must disable intracluster reflection with the no-client-reflect statement. The advertise-external statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p>
<div>  <p>NOTE: When configuring the advertise-external statement for an AS confederation, it is recommended that EBGp peers belonging to different autonomous systems are configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> </div>	
Options	<p>conditional—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The conditional option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the conditional option is configured.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Advertisement on page 3309 • Understanding Route Advertisement on page 3309

- [advertise-inactive on page 3487](#)

advertise-inactive

Syntax	advertise-inactive;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the advertise-inactive statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The advertise-inactive statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the advertise-inactive statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Preference Value for BGP Routes on page 3327 • Example: Configuring BGP Route Preference (Administrative Distance) on page 3326 • Understanding Route Advertisement on page 3309 • advertise-external on page 3486

advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable the default behavior of suppressing AS routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BGP Route Advertisement on page 3309• Understanding Route Advertisement on page 3309• no-advertise-peer-as on page 3557

algorithm (BGP BFD Authentication)

Syntax	<code>algorithm <i>algorithm-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the algorithm used to authenticate the specified BFD session.
Options	<p><i>algorithm-name</i>—Authentication algorithm name: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>simple-password—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.</p> <p>keyed-md5—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.</p>

meticulous-keyed-md5—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method can take additional time to authenticate the session.

keyed-sha-1—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.

meticulous-keyed-sha-1—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method can take additional time to authenticate the session.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD Authentication for Static Routes on page 2848• Example: Configuring BGP Route Authentication on page 3429• Example: Configuring EBGp Multihop Sessions on page 3317• Understanding Route Authentication on page 3429• authentication on page 3492• bfd-liveness-detection on page 3497• key-chain on page 3529• loose-check on page 3540
------------------------------	--

apply-groups

Syntax	<code>apply-groups [<i>group-names</i>];</code>
Hierarchy Level	All hierarchy levels
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.</p> <p>You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.</p>
Options	<i>group-names</i> —One or more names specified in the groups statement.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none"> • <i>Applying a Junos Configuration Group</i> • <i>groups</i>

apply-groups-except

Syntax	<code>apply-groups-except [<i>group-names</i>];</code>
Hierarchy Level	All hierarchy levels except the top level
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable inheritance of a configuration group.
Options	<i>group-names</i> —One or more names specified in the groups statement.
Required Privilege Level	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
Related Documentation	<ul style="list-style-type: none"> • <i>groups</i> • <i>Disabling Inheritance of a Junos OS Configuration Group</i>

authentication (BGP BFD Liveness Detection)

Syntax	<pre>authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check ; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit protocols bgp bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Specify the router and route authentication to mitigate the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, the receiving router accepts the route.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• Example: Configuring BFD Authentication for Static Routes on page 2848• Example: Configuring BGP Route Authentication on page 3429• Understanding Route Authentication on page 3429

- [algorithm on page 3489](#)
- [bfd-liveness-detection on page 3497](#)
- [key-chain on page 3529](#)
- [loose-check on page 3540](#)

authentication-algorithm

Syntax	<code>authentication-algorithm <i>algorithm</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp <i>group</i> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>ldp session <i>session-address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit protocols ldp session <i>session-address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i></code> <code>neighbor <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced for BGP in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure an authentication algorithm type.
Options	<i>algorithm</i> —Specify one of the following types of authentication algorithms: <ul style="list-style-type: none">• aes-128-cmac-96—Cipher-based message authentication code (AES128, 96 bits).• hmac-sha-1-96—Hash-based message authentication code (SHA1, 96 bits).• md5—Message digest 5.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Route Authentication on page 3429• Example: Configuring Route Authentication for BGP on page 3430

authentication-key (Protocols BGP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp group <i>group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.
Options	<i>key</i> —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Authentication for BGP on page 3430

authentication-key-chain (Protocols BGP)

Syntax	authentication-key-chain <i>key-chain</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for BGP, you cannot commit the 0.0.0.0/allow statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Route Authentication for BGP on page 3430• Example: Configuring BFD Authentication for Static Routes on page 2848• Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

bfd-liveness-detection (Protocols BGP)

Syntax

```

bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    hold-down-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    session-mode (automatic | multihop | single-hop);
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor
address]

```

Release Information

Statement introduced in Junos OS Release 8.1.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2

Support for logical routers introduced in Junos OS Release 8.3.

Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.

holddown-interval statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the **[edit protocols bgp group *group-name* neighbor *address*]** hierarchy level.

no-adaptation statement introduced in Junos OS Release 9.0.

Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Range: 0 through 255,000

Default: 0

minimum-interval *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

Range: 1 through 255,000

minimum-receive-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier *number* (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version (Optional)—Configure the BFD version to detect.

Range: 1 or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• Example: Configuring BFD Authentication for Static Routes on page 2848• Example: Configuring BFD on Internal BGP Peer Sessions on page 3349• Example: Configuring BFD Authentication for BGP on page 3359• Understanding BFD for BGP on page 3348
------------------------------	--

bgp

Syntax	<code>bgp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable BGP on the routing device or for a routing instance.
Default	BGP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<p> NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>
Default	Disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3313](#)

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.



CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



NOTE: If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options	<i>cluster-identifier</i> —4-byte number (such as an IPv4 address).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BGP Route Reflectors on page 3405• Understanding External BGP Peering Sessions on page 3149• no-client-reflect on page 3559

damping (Protocols BGP)

Syntax	damping;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>
Description	<p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default Flap damping is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Examples: Configuring BGP Flap Damping on page 3456](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3466](#)

description (Protocols BGP)

Syntax	<code>description text-description;</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name protocols bgp],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp group group-name],</code> <code>[edit protocols bgp group group-name neighbor address],</code> <code>[edit routing-instances routing-instance-name protocols bgp],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the show command and has no effect on the configuration.
Options	<i>text-description</i> —Text description of the configuration. It is limited to 255 characters.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

detection-time (BFD for BGP)

Syntax	<pre> detection-time { threshold milliseconds; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the clear bfd adaptation command to return BFD interval timers to their configured values. The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for BGP on page 3348 • bfd-liveness-detection on page 3497

- [threshold on page 3574](#)

disable (Protocols BGP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable BGP on the system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

disable (BGP Graceful Restart)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart], [edit protocols bgp graceful-restart], [edit protocols bgp group <i>group-name</i> graceful-restart], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.



NOTE: When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2285 • Configuring Graceful Restart for QFabric Systems on page 1375 • graceful-restart on page 1388

export (Protocols BGP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to routes being exported from the routing table into BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Route Advertisement on page 3309• Routing Policy Configuration Guide• import on page 3525

family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
      }
      add-path {
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
        receive;
      }
      aigp [disable];
      damping;
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>inet-mvpn and inet6-mvpn statements introduced in Junos OS Release 8.4.</p> <p>inet-mdt statement introduced in Junos OS Release 9.4.</p> <p>Support for the loops statement introduced in Junos OS Release 9.6.</p>
Description	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>


- Options**
- any**—Configure the family type to be both unicast and multicast.
 - inet**—Configure NLRI parameters for IPv4.
 - inet6**—Configure NLRI parameters for IPv6.
 - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
 - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
 - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
 - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
 - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
 - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
 - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
 - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.

The remaining statements are explained separately.

- Required Privilege Level**
- routing—To view this statement in the configuration.
 - routing-control—To add this statement to the configuration.

- Related Documentation**
- [autonomous-system on page 2862](#)
 - [local-as on page 3533](#)
 - *Understanding Multiprotocol BGP*

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time seconds; stale-routes-time seconds; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2285 • Configuring Graceful Restart for QFabric Systems on page 1375 • <i>Junos OS High Availability Configuration Guide</i>

group (Protocols BGP)

```
Syntax  group group-name {  
    advertise-inactive;  
    allow [ network/mask-length ];  
    authentication-key key;  
    cluster cluster-identifier;  
    damping;  
    description text-description;  
    export [ policy-names ];  
    family {  
        (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {  
            (any | multicast | unicast | signaling) {  
                accepted-prefix-limit {  
                    maximum number;  
                    teardown <percentage> <idle-timeout (forever | minutes)>;  
                }  
            }  
            add-path {  
                send {  
                    path-count number;  
                    prefix-policy [ policy-names ];  
                }  
                receive;  
            }  
            aigp [disable];  
            damping;  
            prefix-limit {  
                maximum number;  
                teardown <percentage> <idle-timeout (forever | minutes)>;  
            }  
            rib-group group-name;  
            topology name {  
                community {  
                    target identifier;  
                }  
            }  
        }  
    }  
    flow {  
        no-validate policy-name;  
    }  
    labeled-unicast {  
        accepted-prefix-limit {  
            maximum number;  
            teardown <percentage> <idle-timeout (forever | minutes)>;  
        }  
        explicit-null {  
            connected-only;  
        }  
        prefix-limit {  
            maximum number;  
            teardown <percentage> <idle-timeout (forever | minutes)>;  
        }  
        resolve-vpn;  
        rib inet.3;  
    }  
}
```

```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
mvpn-iana-rt-import;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
bgp],
[edit protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp]

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the group statement.</p> <p>The group statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p>
Options	<p>group-name—Name of the BGP group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

hold-down-interval (BGP BFD Liveness Detection)

Syntax	<code>holddown-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.</p> <p>When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes the BGP session is down and does not send a state change notification. The holddown-interval statement is supported only for EBGp peers at the [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>] hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the local-address statement at the [edit protocols bgp group <i>group-name</i>] hierarchy level.</p>
Options	<p><i>milliseconds</i>—Specify the hold-down interval value.</p> <p>Range: 0 through 255,000</p> <p>Default: 0</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BFD for Static Routes on page 2835](#)
 - [bfd-liveness-detection on page 3497](#)

hold-time (Protocols BGP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p>
Options	<p>seconds—Hold time.</p> <p>Range: 10 through 65,535 seconds (or 0 for infinite hold time)</p> <p>Default: 90 seconds</p>



TIP: When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Messages Overview on page 3147• <i>precision-timers</i>

import (Protocols BGP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Interactions with IGP's on page 3305 • Understanding Route Advertisement on page 3309 • Importing and Exporting Routes • Routing Policy Configuration Guide • export on page 3512

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable multiprotocol updates to contain next-hop reachability information.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Configuring Multiprotocol BGP</i>

keep

Syntax	keep (all none);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Control whether or not Junos OS keeps in memory and hides certain routes.</p> <p>If the keep none statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The keep none statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the keep none statement has no effect on this route. This rejected route is retained in memory and hidden even though keep none is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.</p> <p>The routing table can retain the route information learned from BGP in one of the following ways:</p> <ul style="list-style-type: none"> • Default (omit the keep statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS. • keep all—Keep all route information that was learned from BGP. • keep none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure keep none for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.



CAUTION: When you configure **keep (all | none)**, the associated BGP sessions are restarted.

Default By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the **keep** statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.

Options **all**—Retain all routes.

none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When **keep none** is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding Route Advertisement on page 3309](#)
- [out-delay on page 3560](#)

key-chain (BGP BFD Authentication)

Syntax	<code>key-chain <i>key-chain-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Associate a security key with the specified BFD session using the name of the security keychain. Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as <i>hitless</i> because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.
Options	<i>key-chain-name</i> —Name of the authentication keychain. The keychain name must match one of the keychains configured with the key-chain <i>key-chain-name</i> statement at the [edit security authentication-key-chain] hierarchy level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes on page 2835 • Example: Configuring BFD Authentication for Static Routes on page 2848 • Example: Configuring BFD on Internal BGP Peer Sessions on page 3349 • Example: Configuring BGP Route Authentication on page 3429

- [Example: Configuring EBGp Multihop Sessions on page 3317](#)
- [Understanding Route Authentication on page 3429](#)

local-address (Protocols BGP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The local-address statement enables you to specify the source information in BGP update messages. If you omit the local-address statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p>

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



NOTE: Although a BGP session can be established when only one of the paired routers has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routers for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options **address**—IPv6 or IPv4 address of the local end of the connection.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3184](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3173](#)
- [Understanding Internal BGP Peering Sessions on page 3172](#)
- [router-id on page 2938](#)

local-as

Syntax	<code>local-as <i>autonomous-system</i> <loops <i>number</i>> <private alias> <no-prepend-global-as>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><i>alias</i> option introduced in Junos OS Release 9.5.</p> <p><i>no-prepend-global-as</i> option introduced in Junos OS Release 9.6.</p>
Description	<p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occur, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>



NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGp or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Range: 1 through 10

Default: 1

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring BGP Local AS on page 3246• Example: Configuring a Local AS for EBGp Sessions on page 3251• autonomous-system on page 2862• family on page 3513

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<p>local-preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Local Preference Value for BGP Routes on page 3195 • Understanding Internal BGP Peering Sessions on page 3172

- [preference on page 3567](#)

log-updown (Protocols BGP)

Syntax	log-updown;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Preventing BGP Session Resets on page 3449 • <i>Junos OS System Basics Configuration Guide</i> • traceoptions on page 3577

loops

Syntax	<code>loops <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> local-as],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> local-as]</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp local-as],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options autonomous-system <i>as-number</i>],</p> <p>[edit protocols bgp family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> local-as],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> local-as]</p> <p>[edit protocols bgp local-as],</p> <p>[edit routing-options autonomous-system <i>as-number</i>]</p>
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Globally, for the local-AS BGP attribute, or the specified address family, allow the local device's AS number to be in the received AS paths, and specify the number of times detection of the local device's AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure loops 1, the route is hidden if the local device's AS number is detected in the path one or more times. This prevents routing loops and is the default behavior. If you configure loops 2, the route is hidden if the local device's AS number is detected in the path two or more times.</p> <p>Some examples of BGP address families are as follows:</p> <ul style="list-style-type: none"> • inet unicast • inet-vpn multicast • inet6 any • l2vpn auto-discovery-only • ... <p>This list is truncated for brevity. For a complete list of protocol families for which you can specify the loops statement, enter the help apropos loops configuration command at the [edit protocols bgp] hierarchy level on your device.</p> <pre>[edit protocols bgp] user@host# help apropos loops set family inet unicast loops Allow local AS in received AS paths set family inet unicast loops <loops> AS-Path loop count set family inet multicast loops</pre>

```

    Allow local AS in received AS paths
set family inet multicast loops <loops>
    AS-Path loop count
set family inet flow loops
    Allow local AS in received AS paths
set family inet flow loops <loops>
    AS-Path loop count
set family inet any loops
    Allow local AS in received AS paths
set family inet any loops <loops>
    AS-Path loop count
set family inet labeled-unicast loops
    Allow local AS in received AS paths
...

```



NOTE: When you configure the `loops` statement for a specific BGP address family, that value is used to evaluate the AS path for routes received by a BGP peer for the specified address family, rather than the `loops` value configured for the global AS number with the `loops` statement at the `[edit routing-options autonomous-system as-number]` hierarchy level.

Options *number*—Number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden.

Range: 1 through 10

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [autonomous-system on page 2862](#)
- [family on page 3513](#)
- [local-as on page 3533](#)

loose-check (BGP BFD Authentication)

Syntax	<code>loose-check ;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• Example: Configuring BFD Authentication for Static Routes on page 2848• Example: Configuring BFD on Internal BGP Peer Sessions on page 3349• Example: Configuring BGP Route Authentication on page 3429• Example: Configuring EBGp Multihop Sessions on page 3317

- [Understanding Route Authentication on page 3429](#)

metric-out (Protocols BGP)

Syntax	<code>metric-out (<i>metric</i> minimum-igp <i>offset</i> igp (delay-med-update <i>offset</i>);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option delay-med-update introduced in Junos OS Release 9.0.</p>
Description	<p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the metric option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or minimum-igp statement) by specifying a value for offset. The metric is increased by specifying a positive value for offset, and decreased by specifying a negative value for offset.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the med-igp-update-interval minutes statement at the [edit routing-options] hierarchy level.</p>
Options	<p>delay-med-update—Specify that a BGP group or peer configured with the metric-out igp statement not advertise MED updates unless the current MED value is lower than</p>

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



NOTE: You cannot configure the `delay-med-update` statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3236 • Examples: Configuring BGP MED on page 3208 • Example: Configuring the MED Attribute Directly on page 3210 • Understanding the MED Attribute on page 3208 • med-igp-update-interval on page 2905
------------------------------	---

minimum-interval (BGP BFD Liveness Detection)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</code> <code>[edit protocols bgp bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>address</i> bfd-liveness-detection]</code>
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the transmit-interval minimum-interval and minimum-receive-interval statements.
Options	<i>milliseconds</i> —Specify the minimum interval value for BGP BFD liveliness detection. Range: 1 through 255,000
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• bfd-liveness-detection on page 3497• minimum-receive-interval on page 3547• transmit-interval on page 3580

minimum-interval (BGP BFD Transmit Interval)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the minimum-interval statement.</p>
Options	<p><i>milliseconds</i>—Minimum transmit interval value.</p> <p>Range: 1 through 255,000</p>
<div>  <p>NOTE: The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

**Related
Documentation**

- [Example: Configuring BFD for Static Routes on page 2835](#)
- [bfd-liveness-detection on page 3497](#)
- [minimum-interval on page 3544](#)
- [threshold on page 3576](#)

minimum-receive-interval (BGP BFD Liveness Detection)

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the minimum-interval statement.
Options	<p><i>milliseconds</i>—Specify the minimum receive interval value.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes on page 2835 • bfd-liveness-detection on page 3497 • minimum-interval on page 3544 • transmit-interval on page 3580

mtu-discovery

Syntax	mtu-discovery;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Limiting TCP Segment Size for BGP on page 3444](#)
 - *Configuring the Junos OS for IPv6 Path MTU Discovery*
 - *Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections*

multihop

Syntax multihop {
 no-nexthop-change;
 ttl *ttl-value*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp *group* *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp *group* *group-name*
 neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp *group* *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp *group* *group-name* *neighbor* *address*],
 [edit protocols bgp],
 [edit protocols bgp *group* *group-name*],
 [edit protocols bgp *group* *group-name* *neighbor* *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp *group* *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp *group* *group-name*
 neighbor *address*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure an EBGp multihop session.

For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routers. This allows you to configure one or more routers between the PE and CE routers.

An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.

If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.



NOTE: You cannot configure the `accept-remote-nexthop` statement at the same time.

Default If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring EBGp Multihop Sessions on page 3317• <i>Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs</i>• accept-remote-nextthop on page 3485• <i>no-nextthop-change</i>• <i>tth</i>

multiplier (BGP BFD Liveness Detection)

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i></code> <code> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</code> <code>[edit protocols bgp bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i></code> <code> bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>address</i> bfd-liveness-detection]</code>
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.
Options	<i>number</i> —Number of hello packets. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• bfd-liveness-detection on page 3497

neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```

```
        rib-group group-name;  
        topology name {  
            community {  
                target identifier;  
            }  
        }  
    }  
}  
route-target {  
    advertise-default;  
    external-paths number;  
    accepted-prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
    prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
}  
signaling {  
    prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
}  
}  
graceful-restart {  
    disable;  
    restart-time seconds;  
    stale-routes-time seconds;  
}  
hold-time seconds;  
import [ policy-names ];  
ipsec-sa ipsec-sa;  
keep (all | none);  
local-address address;  
local-as autonomous-system <private>;  
local-interface interface-name;  
local-preference preference;  
log-updown;  
metric-out (metric | minimum-igp <offset> | igp <offset>);  
mtu-discovery;  
multihop <ttl-value>;  
multipath {  
    multiple-as;  
}  
no-aggregator-id;  
no-client-reflect;  
out-delay seconds;  
passive;  
peer-as autonomous-system;  
preference preference;  
tcp-mss segment-size;  
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;
```

```

        flag flag <flag-modifier> <disable>;
    }
    vpn-apply-export;
}

```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple neighbor statements.</p> <p>By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the neighbor statement.</p> <p>The neighbor statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an allow all statement in place of a neighbor statement.)</p>
Options	<p>address—IPv6 or IPv4 address of a single peer.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>BGP Configuration Guide</i>

no-adaptation (BGP BFD Liveness Detection)

Syntax	no-adaptation;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0 Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• bfd-liveness-detection on page 3497

no advertise-peer-as

Syntax	no-advertise-peer-as;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable the default behavior of suppressing AS routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Advertisement on page 3309 • Understanding Route Advertisement on page 3309 • advertise-peer-as on page 3488

no-aggregator-id

Syntax	no-aggregator-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Prevent different routers within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The no-aggregator-id statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>
Default	If you omit this statement, the router ID is included in the BGP aggregator path attribute.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Update Messages on page 3148


no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Reflectors on page 3405 • cluster on page 3504

out-delay

Syntax	<code>out-delay seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Specify how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates. When configured, the out-delay value displays as Outbound Timer when using show bgp group or show bgp group neighbor commands.
Default	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
Options	seconds —Output delay time. Range: 0 through 65,535 seconds Default: 0 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Route Advertisement on page 3309

outbound-route-filter

Syntax	<pre> outbound-route-filter { bgp-orf-cisco-mode; prefix-based { accept { (inet inet6); } } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a BGP peer to accept outbound route filters from a remote peer.
Options	<p>accept—Specify that outbound route filters from a BGP peer be accepted.</p> <p>inet—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p>inet6—Specify that IPv6 prefix-based outbound route filters be accepted.</p>
	<p> NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p>
	<p>prefix-based—Specify that prefix-based filters be accepted.</p> <p>The bgp-orf-cisco-mode statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3313](#)

passive (Protocols BGP)

Syntax	passive;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the router so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3449

path-selection

Syntax	<pre> path-selection { (always-compare-med cisco-non-deterministic external-router-id); as-path-ignore; l2vpn-use-bgp-rules; med-plus-igp { igp-multiplier <i>number</i>; med-multiplier <i>number</i>; } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>med-plus-igp option introduced in Junos OS Release 8.1.</p> <p>as-path-ignore and l2vpn-use-bgp-rules options introduced in Junos OS Release 10.2.</p>
Description	Configure BGP path selection.
Default	If the path-selection statement is not included in the configuration, only the multiple exit discriminators (MEDs) of routes that have the same peer ASs are compared.
Options	always-compare-med —Always compare MEDs whether or not the peer ASs of the compared routes are the same.



NOTE: We recommend that you configure the **always-compare-med** option.

as-path-ignore—In the best-path algorithm, skip the step that compares the autonomous system (AS) path lengths. By default, the best-path algorithm evaluates the length of the AS paths and prefers the route with the shortest AS path length.



NOTE: The **as-path-ignore** statement is not supported with routing instances.

cisco-non-deterministic—Emulate the Cisco IOS default behavior. This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With **cisco-non-deterministic** mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order

in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGp; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGp peer. This allows the routing device to install path 1 as the active path for the route.



NOTE: We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

external-router-id—Compare the router ID between external BGP paths to determine the active path.

igp-multiplier *number*—The multiplier value for the IGP cost to a next-hop address. This option is useful for making the MED and IGP cost comparable.

Range: 1 through 1000

Default: 1

med-multiplier *number*—The multiplier value for the MED calculation. This option is useful for making the MED and IGP cost comparable.

Range: 1 through 1000

Default: 1

med-plus-igp—Add the IGP cost to the indirect next-hop destination to the MED before comparing MED values for path selection. This statement only affects best-path selection. It does not affect the advertised MED.

The other option is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Understanding BGP Path Selection on page 3332• Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3336
------------------------------	---

peer-as (Protocols BGP)

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the peer-as statement must be different from the local router's AS number, which you specify in the autonomous-system statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the autonomous-system and peer-as statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p>

- *Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 2-Byte AS Number* in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.
- *Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 4-Byte AS Number* in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

Options *autonomous-system*—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege routing—To view this statement in the configuration.

Level routing-control—To add this statement to the configuration.

- Related Documentation**
- *4-Byte Autonomous System Numbers Overview* in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)
 - *Juniper Networks Implementation of 4-Byte Autonomous System Numbers* in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

preference (Protocols BGP)

Syntax	<code>preference preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>
Options	<p>preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 170 for the primary preference</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • local-preference on page 3536 • Example: Configuring the Preference Value for BGP Routes on page 3327

remove-private

Syntax remove-private all replace nearest;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp *group* *group-name* **neighbor** *address*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp *group* *group-name* **neighbor** *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp *group* *group-name* **neighbor**
 address]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Options **all**—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.

nearest—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

replace—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Example: Removing Private AS Numbers from AS Paths on page 3343
------------------------------	---

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time seconds;</code>
Hierarchy Level	<code>[edit protocols (bgp rip ripng) graceful-restart]</code> , <code>[edit logical-systems logical-system-name protocols (bgp rip ripng) graceful-restart (Enabling Globally)]</code> , <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp graceful-restart]</code> , <code>[edit routing-instances routing-instance-name protocols bgp graceful-restart]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	seconds —Length of time for the graceful restart period. Range: 1 through 600 seconds Default: Varies by protocol: <ul style="list-style-type: none">• BGP—120 seconds• RIP and RIPng—60 seconds
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 2285• Configuring Graceful Restart Options for RIP and RIPng on page 2289• Configuring Graceful Restart for QFabric Systems on page 1375• stale-routes-time on page 2318

session-mode

Syntax	<code>session-mode (automatic multihop single-hop);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.1.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface. BGP uses multihop BFD sessions if the peer is not directly connected to the router's interface. If the peer session's local-address option is configured, the directly connected check is based partly on the source address that would be used for BGP and BFD.</p> <p>For backward compatibility, you can override the default behavior by configuring the single-hop or multihop option. Before Junos OS Release 11.1, the behavior was to assume that IBGP peer sessions were multihop.</p>
Options	<p>automatic—Configure BGP to use single-hop BFD sessions if the peer is directly connected to the router's interface, and multihop BFD sessions if the peer is not directly connected to the router's interface</p> <p>multihop—Configure BGP to use multihop BFD sessions.</p> <p>single-hop—Configure BGP to use single-hop BFD sessions.</p> <p>Default: automatic</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD Authentication for BGP on page 3359 • Example: Configuring BFD on Internal BGP Peer Sessions on page 3349

- [Example: Configuring BFD Authentication for BGP on page 3359](#)
- [Understanding BFD Authentication for BGP on page 3357](#)


stale-routes-time

Syntax	stale-routes-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• restart-time (BGP Graceful Restart) on page 2317

tcp-mss (Protocols BGP)


Syntax	<code>tcp-mss <i>segment-size</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit protocols bgp], [edit protocol bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors.</p> <p>The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.</p>
Options	<p><i>segment-size</i>—MSS for the TCP connection.</p> <p>Range: 1 through 4096</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Limiting TCP Segment Size for BGP on page 3444

threshold (BGP BFD Detection Time)


Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div>  <p>NOTE: The threshold value must be equal to or greater than the transmit interval.</p> <p>The threshold time must be equal to or greater than the value specified in the <code>minimum-interval</code> or the <code>minimum-receive-interval</code> statement.</p> </div>	
Options	<p><i>milliseconds</i>—Value for the detection time adaptation threshold.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BFD for Static Routes on page 2835](#)

threshold (BGP BFD Transmit Interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
Options	<p><i>milliseconds</i>—Value for the transmit interval adaptation threshold.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p>
<div>  <p>NOTE: The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes on page 2835 • bfd-liveness-detection on page 3497

traceoptions (Protocols BGP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>4byte-as statement introduced in Junos OS Release 9.2.</p> <p>4byte-as statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.</p>
<div>  NOTE: The traceoptions statement is not supported on QFabric systems. </div>	
Default	<p>The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.</p>
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place BGP tracing output in the file bgp-log.</p>

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route**, **damping**, and **update** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	<p>routing and trace—To view this statement in the configuration.</p> <p>routing-control and trace-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • log-updown on page 3537 statement • Understanding Trace Operations for BGP Protocol Traffic on page 3476 • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 3860

transmit-interval (BGP BFD Liveness Detection)

Syntax	<pre>transmit-interval { minimum-interval milliseconds; threshold milliseconds; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Specify the transmit interval for the bfd-liveness-detection statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum time that it requires between packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 2835• bfd-liveness-detection on page 3497• threshold on page 3576• minimum-interval on page 3545• minimum-receive-interval on page 3547

version (BGP BFD Liveness Detection)

Syntax	version (1 automatic);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the BFD version for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version, which is either 0 or 1.
Options	<p>Configure the BFD version to detect: 1 (BFD version 1) or automatic (autodetect the BFD version)</p> <p>Default: automatic</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD Authentication for BGP on page 3359 • Example: Configuring BFD on Internal BGP Peer Sessions on page 3349 • Example: Configuring BFD Authentication for BGP on page 3359 • Understanding BFD Authentication for BGP on page 3357

CHAPTER 38

Administration

- [Routine Monitoring on page 3583](#)
- [Operational Commands on page 3583](#)

Routine Monitoring

- [Monitoring BGP Routing Information on page 3583](#)

Monitoring BGP Routing Information

Purpose Use the monitoring functionality to monitor BGP routing information on the routing device.

Action To view BGP routing information in the CLI, enter the following commands:

- `show bgp summary`
- `show bgp neighbor`

Related Documentation

- [show bgp neighbor on page 2346](#)
- [show bgp summary on page 3610](#)

Operational Commands

- `clear bgp damping`
- `clear bgp neighbor`
- `clear bgp table`
- `show bgp group`
- `show bgp neighbor`
- `show bgp summary`
- `show policy damping`
- `show route damping`

clear bgp damping

List of Syntax	Syntax on page 3584 Syntax (EX Series Switch and QFX Series) on page 3584
Syntax	<code>clear bgp damping</code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><<i>prefix</i>></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp damping</code> <code><<i>prefix</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear BGP route flap damping information.
Options	none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>prefix</i> —(Optional) Clear route flap damping information for only the specified destination prefix.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show policy damping on page 3615• show route damping on page 2997
List of Sample Output	clear bgp damping on page 3584
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

List of Syntax	Syntax on page 3585 Syntax (EX Series Switch and QFX Series) on page 3585
Syntax	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> • Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection. • (soft or soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.
Options	<p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.</p> <p>soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.</p> <p>soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.</p>

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation • [show bgp neighbor on page 2346](#)

List of Sample Output [clear bgp neighbor on page 3586](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear bgp neighbor](#)

```
user@host> clear bgp neighbor
```

clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	<code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>table-name</code> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 3587 clear bgp table inet.6 logical-system all on page 3587 clear bgp table private.inet.6 logical-system ls1 on page 3587 clear bgp table logical-system all inet.0 on page 3587 clear bgp table logical-system ls2 private.inet.0 on page 3588
Output Fields	This command produces no output.

Sample Output

`clear bgp table private.inet.0`

```
user@host> clear bgp table private.inet.0
```

`clear bgp table inet.6 logical-system all`

```
user@host> clear bgp table inet.6 logical-system all
```

`clear bgp table private.inet.6 logical-system ls1`

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

`clear bgp table logical-system all inet.0`

```
user@host> clear bgp table logical-system all inet.0
```

clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```


show bgp group

List of Syntax	Syntax on page 3589 Syntax (EX Series Switch and QFX Series) on page 3589
Syntax	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about the configured BGP groups.
Options	<p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p>
Required Privilege Level	view
List of Sample Output	show bgp group on page 3593 show bgp group brief on page 3593 show bgp group detail on page 3594

[show bgp group rtf detail on page 3595](#)
[show bgp group summary on page 3595](#)

Output Fields [Table 287 on page 3590](#) describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 287: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none
Total peers	Total number of peers in the group.	brief detail none

Table 287: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Established	Number of peers in the group that are in the established state.	All levels
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 287: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table inet.number	Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 287: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2  External: 0  Internal: 2  Down peers: 1  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

bgp.rtarget.0
          2         0         0         0         0         0

```

show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2  External: 0  Internal: 2  Down peers: 1  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

bgp.rtarget.0
          2         0         0         0         0         0

```

show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal   AS: 1                               Local AS: 1
Name: ibgp             Index: 0                             Flags: <Export Eval>
Holdtime: 0
Total peers: 3         Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0

```

```

Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target
    100:100/64
    200:201/64
    Mask
    00000002
    (Group)
    Entries: 2
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target
    200:201/64
    Mask
    (Group)
    Entries: 1

```

show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers      Established      Active/Received/Accepted/Damped
ibgp       Internal  3           0
Groups: 1  Peers: 3      External: 0      Internal: 3      Down peers: 3      Flaps: 3
  bgp.l3vpn.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  bgp.mdt.0        : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.inet.0     : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.mdt.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

show bgp neighbor

List of Syntax	Syntax on page 3596 Syntax (EX Series Switch and QFX Series) on page 3596
Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Configuration Guide</i> .
Required Privilege Level	view

Related Documentation • [clear bgp neighbor on page 3585](#)

List of Sample Output [show bgp neighbor on page 3603](#)
[show bgp neighbor \(CLNS\) on page 3604](#)
[show bgp neighbor \(Layer 2 VPN\) on page 3605](#)
[show bgp neighbor \(Layer 3 VPN\) on page 3607](#)
[show bgp neighbor neighbor-address on page 3607](#)
[show bgp neighbor neighbor-address on page 3608](#)
[show bgp neighbor orf neighbor-address detail on page 3609](#)

Output Fields [Table 224 on page 2347](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 288: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	Current state of the BGP session: <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Local Address	Address of the local routing device.
Remove-private options	Options associated with the <code>remove-private</code> statement.
Holdtime	Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the <code>preference</code> statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <code>out-delay</code> parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .
Table inet.number	Information about the routing table: <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	Information about dropped path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	Information about ignored path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.

Table 288: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)

```

```

Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```


show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000

```

```

RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.13vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 15      Sent 20      Checked 20
  Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
  Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal      State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None

```

```

Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6    Local ID: 10.255.245.5    Active Holdtime: 60000
Keepalive Interval: 20000 Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)

```

```

Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           0
  Received prefixes:         0
  Accepted prefixes:         0
  Suppressed due to damping: 0
  Advertised prefixes:       0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:           1 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:           0 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show bgp summary

List of Syntax	Syntax on page 3610 Syntax (EX Series Switch and QFX Series) on page 3610
Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show bgp summary (When a Peer Is Not Established) on page 3613 show bgp summary (When a Peer Is Established) on page 3613 show bgp summary (CLNS) on page 3613 show bgp summary (Layer 2 VPN) on page 3613 show bgp summary (Layer 3 VPN) on page 3614
Output Fields	<p>Table 289 on page 3610 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 289: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.

Table 289: show bgp summary Output Fields (*continued*)

Field Name	Field Description
Peers	Number of BGP peers.
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 289: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p>

Sample Output

show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002      86      90       0        2      42:54 0/0/0

0/0/0
10.0.0.4        65002      90      91       0        1      42:54 0/2/0

0/0/0
10.0.0.6        65002      87      90       0        3          3 Active
10.1.12.1       65001      89      89       0        1      42:54 4/4/0

0/0/0

```

show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002    88675    88652     0        2      42:38 2/4/0

0/0/0
10.0.0.3        65002    54528    54532     0        1     2w4d22h 0/0/0

0/0/0
10.0.0.4        65002    51597    51584     0        0     2w3d22h 2/2/0

0/0/0

```

show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer           AS      InPkt   OutPkt   OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1    200     1735    1737     0        0    14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0

```

show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.l2vpn.0      1          1          0          0          0          0          0
inet.0           0          0          0          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35   65299      72      74       0        1     19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0

```

```

10.255.245.36 65299      2164      2423      0        4        19:50 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.37 65299      36         37        0        4        17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299     138        168        0        6        53:48 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.69 65299     134        140        0        6        53:42 Establ
  inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State Pending
bgp.13vpn.0      2          2          0          0      0      0      0
Peer          AS      InPkt      OutPkt      OutQ      Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5      2          21         22          0          0      6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15   1          19         21          0          0      6:17 Establ
  bgp.13vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```

show policy damping

List of Syntax	Syntax on page 3615 Syntax (EX Series Switch and QFX Series) on page 3615
Syntax	<pre>show policy damping <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about BGP route flap damping parameters.
Options	<p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the <i>Routing Policy Configuration Guide</i> • clear bgp damping on page 3584 • show route damping on page 2997
List of Sample Output	show policy damping on page 3616
Output Fields	<p>Table 290 on page 3616 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.</p>

Table 290: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```

show route damping

List of Syntax	Syntax on page 3617 Syntax (EX Series Switch and QFX Series) on page 3617
Syntax	show route damping (decayed history suppressed) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show route damping (decayed history suppressed) <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>decayed—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p>history—Display entries that have already been withdrawn, but have been logged.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp damping on page 3584 • show policy damping on page 3615
List of Sample Output	show route damping decayed detail on page 3620 show route damping history on page 3621 show route damping history detail on page 3621
Output Fields	Table 269 on page 2998 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.

Table 291: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0 .	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holdddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
<i>destination-prefix (entry, announced)</i>	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
<i>[protocol, preference]</i>	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive

Table 291: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive

Table 291: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP    Preference: 170/-101
           Next-hop reference count: 151973
           Source: 172.23.2.129
           Next hop: via so-1/2/0.0
           Next hop: via so-5/1/0.0, selected
           Next hop: via so-6/0/0.0
           Protocol next hop: 172.23.2.129
           Indirect next hop: 89a1a00 264185
           State: <Active Ext>
           Local AS: 65000 Peer AS: 65490
           Age: 3:28      Metric2: 0
           Task: BGP_65490.172.23.2.129+179
           Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

  6-Resolve tree 2 7-Resolve tree 3
    AS path: 65490 65520 65525 65525 65525 65525 I ()
    Communities: 65501:390 65501:2000 65501:3000 65504:701
    Localpref: 100
    Router ID: 172.23.2.129
    Merit (last update/now): 1934/1790
    damping-parameters: damping-high

```



```

Last update:      00:03:28 First update:      00:06:40
Flaps: 2

```

show route damping history

```

user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0

```

show route damping history detail

```

user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:      00:01:05 First update:      00:01:05
        Flaps: 1

```


PART 13

Intermediate System to Intermediate System

- [Overview on page 3625](#)
- [Configuration on page 3633](#)
- [Administration on page 3755](#)

CHAPTER 39

Overview

- [IS-IS Overview on page 3625](#)

IS-IS Overview

- [IS-IS Overview on page 3625](#)
- [Understanding BFD Authentication for IS-IS on page 3630](#)
- [Understanding Hitless Authentication Key Rollover for IS-IS on page 3631](#)

IS-IS Overview

The IS-IS protocol is an interior gateway protocol (IGP) that uses link-state information to make routing decisions.

IS-IS is a link-state IGP that uses the shortest-path-first (SPF) algorithm to determine routes. IS-IS evaluates the topology changes and determines whether to perform a full SPF recalculation or a partial route calculation (PRC). This protocol originally was developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets.

Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected. IS-IS uses the SPF algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required.



NOTE: Because IS-IS uses ISO addresses, the configuration of IP version 6 (IPv6) and IP version 4 (IPv4) implementations of IS-IS is identical.

This section discusses the following topics:

- [IS-IS Terminology on page 3626](#)
- [ISO Network Addresses on page 3626](#)
- [IS-IS Packets on page 3628](#)
- [Persistent Route Reachability on page 3629](#)

- [IS-IS Support for Multipoint Network Clouds on page 3629](#)
- [Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels on page 3629](#)

IS-IS Terminology

An IS-IS network is a single autonomous system (AS), also called a *routing domain*, that consists of *end systems* and *intermediate systems*. End systems are network entities that send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring *Level 1* and *Level 2* intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

ISO Network Addresses

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*.

IS-IS supports multiple NSAP addresses on the loopback lo0 interface.

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title (NET)*. Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

```
49.0001.00a0.c96b.c490.00
49.0001.2081.9716.9018.00
```

NETs take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists

of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID
- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain (area) identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next six bytes form the system identifier. The system identifier can be any six bytes that are unique throughout the entire domain. The system identifier commonly is the media access control (MAC) address (as in the first example, 00a0.c96b.c490) or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 2081.9716.9018, which corresponds to IP address 208.197.169.18). The last byte (00) is the n-selector.



NOTE: The system identifier cannot be 0000.0000.0000. All 0s is an illegal setting, and the adjacency is not formed with this setting.

To provide help with IS-IS debugging, the Junos[®] operating system (Junos OS) supports dynamic mapping of ISO system identifiers to the hostname. Each system can be configured with a hostname, which allows the system identifier-to-hostname mapping to be carried in a dynamic hostname type, length, and value (TLV) tuple in IS-IS link-state PDUs. This enables intermediate systems in the routing domain to learn about the ISO system identifier of a particular intermediate system.

IS-IS Packets

Each IS-IS PDU shares a common header. IS-IS uses the following PDUs to exchange protocol information:

- IS-IS hello (IIH) PDUs—Broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

- Link-state PDUs—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.

Also included is metric and IS-IS neighbor information. Each link-state PDU must be refreshed periodically on the network and is acknowledged by information within a sequence number PDU.

On point-to-point links, each link-state PDU is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU (CSNP) is sent out over the network. Any router that finds newer link-state PDU information in the CSNP then purges the out-of-date entry and updates the link-state database.

Link-state PDUs support variable-length subnet mask addressing.

- Complete sequence number PDUs (CSNPs)—Contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.

Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP).

- Partial sequence number PDUs (PSNPs)—Sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively

requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.

A PSNP is used by an IS-IS router to request link-state PDU information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of a link-state PDU on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a device compares a CSNP to its local database and determines that a link-state PDU is missing, the router issues a PSNP for the missing link-state PDU, which is returned in a link-state PDU from the router sending the CSNP. The received link-state PDU is then stored in the local database, and an acknowledgment is sent back to the originating router.

Persistent Route Reachability

IPv4 and IPv6 route reachability information in IS-IS link-state PDUs is preserved when you commit a configuration. IP prefixes are preserved with their original packet fragment upon link-state PDU regeneration.


IS-IS Support for Multipoint Network Clouds

IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.

Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B). See [Figure 110 on page 3629](#).

Figure 110: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2



Related Documentation • *IS-IS Configuration Guide*

Understanding BFD Authentication for IS-IS

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IS-IS. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3630](#)
- [Security Authentication Keychains on page 3631](#)
- [Strict Versus Loose Authentication on page 3631](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords might be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm 1 for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method,

packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.

- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Related Documentation

- [Example: Configuring BFD Authentication for IS-IS on page 3665](#)

Understanding Hitless Authentication Key Rollover for IS-IS

IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in routing. By default, authentication is disabled. The authentication algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

If you configure authentication for all peers, each peer in that group inherits the group's authentication.

You can update authentication keys without resetting any IS-IS neighbor sessions. This is referred to as *hitless authentication key rollover*.

Hitless authentication key rollover uses authentication keychains, which consist of the authentication keys that are being updated. The keychain includes multiple keys. Each key in the keychain has a unique start time. At the next key's start time, a rollover occurs from the current key to the next key, and the next key becomes the current key.

You can choose the algorithm through which authentication is established. You can configure MD5 or SHA-1 authentication. You associate a keychain and the authentication algorithm with an IS-IS neighboring session. Each key contains an identifier and a secret password.

The sending peer chooses the active key based on the system time and the start times of the keys in the keychain. The receiving peer determines the key with which it authenticates based on the incoming key identifier.

You can configure either RFC 5304-based encoding or RFC 5310-based encoding for the IS-IS protocol transmission encoding format.

**Related
Documentation**

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647](#)

CHAPTER 40

Configuration

- [Configuration Guidelines on page 3633](#)
- [Configuration Examples on page 3638](#)
- [Configuration Tasks on page 3690](#)
- [Configuration Statements on page 3692](#)

Configuration Guidelines

- [Example: Configuring IS-IS on page 3633](#)

Example: Configuring IS-IS

This example shows how to configure IS-IS.

- [Requirements on page 3633](#)
- [Overview on page 3633](#)
- [Configuration on page 3634](#)
- [Verification on page 3636](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, you configure the two IS-IS routing devices in a single area. The devices have NET addresses 49.0002.0192.0168.0001.00 and 49.0002.0192.0168.0002.00 on the lo0 interfaces. Additionally, you configure the ISO family on the IS-IS interfaces.

For Junos OS security devices only, you configure the **mode packet-based** statement at the **[edit security forwarding-options family iso]** hierarchy level.

[Figure 111 on page 3634](#) shows the topology used in this example.

Figure 111: Simple IS-IS Topology



“CLI Quick Configuration” on page 3634 shows the configuration for both of the devices in Figure 111 on page 3634. The section “Step-by-Step Procedure” on page 3634 describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R2
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0
```

Device R2

```
set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R1
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS:

1. Enable IS-IS if your router is in secure context.

```
[edit security forwarding-options family iso]
user@R1# set mode packet-based
```
2. Create the interface that connects to Device R2, and configure the ISO family on the interface.

```
[edit interfaces ge-1/2/0 unit 0]
user@R1# set description to-R2
user@R1# set family inet address 10.0.0.1/30
user@R1# set family iso
```

3. Create the loopback interface, set the IP address, and set the NET address.

```
[edit interfaces lo0 unit 0]
user@R1# set family inet address 192.168.0.1/32
user@R1# set family iso address 49.0002.0192.0168.0001.00
```

4. Enable IS-IS on the interfaces.

```
[edit protocols isis]
user@R1# set interface ge-1/2/0.0
user@R1# set interface lo0.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show security
forwarding-options {
  family iso {
    mode packet-based;
  }
}

user@R1# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
    }
  }
}

user@R1# show protocols
isis {
  interface ge-1/2/0.0;
  interface lo0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying IS-IS Interface Configuration on page 3636](#)
- [Verifying IS-IS Interface Configuration in Detail on page 3636](#)
- [Verifying IS-IS Adjacencies on page 3637](#)
- [Verifying IS-IS Adjacencies in Detail on page 3637](#)

Verifying IS-IS Interface Configuration

Purpose Verify the status of the IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface brief** command.

```
user@R1> show isis interface brief
IS-IS interface database:
Interface          L CirID Level 1 DR          Level 2 DR          L1/L2 Metric
lo0.0              0  0x1 Passive                Passive              0/0
ge-1/2/0.0         3  0x1 R2.02                  R2.02                10/10
```

Meaning Verify that the output shows the intended configuration of the interfaces on which IS-IS is enabled.

Verifying IS-IS Interface Configuration in Detail

Purpose Verify the details of IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface detail** command.

```
user@R1> show isis interface detail
IS-IS interface database:
lo0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                0        64      0 Passive
    2                0        64      0 Passive
ge-1/2/0.0
  Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                1        64     10    9.000    27 R2.02 (not us)
    2                1        64     10    9.000    27 R2.02 (not us)
```

Meaning Check the following output fields and verify that the output shows the intended configuration of IS-IS-enabled interfaces:

- Interface—Interface configured for IS-IS.
- State—Internal implementation information.
- Circuit id—Circuit identifier.

- Circuit type—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- link-state PDU interval—Time between IS-IS information messages.
- L or Level—Type of adjacency:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- Adjacencies—Adjacencies established on the interface.
- Priority—Priority value established on the interface.
- Metric—Metric value for the interface.
- Hello(s)—Intervals between hello PDUs.
- Hold(s)—Hold time on the interface.

Verifying IS-IS Adjacencies

Purpose Display brief information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency brief** command.

```
user@R1> show isis adjacency brief
Interface      System      L State      Hold (secs) SNPA
ge-1/2/0.0     R2          1 Up          6 0:5:85:8f:c8:bd
ge-1/2/0.0     R2          2 Up          6 0:5:85:8f:c8:bd
```

Meaning Verify the adjacent routers in the IS-IS database.

Verifying IS-IS Adjacencies in Detail

Purpose Display extensive information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency extensive** command.

```
user@R1> show isis adjacency extensive
R2
Interface: ge-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When              State      Event          Down reason
Thu May 31 11:18:48 Up         Seenself
```

R2

```
Interface: ge-1/2/0.0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When                State      Event      Down reason
Thu May 31 11:18:48 Up         SeenseIf
```

Meaning Check the following fields and verify the adjacency information about IS-IS neighbors:

- Interface—Interface through which the neighbor is reachable.
- L or Level—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2

An exclamation point before the level number indicates that the adjacency is missing an IP address.

- State—Status of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.
- Event—Message that identifies the cause of a state.
- Down reason—Reason the adjacency is down.
- Restart capable—A neighbor is configured for graceful restart.
- Transition log—List of transitions including **When**, **State**, and **Reason**.

**Related
Documentation**

- *Understanding IS-IS Configuration*
- [Example: Configuring Designated Router Election Priority for IS-IS on page 3687](#)
- *Junos OS Feature Support Reference for SRX Series and J Series Devices*

Configuration Examples

- [Example: Configuring Multi-Level IS-IS on page 3639](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 3651](#)
- [Example: Configuring BFD for IS-IS on page 3659](#)
- [Example: Configuring BFD Authentication for IS-IS on page 3665](#)
- [Example: Configuring IS-IS Multicast Topology on page 3668](#)
- [Example: Configuring IS-IS for CLNS on page 3684](#)
- [Example: Configuring IS-IS Designated Routers on page 3686](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces on page 3687](#)

Example: Configuring Multi-Level IS-IS

This example shows how to configure a multi-level IS-IS topology.

- [Requirements on page 3639](#)
- [Overview on page 3639](#)
- [Configuration on page 3640](#)
- [Verification on page 3644](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


Like OSPF, the IS-IS protocol supports the partitioning of a routing domain into multiple areas with levels that control interarea flooding. The use of multiple levels improves protocol scalability, as Level 2 (backbone) link-state PDUs are normally not flooded into a Level 1 area.

An IS-IS Level 2 area is analogous to the OSPF backbone area (0), while a Level 1 area operates much like an OSPF totally stubby area, in that a default route is normally used to reach both inter-level and AS external routes.

Unlike OSPF, IS-IS area boundaries occur between routers, such that a given routing device is always wholly contained within a particular area. Level 1 adjacencies can be formed between routers that share a common area number, while a Level 2 adjacency can be formed between routers that might or might not share an area number.

[Figure 112 on page 3639](#) shows the topology used in this example.

Figure 112: IS-IS Multi-Level Topology



[“CLI Quick Configuration” on page 3640](#) shows the configuration for all of the devices in [Figure 112 on page 3639](#). The section [“Step-by-Step Procedure” on page 3641](#) describes the steps on Device R5.

This example has the following characteristics:

- Device R5 functions as a Level 1/Level 2 router to interconnect the Level 2 backbone area 49.001 and the Level 1 area 49.002 containing Device R6 and Device R7.
- The system ID is based on the devices' IPv4 lo0 addresses.
- Loss of any individual interface does not totally disrupt the IS-IS operation.
- The IPv4 lo0 addresses of all routers are reachable through IS-IS.
- The link between Device R3 and Device S1 appears in area 49.001 as an intra-area route. No IS-IS adjacencies can be established on this interface. This is accomplished by configuring the [passive](#) statement on Device R3's interface to Device S1.
- The loopback addresses of Level 2 devices do not appear in a Level 1 area.
- There is only one adjacency for each device pairing.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R3

```
set interfaces fe-1/2/0 unit 0 description to-R4
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.41/30
set interfaces fe-1/2/2 unit 0 description to-S1
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis interface fe-1/2/2.0 passive
```

Device R4

```
set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
```

Device R5	<pre> set interfaces fe-1/2/0 unit 0 description to-R3 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R4 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/2 unit 0 description to-R6 set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30 set interfaces fe-1/2/2 unit 0 family iso set interfaces fe-1/2/3 unit 0 description to-R7 set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.38/30 set interfaces fe-1/2/3 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.5/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0005.00 set protocols isis interface fe-1/2/0.0 level 1 disable set protocols isis interface fe-1/2/1.0 level 1 disable set protocols isis interface fe-1/2/2.0 level 2 disable set protocols isis interface fe-1/2/3.0 level 2 disable set protocols isis interface lo0.0 level 1 disable </pre>
Device R6	<pre> set interfaces fe-1/2/0 unit 0 description to-R5 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.30/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R7 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.33/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.6/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0006.00 set protocols isis interface fe-1/2/0.0 level 2 disable set protocols isis interface fe-1/2/1.0 level 2 disable set protocols isis interface lo0.0 level 2 disable </pre>
Device R7	<pre> set interfaces fe-1/2/0 unit 0 description to-R6 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.34/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/1 unit 0 description to-R5 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces lo0 unit 0 family inet address 192.168.0.7/32 set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0007.00 set protocols isis interface fe-1/2/0.0 level 2 disable set protocols isis interface fe-1/2/1.0 level 2 disable set protocols isis interface lo0.0 level 2 disable </pre>
Device S1	<pre> set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.42/30 set interfaces fe-1/2/0 unit 0 description to-R3 </pre>
Step-by-Step Procedure	<p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure multi-level IS-IS:</p> <ol style="list-style-type: none"> 1. Configure the network interfaces.

Enable IS-IS on the interfaces by including the ISO address family on each interface.

```
[edit interfaces]
user@R5# set fe-1/2/0 unit 0 description to-R3
user@R5# set fe-1/2/0 unit 0 family inet address 10.0.0.22/30
user@R5# set fe-1/2/0 unit 0 family iso
user@R5# set fe-1/2/1 unit 0 description to-R4
user@R5# set fe-1/2/1 unit 0 family inet address 10.0.0.26/30
user@R5# set fe-1/2/1 unit 0 family iso
user@R5# set fe-1/2/2 unit 0 description to-R6
user@R5# set fe-1/2/2 unit 0 family inet address 10.0.0.29/30
user@R5# set fe-1/2/2 unit 0 family iso
user@R5# set fe-1/2/3 unit 0 description to-R7
user@R5# set fe-1/2/3 unit 0 family inet address 10.0.0.38/30
user@R5# set fe-1/2/3 unit 0 family iso
```

2. Configure two loopback interface addresses.

One address is for IPv4.

The other is for the IS-IS area 49.002 so that Device R5 can form adjacencies with the other Level 1 devices in area 49.002. Even though Device R5's NET identifies itself as belonging to the Level 1 area 49.002, its loopback interface is not configured as a Level 1 interface. Doing so would cause the route to Device R5's loopback to be injected into the Level 1 area.

```
[edit interfaces lo0 unit 0]
user@R5# set family inet address 192.168.0.5/32
user@R5# set family iso address 49.002.0192.0168.0005.00
```

3. Specify the IS-IS level on a per-interface basis.

Device R5 becomes adjacent to the other routing devices on the same level on each link.

By default, IS-IS is enabled for IS-IS areas on all interfaces on which the ISO protocol family is enabled (at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level). To disable IS-IS at any particular level on an interface, include the **disable** statement.

Device R5's loopback interface is configured to run Level 2 only. If Level 1 operation were enabled on lo0.0, Device R5 would include its loopback address in its Level 1 link-state PDU, which is incorrect for this example in which the loopback addresses of Level 2 devices must not appear in a Level 1 area.

Unlike OSPF, you must explicitly list the router's lo0 interface at the **[edit protocols isis]** hierarchy level, because this interface is the source of the router's NET, and therefore must be configured as an IS-IS interface. In IS-IS, the lo0 interface operates in the passive mode by default, which is ideal because adjacency formation can never occur on a virtual interface.

```
[edit protocols isis]
user@R5# set interface fe-1/2/0.0 level 1 disable
user@R5# set interface fe-1/2/1.0 level 1 disable
user@R5# set interface fe-1/2/0.0 level 2 disable
user@R5# set interface fe-1/2/3.0 level 2 disable
user@R5# set interface lo0.0 level 1 disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R5# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-R3;
    family inet {
      address 10.0.0.22/30;
    }
    family iso;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-R4;
    family inet {
      address 10.0.0.26/30;
    }
    family iso;
  }
}
fe-1/2/2 {
  unit 0 {
    description to-R6;
    family inet {
      address 10.0.0.29/30;
    }
    family iso;
  }
}
fe-1/2/3 {
  unit 0 {
    description to-R7;
    family inet {
      address 10.0.0.38/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.5/32;
    }
    family iso {
      address 49.002.0192.0168.0005.00;
    }
  }
}
user@R5# show protocols
isis {
  interface fe-1/2/0.0 {
    level 1 disable;
  }
}

```

```
}
interface fe-1/2/1.0 {
  level 1 disable;
}
interface fe-1/2/0.0 {
  level 2 disable;
}
interface fe-1/2/3.0 {
  level 2 disable;
}
interface lo0.0 {
  level 1 disable;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking Interface-to-Area Associations on page 3644](#)
- [Verifying IS-IS Adjacencies on page 3644](#)
- [Examining the IS-IS Database on page 3645](#)

Checking Interface-to-Area Associations

Purpose Make sure that the interface-to-area associations are configured as expected.

Action From operational mode, enter the **show isis interface** command.

```
user@R5> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0  0x1 Disabled           Passive          0/0
fe-1/2/0.0         2  0x3 Disabled          R5.03           10/10
fe-1/2/1.0         2  0x2 Disabled          R5.02           10/10
fe-1/2/0.0         1  0x1 R6.02             Disabled         10/10
fe-1/2/3.0         1  0x4 R5.04             Disabled         10/10
```

Meaning The output shows that Device R5's interfaces have been correctly configured with the ISO family, and that the interfaces have been placed into the correct levels.

You can also see that Device R5 has elected itself as the designated intermediate system (DIS) on its broadcast-capable IS-IS interfaces.

Verifying IS-IS Adjacencies

Purpose Verify that the expected adjacencies have formed between Device R5 and its IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R5> show isis adjacency detail
```



```

R3
Interface: fe-1/2/0.0, Level: 2, State: Up, Expires in 25 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:31 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.03, IP addresses: 10.0.0.21

R4
Interface: fe-1/2/1.0, Level: 2, State: Up, Expires in 24 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:36 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.02, IP addresses: 10.0.0.25

R6
Interface: fe-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:20:24 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R6.02, IP addresses: 10.0.0.30

R7
Interface: fe-1/2/3.0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:29 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.04, IP addresses: 10.0.0.37

```

Meaning These results confirm that Device R5 has two Level 2 adjacencies and two Level 1 adjacencies.

Examining the IS-IS Database

Purpose Because Device R5 is a L1/L2 attached router, examine the Level 1 link-state database associated with area 49.002 to confirm that loopback addresses from backbone routers are not being advertised into the Level 1 area.

Action From operational mode, enter the **show isis database detail** command.

```

user@R5> show isis database detail
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0x19, Checksum: 0x7488, Lifetime: 727 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R6.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up

R5.04-00 Sequence: 0x14, Checksum: 0x2668, Lifetime: 821 secs
  IS neighbor: R5.00                      Metric:       0
  IS neighbor: R7.00                      Metric:       0

R6.00-00 Sequence: 0x17, Checksum: 0xa65, Lifetime: 774 secs
  IS neighbor: R6.02                      Metric:      10
  IS neighbor: R7.02                      Metric:      10

```

```

IP prefix: 10.0.0.28/30          Metric:      10 Internal Up
IP prefix: 10.0.0.32/30          Metric:      10 Internal Up
IP prefix: 192.168.0.6/32        Metric:      0 Internal Up

R6.02-00 Sequence: 0x13, Checksum: 0xd1c0, Lifetime: 908 secs
IS neighbor: R5.00               Metric:      0
IS neighbor: R6.00               Metric:      0

R7.00-00 Sequence: 0x17, Checksum: 0xe39, Lifetime: 775 secs
IS neighbor: R5.04               Metric:      10
IS neighbor: R7.02               Metric:      10
IP prefix: 10.0.0.32/30          Metric:      10 Internal Up
IP prefix: 10.0.0.36/30          Metric:      10 Internal Up
IP prefix: 192.168.0.7/32        Metric:      0 Internal Up

R7.02-00 Sequence: 0x13, Checksum: 0x404d, Lifetime: 966 secs
IS neighbor: R6.00               Metric:      0
IS neighbor: R7.00               Metric:      0

IS-IS level 2 link-state database:

R3.00-00 Sequence: 0x17, Checksum: 0x5f84, Lifetime: 1085 secs
IS neighbor: R4.02               Metric:      10
IS neighbor: R5.03               Metric:      10
IP prefix: 10.0.0.16/30          Metric:      10 Internal Up
IP prefix: 10.0.0.20/30          Metric:      10 Internal Up
IP prefix: 10.0.0.40/30          Metric:      10 Internal Up
IP prefix: 192.168.0.3/32        Metric:      0 Internal Up

R4.00-00 Sequence: 0x17, Checksum: 0xab3a, Lifetime: 949 secs
IS neighbor: R4.02               Metric:      10
IS neighbor: R5.02               Metric:      10
IP prefix: 10.0.0.16/30          Metric:      10 Internal Up
IP prefix: 10.0.0.24/30          Metric:      10 Internal Up
IP prefix: 192.168.0.4/32        Metric:      0 Internal Up

R4.02-00 Sequence: 0x14, Checksum: 0xf2a8, Lifetime: 1022 secs
IS neighbor: R3.00               Metric:      0
IS neighbor: R4.00               Metric:      0

R5.00-00 Sequence: 0x1f, Checksum: 0x20d7, Lifetime: 821 secs
IS neighbor: R5.02               Metric:      10
IS neighbor: R5.03               Metric:      10
IP prefix: 10.0.0.20/30          Metric:      10 Internal Up
IP prefix: 10.0.0.24/30          Metric:      10 Internal Up
IP prefix: 10.0.0.28/30          Metric:      10 Internal Up
IP prefix: 10.0.0.32/30          Metric:      20 Internal Up
IP prefix: 10.0.0.36/30          Metric:      10 Internal Up
IP prefix: 192.168.0.5/32        Metric:      0 Internal Up
IP prefix: 192.168.0.6/32        Metric:      10 Internal Up
IP prefix: 192.168.0.7/32        Metric:      10 Internal Up

R5.02-00 Sequence: 0x14, Checksum: 0x6135, Lifetime: 977 secs
IS neighbor: R4.00               Metric:      0
IS neighbor: R5.00               Metric:      0

R5.03-00 Sequence: 0x14, Checksum: 0x1483, Lifetime: 1091 secs
IS neighbor: R3.00               Metric:      0
IS neighbor: R5.00               Metric:      0

```

Meaning This display indicates that Device R5's loopback interface is correctly configured to run Level 2 only. Had Level 1 operation been enabled on lo0.0, Device R5 would have then included its loopback address in its Level 1 link-state PDU.

You can also see that Device R5 has Level 2 link-state PDUs, received from its adjacent neighbors.

Like an OSPF totally stubby area, no backbone (Level 2) or external prefixes are leaked into a Level 1 area, by default. Level 1 prefixes are leaked up into the IS-IS backbone, however, as can be seen in Device R5's Level 2 link-state PDU.

Related Documentation

- [Understanding IS-IS Areas](#)

Example: Configuring Hitless Authentication Key Rollover for IS-IS

This example shows how to configure hitless authentication key rollover for IS-IS.

- [Requirements on page 3647](#)
- [Overview on page 3647](#)
- [Configuration on page 3648](#)
- [Verification on page 3651](#)

Requirements

No special configuration beyond device initialization is required before configuring hitless authentication key rollover for IS-IS.

Overview

Authentication guarantees that only trusted routers participate in routing updates. This keychain authentication method is referred to as hitless because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol. Junos OS supports both RFC 5304, *IS-IS Cryptographic Authentication* and RFC 5310, *IS-IS Generic Cryptographic Authentication*.

This example includes the following statements for configuring the keychain:

- **algorithm**—For each key in the keychain, you can specify an encryption algorithm. The algorithm can be SHA-1 or MD-5.
- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.
- **key-chain**—For each keychain, you must specify a name. This example defines two keychains: **base-key-global** and **base-key-inter**.
- **options**—For each key in the keychain, you can specify the encoding for the message authentication code: **isis-enhanced** or **basic**. The basic (RFC 5304) operation is enabled by default.

When you configure the **isis-enhanced** option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.

When you configure **basic** (or do not include the **options** statement in the key configuration), Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.

Because this setting is for IS-IS only, the TCP and the BFD protocols ignore the encoding option configured in the key.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the key chain.

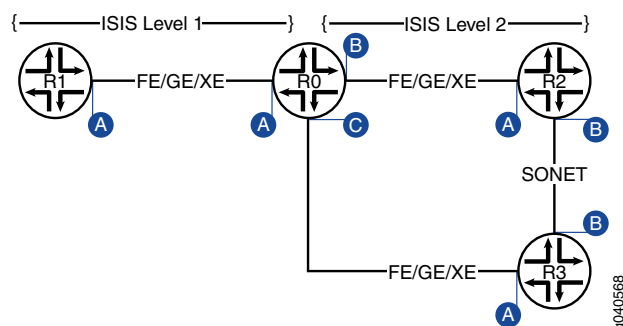
You can apply a keychain globally to all interfaces or more granularly to specific interfaces.

This example includes the following statements for applying the keychain to all interfaces or to particular interfaces:

- **authentication-key-chain**—Enables you to apply a keychain at the global IS-IS level for all Level 1 or all Level 2 interfaces.
- **hello-authentication-key-chain**—Enables you to apply a keychain at the individual IS-IS interface level. The interface configuration overrides the global configuration.

Figure 113 on page 3648 shows the topology used in the example.

Figure 113: Hitless Authentication Key Rollover for IS-IS



This example shows the configuration for Router R0.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set interfaces ge-0/0/0 unit 0 description "interface A"
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address fe80::200:f8ff:fe21:67cf/128
set interfaces ge-0/0/1 unit 0 description "interface B"
set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 10FB::C:ABC:1FOC:44DA/128
set interfaces ge-0/0/2 unit 0 description "interface C"
set interfaces ge-0/0/2 unit 0 family inet address 10.0.0.9/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
set security authentication-key-chains key-chain base-key-global key 63 secret
"$9$jfkdTQnCpBDiCt"
set security authentication-key-chains key-chain base-key-global key 63 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret
"$9$8sgx7Vws4ZDkWLGD"
set security authentication-key-chains key-chain base-key-inter key 0 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-inter key 0 algorithm md5
set security authentication-key-chains key-chain base-key-inter key 0 options basic
set protocols isis level 1 authentication-key-chain base-key-global
set protocols isis interface ge-0/0/0.0 level 1 hello-authentication-key-chain
base-key-inter

```

Step-by-Step Procedure

To configure hitless authentication key rollover for IS-IS:

1. Configure the Router R0 interfaces.

```

[edit interfaces ge-0/0/0 unit 0]
user@R0# set description "interface A"
user@R0# set family inet address 10.0.0.1/30
user@R0# set family iso
user@R0# set family inet6 address fe80::200:f8ff:fe21:67cf/128
[edit interfaces ge-0/0/1 unit 0]
user@R0# set interfaces ge-0/0/1 unit 0 description "interface B"
user@R0# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
user@R0# set interfaces ge-0/0/1 unit 0 family iso
user@R0# set interfaces ge-0/0/1 unit 0 family inet6 address
10FB::C:ABC:1FOC:44DA/128
[edit interfaces ge-0/0/2 unit 0]
user@R0# set description "interface C"
user@R0# set family inet address 10.0.0.9/30
user@R0# set interfaces ge-0/0/2 unit 0 family iso
user@R0# set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128

```

2. Configure one or more authentication keys.

```

[edit security authentication-key-chains key-chain base-key-global]
user@R0# set key 63 secret "$9$jfkdTQnCpBDiCt"
user@R0# set key 63 start-time "2011-8-6.06:54:00-0700"

```

```
user@R0# set key 63 algorithm hmac-sha-1
user@R0# set key 63 options isis-enhanced
[edit security authentication-key-chains key-chain base-key-inter]
user@R0# set key 0 secret "$9$8sgx7Vws4ZDkWLGD"
user@R0# set key 0 start-time "2011-8-6.06:54:00-0700"
user@R0# set key 0 algorithm md5
user@R0# set key 0 options basic
```

3. Apply the base-key-global keychain to all Level 1 IS-IS interfaces on Router R0.

```
[edit protocols isis level 1]
user@R0# set authentication-key-chain base-key-global
```

4. Apply the base-key-inter keychain to the ge-0/0/0.0 interface on Router R0.

```
[edit protocols isis interface ge-0/0/0.0 level 1]
user@R0# set hello-authentication-key-chain base-key-inter
```

5. If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    description "interface A";
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address fe80::200:f8ff:fe21:67cf/128;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description "interface B";
    family inet {
      address 10.0.0.5/30;
    }
    family iso;
    family inet6 {
      address 10fb::c:abc:1f0c:44da/128;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    description "interface C";
    family inet {
```

```

        address 10.0.0.9/30;
    }
    family iso;
    family inet6 {
        address ff06::c3/128;
    }
}

user@R0# show protocols
isis {
    level 1 authentication-key-chain base-key-global;
    interface ge-0/0/0.0 {
        level 1 hello-authentication-key-chain base-key-inter;
    }
}

user@R0# show security
authentication-key-chains {
    key-chain base-key-global {
        key 63 {
            secret "$9$jfkqfTQnCpBDiCt"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm hmac-sha-1;
            options isis-enhanced;
        }
    }
    key-chain base-key-inter {
        key 0 {
            secret "$9$8sgx7Vws4ZDkWLGD"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm md5;
            options basic;
        }
    }
}

```

Verification

To verify the configuration, run the following commands:

- [show isis authentication](#)
- [show security keychain](#)

Related Documentation

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 3631](#)

Example: Redistributing OSPF Routes into IS-IS

This example shows how to redistribute OSPF routes into an IS-IS network.

- [Requirements on page 3652](#)
- [Overview on page 3652](#)

- [Configuration on page 3653](#)
- [Verification on page 3658](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Export policy can be applied to IS-IS to facilitate route redistribution.


Junos OS does not support the application of import policy for link-state routing protocols like IS-IS because such policies can lead to inconsistent link-state database (LSDB) entries, which in turn can result in routing inconsistencies.

In this example, OSPF routes 192.168.0/24 through 192.168.3/24 are redistributed into IS-IS area 49.0002 from Device R2.

In addition, policies are configured to ensure that Device R1 can reach destinations on the 10.0.0.44/30 network, and that Device R3 can reach destinations on the 10.0.0.36/30 network. This enables end-to-end reachability.

[Figure 114 on page 3652](#) shows the topology used in this example.

Figure 114: IS-IS Route Redistribution Topology



“CLI Quick Configuration” on [page 3653](#) shows the configuration for all of the devices in [Figure 114 on page 3652](#). The section “Step-by-Step Procedure” on [page 3654](#) describes the steps on Device R2. “Step-by-Step Procedure” on [page 3655](#) describes the steps on Device R3.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set interfaces fe-1/2/0 unit 0 description to-R7 set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.38/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces lo0 unit 0 family inet address 172.16.3.5/32 set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0305.00 set protocols isis interface fe-1/2/0.38 set protocols isis interface lo0.0 </pre>
Device R2	<pre> set interfaces fe-1/2/1 unit 0 description to-R5 set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/0 unit 0 description to-OSPF-network set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.45/30 set interfaces lo0 unit 0 family inet address 172.16.9.7/32 set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0907.00 set protocols isis export ospf-isis set protocols isis export send-direct-to-isis-neighbors set protocols isis interface fe-1/2/1.0 set protocols isis interface lo0.0 set protocols ospf export send-direct-to-ospf-neighbors set protocols ospf area 0.0.0.1 interface fe-1/2/0.0 set protocols ospf area 0.0.0.1 interface lo0.0 passive set policy-options policy-statement ospf-isis term 1 from protocol ospf set policy-options policy-statement ospf-isis term 1 from route-filter 192.168.0.0/22 longer set policy-options policy-statement ospf-isis term 1 then accept set policy-options policy-statement send-direct-to-isis-neighbors from protocol direct set policy-options policy-statement send-direct-to-isis-neighbors from route-filter 10.0.0.44/30 exact set policy-options policy-statement send-direct-to-isis-neighbors then accept set policy-options policy-statement send-direct-to-ospf-neighbors from protocol direct set policy-options policy-statement send-direct-to-ospf-neighbors from route-filter 10.0.0.36/30 exact set policy-options policy-statement send-direct-to-ospf-neighbors then accept </pre>
Device R3	<pre> set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.46/30 set interfaces lo0 unit 0 family inet address 192.168.1.1/32 set interfaces lo0 unit 0 family inet address 192.168.2.1/32 set interfaces lo0 unit 0 family inet address 192.168.3.1/32 set interfaces lo0 unit 0 family inet address 192.168.0.1/32 set protocols ospf export ospf set protocols ospf area 0.0.0.1 interface fe-1/2/0.0 set protocols ospf area 0.0.0.1 interface lo0.0 passive set policy-options policy-statement ospf term 1 from protocol static set policy-options policy-statement ospf term 1 then accept set routing-options static route 192.168.0.0/24 discard set routing-options static route 192.168.1.0/24 discard set routing-options static route 192.168.3.0/24 discard </pre>

```
set routing-options static route 192.168.2.0/24 discard
```

**Step-by-Step
Procedure**

To configure Device R2:

1. Configure the network interfaces.

```
[edit interfaces]  
user@R2# set fe-1/2/1 unit 0 description to-R5  
user@R2# set fe-1/2/1 unit 0 family inet address 10.0.0.37/30  
user@R2# set fe-1/2/1 unit 0 family iso  
user@R2# set fe-1/2/0 unit 0 description to-OSPF-network  
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.45/30  
user@R2# set lo0 unit 0 family inet address 172.16.9.7/32  
user@R2# set lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
```
2. Configure IS-IS on the interface facing Device R1 and the loopback interface.

```
[edit protocols isis]  
user@R2# set interface fe-1/2/1.0  
user@R2# set interface lo0.0
```
3. Configure the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit policy-options policy-statement send-direct-to-isis-neighbors]  
user@R2# set from protocol direct  
user@R2# set from route-filter 10.0.0.44/30 exact  
user@R2# set then accept
```
4. Apply the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit protocols isis]  
user@R2# set export send-direct-to-isis-neighbors
```
5. Configure OSPF on the interfaces.

```
[edit protocols ospf]  
user@R2# set area 0.0.0.1 interface fe-1/2/0.0  
user@R2# set area 0.0.0.1 interface lo0.0 passive
```
6. Configure the OSPF route redistribution policy.

```
[edit policy-options policy-statement ospf-isis term 1]  
user@R2# set from protocol ospf  
user@R2# set from route-filter 192.168.0.0/22 longer  
user@R2# set then accept
```
7. Apply the OSPF route redistribution policy to the IS-IS instance.

```
[edit protocols isis]  
user@R2# set export ospf-isis
```
8. Configure the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit policy-options policy-statement send-direct-to-ospf-neighbors]  
user@R2# set from protocol direct  
user@R2# set from route-filter 10.0.0.36/30 exact  
user@R2# set then accept
```
9. Apply the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit protocols ospf]  
user@R2# set export send-direct-to-ospf-neighbors
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1. Configure the network interfaces.

Multiple addresses are configured on the loopback interface to simulate multiple route destinations.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.46/30
user@R3# set lo0 unit 0 family inet address 192.168.1.1/32
user@R3# set lo0 unit 0 family inet address 192.168.2.1/32
user@R3# set lo0 unit 0 family inet address 192.168.3.1/32
user@R3# set lo0 unit 0 family inet address 192.168.0.1/32
```

2. Configure static routes to the loopback interface addresses.

These are the routes that are redistributed into IS-IS.

```
[edit routing-options static]
user@R3# set route 192.168.0.0/24 discard
user@R3# set route 192.168.1.0/24 discard
user@R3# set route 192.168.3.0/24 discard
user@R3# set route 192.168.2.0/24 discard
```

3. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.1]
user@R3# set interface fe-1/2/0.0
user@R3# set interface lo0.0 passive
```

4. Configure the OSPF policy to export the static routes.

```
[edit policy-options policy-statement ospf term 1]
user@R3# set from protocol static
user@R3# set then accept
```

5. Apply the OSPF export policy.

```
[edit protocols ospf]
user@R3# set export ospf
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device R2 user@R2# show interfaces
fe-1/2/1 {
  unit 0 {
    description to-R5;
    family inet {
      address 10.0.0.37/30;
    }
    family iso;
  }
}
```

```
}
fe-1/2/0 {
  unit 0 {
    description to-OSPF-network;
    family inet {
      address 10.0.0.45/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.9.7/32;
    }
    family iso {
      address 49.0002.0172.0016.0907.00;
    }
  }
}

user@R2# show protocols
isis {
  export [ ospf-isis send-direct-to-isis-neighbors ];
  interface fe-1/2/1.0;
  interface lo0.0;
}
ospf {
  export send-direct-to-ospf-neighbors;
  area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

user@R2# show policy-options
policy-statement ospf-isis {
  term 1 {
    from {
      protocol ospf;
      route-filter 192.168.0.0/22 longer;
    }
    then accept;
  }
}
policy-statement send-direct-to-isis-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.44/30 exact;
  }
  then accept;
}
policy-statement send-direct-to-ospf-neighbors {
  from {
    protocol direct;
```

```

        route-filter 10.0.0.36/30 exact;
    }
    then accept;
}

Device R3 user@R3# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet {
            address 10.0.0.46/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.1.1/32;
            address 192.168.2.1/32;
            address 192.168.3.1/32;
            address 192.168.0.1/32;
        }
    }
}

user@R3# show protocols
ospf {
    export ospf;
    area 0.0.0.1 {
        interface fe-1/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}

user@R3# show policy-options
policy-statement ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show routing-options
static {
    route 192.168.0.0/24 discard;
    route 192.168.1.0/24 discard;
    route 192.168.3.0/24 discard;
    route 192.168.2.0/24 discard;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Route Advertisement on page 3658](#)
- [Verifying Route Redistribution on page 3658](#)
- [Verifying Connectivity on page 3659](#)

Verifying OSPF Route Advertisement

Purpose Make sure that the expected routes are advertised by OSPF.

Action From operational mode on Device R2, enter the **show route protocol ospf** command.

```
user@R2> show route protocol ospf
```

```
inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.0.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.0.1/32    *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.1.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.1.1/32    *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.2.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.2.1/32    *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.3.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                  > to 10.0.0.46 via fe-1/2/0.0
192.168.3.1/32    *[OSPF/10] 03:54:21, metric 1
                  > to 10.0.0.46 via fe-1/2/0.0
224.0.0.5/32     *[OSPF/10] 03:56:03, metric 1
                  MultiRecv
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning The 192.168/16 routes are advertised by OSPF.

Verifying Route Redistribution

Purpose Make sure that the expected routes are redistributed from OSPF into IS-IS.

Action From operational mode on Device R1, enter the **show route protocol isis** command.

```
user@R1> show route protocol isis
```

```
inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.44/30     *[IS-IS/160] 03:45:24, metric 20
                  > to 10.0.0.37 via fe-1/2/0.0
```

```

172.16.9.7/32      *[IS-IS/15] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.0.0/24    *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.0.1/32    *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.1.0/24    *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.1.1/32    *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.2.0/24    *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.2.1/32    *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.3.0/24    *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.3.1/32    *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

Meaning The 192.168/16 routes are redistributed into IS-IS.

Verifying Connectivity

Purpose Check that Device R1 can reach the destinations on Device R3.

Action From operational mode, enter the **ping** command.

```

user@R1> ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=63 time=2.089 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=63 time=1.270 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=63 time=2.135 ms

```

Meaning These results confirm that Device R1 can reach the destinations in the OSPF network.

Related Documentation

- [Understanding Routing Policies](#)

Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.

- [Requirements on page 3659](#)
- [Overview on page 3660](#)
- [Configuration on page 3660](#)
- [Verification on page 3663](#)

Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3633](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

Figure 115 on page 3660 shows the sample network.

Figure 115: Configuring BFD for IS-IS



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

Router R2

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```


Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.



NOTE: To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



NOTE: You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.


```
[edit protocols isis]
user@R1# set interface so-0/0/0 bfd-liveness-detection

[edit protocols isis]
user@R2# set interface so-0/0/0 bfd-liveness-detection
```
2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set detection-time threshold 5

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set detection-time threshold 6
```
3. Configure the minimum transmit and receive intervals for failure detection.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-interval 2

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-interval 3
```
4. Configure only the minimum receive interval for failure detection.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-receive-interval 1

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-receive-interval 1
```
5. Disable BFD adaptation.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols isis interface so-0/0/0
```

```
bfd-liveness-detection {
  version automatic;
  minimum-interval 2;
  minimum-receive-interval 1;
  multiplier 2;
  no-adaptation;
  transmit-interval {
    minimum-interval 1;
    threshold 3;
  }
  detection-time {
    threshold 5;
  }
}
```

```
...
```

```
user@R2# show protocols isis interface so-0/0/0
```

```
    bfd-liveness-detection {
      version automatic;
      minimum-interval 3;
      minimum-receive-interval 1;
      multiplier 2;
      no-adaptation;
      transmit-interval {
        minimum-interval 1;
        threshold 4;
      }
      detection-time {
        threshold 6;
      }
    }
  ...
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 3663](#)
- [Verifying That IS-IS Is Configured on page 3664](#)
- [Verifying That BFD Is configured on page 3664](#)

Verifying the Connection Between Routers R1 and R2

Purpose Make sure that Routers R1 and R2 are connected to each other.

Action Ping the other router to check the connectivity between the two routers as per the network topology.

```
user@R1> ping 10.0.0.2
```

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```

```
user@R2> ping 10.0.0.1
```

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms
```

Meaning Routers R1 and R2 are connected to each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on both routers.

Action Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

user@R1> show isis database

```
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a571  0x30c5    1195 L1 L2
R2.00-00        0x4a586  0x4b7e    1195 L1 L2
R2.02-00        0x330ca1 0x3492    1196 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a856  0x5db0    1194 L1 L2
R2.00-00        0x4a89d  0x149b    1194 L1 L2
R2.02-00        0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs
```

user@R2> show isis database

```
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4b707  0xcc80    1195 L1 L2
R2.00-00        0x4b71b  0xeb37    1198 L1 L2
R2.02-00        0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4b9f2  0xee70    1192 L1 L2
R2.00-00        0x4ba41  0x9862    1197 L1 L2
R2.02-00        0x3      0x6242    1198 L1 L2
  3 LSPs
```

Meaning IS-IS is configured on both routers, R1 and R2.

Verifying That BFD Is configured

Purpose Make sure that the BFD instance is running on both routers, R1 and R2.

Action Use the **show bfd session detail** statement to check if BFD instance is running on the routers.

user@R1> show bfd session detail

```
Address          State      Interface      Detect   Transmit
10.0.0.2          Up         so-0/0/0       Time    Interval  Multiplier
                  Up         so-0/0/0       2.000   1.000     2
Client ISIS R2, TX interval 0.001, RX interval 0.001
Client ISIS R1, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:15
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 3, routing table index 17
```

1 sessions, 2 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

user@R2> show bfd session detail

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.1	Up	so-0/0/0	2.000	1.000	2

Client ISIS R2, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:05
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 2, routing table index 15

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

Meaning BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

Related Documentation

- [Understanding BFD for IS-IS](#)

Example: Configuring BFD Authentication for IS-IS

This example shows how to configure BFD authentication for IS-IS.

- [Requirements on page 3665](#)
- [Overview on page 3665](#)
- [Configuration on page 3666](#)
- [Verification on page 3667](#)

Requirements


Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3633](#) for information about the required IS-IS configuration.

Overview

In this example, a BFD authentication keychain is configured with meticulous keyed MD5 authentication.

[Figure 116 on page 3665](#) shows the topology used in this example.

Figure 116: IS-IS BFD Authentication Topology



[“CLI Quick Configuration” on page 3666](#) shows the configuration for both of the devices in [Figure 116 on page 3665](#). The section [“Step-by-Step Procedure” on page 3666](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set security authentication-key-chains key-chain secret123 description for-isis-bfd
set security authentication-key-chains key-chain secret123 key 1 secret "$9$cW-yrv"
set security authentication-key-chains key-chain secret123 key 1 start-time
  "2012-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 2 secret "$9$m5T3"
set security authentication-key-chains key-chain secret123 key 2 start-time
  "2013-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"
set security authentication-key-chains key-chain secret123 key 3 start-time
  "2014-5-31.13:00:00 -0700"
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain
secret123
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm
meticulous-keyed-md5
```

Device R2

```
set security authentication-key-chains key-chain secret123 description for-isis-bfd
set security authentication-key-chains key-chain secret123 key 1 secret "$9$cW-yrv"
set security authentication-key-chains key-chain secret123 key 1 start-time
  "2012-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 2 secret "$9$m5T3"
set security authentication-key-chains key-chain secret123 key 2 start-time
  "2013-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"
set security authentication-key-chains key-chain secret123 key 3 start-time
  "2014-5-31.13:00:00 -0700"
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain
secret123
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm
meticulous-keyed-md5
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS BFD authentication:

1. Configure the authentication keychain.

```
[edit security authentication-key-chains key-chain secret123]
user@R1# set description for-isis-bfd
user@R1# set key 1 secret "$9$cW-yrv"
user@R1# set key 1 start-time "2012-5-31.13:00:00 -0700"
user@R1# set key 2 secret "$9$m5T3"
user@R1# set key 2 start-time "2013-5-31.13:00:00 -0700"
user@R1# set key 3 secret "$9$mTQn"
user@R1# set key 3 start-time "2014-5-31.13:00:00 -0700"
```

2. Enable BFD.


```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set minimum-interval 100
```
3. Apply the authentication keychain.


```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set authentication key-chain secret123
```
4. Set the authentication type.


```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set authentication algorithm meticulous-keyed-md5
```

Results From configuration mode, confirm your configuration by entering the **show protocols** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
isis {
  interface ge-1/2/0.0 {
    bfd-liveness-detection {
      minimum-interval 100;
      authentication {
        key-chain secret123;
        algorithm meticulous-keyed-md5;
      }
    }
  }
}

user@R1# show security
authentication-key-chains {
  key-chain secret123 {
    description for-isis-bfd;
    key 1 {
      secret "$9$cW-yrv"; ## SECRET-DATA
      start-time "2012-5-31.13:00:00 -0700";
    }
    key 2 {
      secret "$9$m5T3"; ## SECRET-DATA
      start-time "2013-5-31.13:00:00 -0700";
    }
    key 3 {
      secret "$9$mTQn"; ## SECRET-DATA
      start-time "2014-5-31.13:00:00 -0700";
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying IS-IS BFD Authentication

Purpose Verify the status of IS-IS BFD authentication.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Down	ge-1/2/0.0	0.300	1.000	3

Client ISIS L1, TX interval 0.100, RX interval 0.100, Authenticate
keychain secret123, algo meticulous-keyed-md5, mode strict
Client ISIS L2, TX interval 0.100, RX interval 0.100, Authenticate
keychain secret123, algo meticulous-keyed-md5, mode strict
Session down time 00:35:13, previous up time 00:12:17
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 2, routing table index 85
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 0.100, min RX interval 0.100, multiplier 3
Local discriminator 2, remote discriminator 1
Echo mode disabled/inactive, no-absorb, no-refresh
Authentication enabled/active, keychain secret123, algo meticulous-keyed-md5, mode strict
Session ID: 0x100101

1 sessions, 2 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 10.0 pps

Meaning The output shows that BFD authentication is enabled on IS-IS Level 1 and Level 2.

Related Documentation

- [Configuring BFD Authentication for IS-IS](#)
- [Example: Configuring BFD for IS-IS](#)

Example: Configuring IS-IS Multicast Topology

- [IS-IS Multicast Topologies Overview on page 3669](#)
- [Example: Configuring IS-IS Multicast Topology on page 3670](#)

IS-IS Multicast Topologies Overview

Most multicast routing protocols perform a reverse-path forwarding (RPF) check on the source of multicast data packets. If a packet comes in on the interface that is used to send data to the source, the packet is accepted and forwarded to one or more downstream interfaces. Otherwise, the packet is discarded and a notification is sent to the multicast routing protocol running on the interface.

In certain instances, the unicast routing table used for the RPF check is also the table used for forwarding unicast data packets. Thus, unicast and multicast routing are congruent. In other cases, where it is preferred that multicast routing be independent of unicast routing, the multicast routing protocols are configured to perform the RPF check using an alternate unicast routing table `inet.2`.

You can configure IS-IS to calculate an alternate IPv4 multicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to `inet.2`. The IS-IS interface metrics for the multicast topology can be configured independently of the unicast metrics. You can also selectively disable interfaces from participating in the multicast topology while continuing to participate in the regular unicast topology. This enables you to exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths. You can also configure IS-IS to calculate an alternate IPv6 multicast topology, in addition to the normal IPv6 unicast topology.



NOTE: IS-IS only starts advertising the routes when the interface routes are in `inet.2`.

Table 292 on page 3669 lists the various IPv4 statements you can use to configure IS-IS topologies.

Table 292: IPv4 Statements

Statement	Description
<code>ipv4-multicast</code>	Enables an alternate IPv4 multicast topology.
<code>ipv4-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv4 multicast topology.
<code>no-ipv4-multicast</code>	Excludes an interface from the IPv4 multicast topology.
<code>no-unicast-topology</code>	Excludes an interface from the IPv4 unicast topologies.

Table 293 on page 3669 lists the various IPv6 statements you can use to configure IS-IS topologies.

Table 293: IPv6 Statements

Statement	Description
<code>ipv6-multicast</code>	Enables an alternate IPv6 multicast topology.

Table 293: IPv6 Statements (*continued*)

Statement	Description
<code>ipv6-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv6 multicast topology.
<code>ipv6-unicast-metric <i>number</i></code>	Configures the unicast metric for an alternate IPv6 multicast topology.
<code>no-ipv6-multicast</code>	Excludes an interface from the IPv6 multicast topology.
<code>no-ipv6-unicast</code>	Excludes an interface from the IPv6 unicast topologies.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Example: Configuring IS-IS Multicast Topology

This example shows how to configure a multicast topology for an IS-IS network.

- [Requirements on page 3670](#)
- [Overview on page 3670](#)
- [Configuration on page 3671](#)
- [Verification on page 3675](#)

Requirements

Before you begin, configure IS-IS on all routers. See [“Example: Configuring IS-IS” on page 3633](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows an IS-IS multicast topology configuration. Three routers are connected to each other. A loopback interface is configured on each router.

[Figure 117 on page 3671](#) shows the sample network.

Figure 117: Configuring IS-IS Multicast Topology



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 15
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-0/0/0 level 2 metric 20
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 14
set protocols isis interface so-1/0/0 level 1 metric 13
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-1/0/0 level 2 metric 29
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface fxp0.0 disable
```

Router R2

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 13
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-0/0/0 level 2 metric 29
```

```
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface so-1/0/0 level 1 metric 14
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-1/0/0 level 2 metric 32
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 26
set protocols isis interface fxp0.0 disable
```

Router R3

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 19
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 11
set protocols isis interface so-0/0/0 level 2 metric 27
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 21
set protocols isis interface so-1/0/0 level 1 metric 16
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 26
set protocols isis interface so-1/0/0 level 2 metric 30
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 20
set protocols isis interface fxp0.0 disable
```

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS multicast topologies:

1. Enable the multicast topology for IS-IS by using the **ipv4-multicast** statement.

Routers R1, R2, and R3

```
[edit protocols isis]
user@host# set traceoptions file isis size 5m world-readable
user@host# set traceoptions flag error
user@host# set topologies ipv4-multicast
```

2. Enable multicast metrics on the first SONET/SDH Interface by using the **ipv4-multicast-metric** statement.

Router R1

```
[edit protocols isis interface so-0/0/0 ]
user@R1# set level 1 metric 15
user@R1# set level 1 ipv4-multicast-metric 18
user@R1# set level 2 metric 20
user@R1# set level 2 ipv4-multicast-metric 14
```

Router R2

```
[edit protocols isis interface so-0/0/0]
user@R2# set level 1 metric 13
user@R2# set level 1 ipv4-multicast-metric 12
user@R2# set level 2 metric 29
user@R2# set level 2 ipv4-multicast-metric 23
```

Router R3

```
[edit protocols isis interface so-0/0/0]
```

```

user@R3# set level 1 metric 19
user@R3# set level 1 ipv4-multicast-metric 11
user@R3# set level 2 metric 27
user@R3# set level 2 ipv4-multicast-metric 21

```

3. Enable multicast metrics on a second sonet Interface by using the **ipv4-multicast-metric** statement.

Router R1

```

[edit protocols isis interface so-1/0/0]
user@R1# set level 1 metric 13
user@R1# set level 1 ipv4-multicast-metric 12
user@R1# set level 2 metric 29
user@R1# set level 2 ipv4-multicast-metric 23

```

Router R2

```

[edit protocols isis interface so-1/0/0]
user@R2# set level 1 metric 14
user@R2# set level 1 ipv4-multicast-metric 18
user@R2# set level 2 metric 32
user@R2# set level 2 ipv4-multicast-metric 26

```

Router R3

```

[edit protocols isis interface so-1/0/0]
user@R3# set level 1 metric 16
user@R3# set level 1 ipv4-multicast-metric 26
user@R3# set level 2 metric 30
user@R3# set level 2 ipv4-multicast-metric 20

```

4. Disable the out-of-band management port, fxp0.

Routers R1, R2, and R3

```

[edit protocols isis]
user@host# set interface fxp0.0 disable

```

5. If you are done configuring the routers, commit the configuration.

Routers R1, R2, and R3

```

[edit]
user@host# commit

```

Results From configuration mode, confirm your configuration by using the **show protocols isis** statement. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Router R1

```

user@R1# show protocols isis

traceoptions {
  file isis size 5m world-readable;
  flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
  level 1 {

```

```
        metric 15;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 20;
        ipv4-multicast-metric 14;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface fxp0.0 {
    disable;
}
```

Router R2

user@R2# show protocols isis

```
traceoptions {
    file isis size 5m world-readable;
    flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 14;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 32;
        ipv4-multicast-metric 26;
    }
}
interface fxp0.0 {
    disable;
}
```

Router R3

user@R3# show protocols isis

```
traceoptions {
    file isis size 5m world-readable;
    flag error;
```

```

}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 19;
        ipv4-multicast-metric 11;
    }
    level 2 {
        metric 27;
        ipv4-multicast-metric 21;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 16;
        ipv4-multicast-metric 26;
    }
    level 2 {
        metric 30;
        ipv4-multicast-metric 20;
    }
}
interface fxp0.0 {
    disable;
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1, R2, and R3 on page 3675](#)
- [Verifying That IS-IS Is Configured on page 3677](#)
- [Verifying the Configured Multicast Metric Values on page 3679](#)
- [Verifying the Configuration of the Multicast Topology on page 3680](#)

Verifying the Connection Between Routers R1, R2, and R3

Purpose Make sure that Routers R1, R2, and R3 are connected to each other.

Action Ping the other two routers from any router, to check the connectivity between the three routers as per the network topology.

```
user@R1> ping 10.0.3.9
```

```

PING 10.0.3.9 (10.0.3.9): 56 data bytes
64 bytes from 10.0.3.9: icmp_seq=0 ttl=64 time=1.299 ms
64 bytes from 10.0.3.9: icmp_seq=1 ttl=64 time=52.304 ms
64 bytes from 10.0.3.9: icmp_seq=2 ttl=64 time=1.271 ms
64 bytes from 10.0.3.9: icmp_seq=3 ttl=64 time=1.343 ms
64 bytes from 10.0.3.9: icmp_seq=4 ttl=64 time=1.434 ms
64 bytes from 10.0.3.9: icmp_seq=5 ttl=64 time=1.306 ms
^C
--- 10.0.3.9 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.271/9.826/52.304/18.997 ms

```

```
user@R1> ping 10.0.3.10
```

```
PING 10.0.3.10 (10.0.3.10): 56 data bytes
64 bytes from 10.0.3.10: icmp_seq=0 ttl=64 time=1.431 ms
64 bytes from 10.0.3.10: icmp_seq=1 ttl=64 time=1.296 ms
64 bytes from 10.0.3.10: icmp_seq=2 ttl=64 time=1.887 ms
^C
--- 10.0.3.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.296/1.538/1.887/0.253 ms
```

```
user@R2> ping 10.0.2.9
```

```
PING 10.0.2.9 (10.0.2.9): 56 data bytes
64 bytes from 10.0.2.9: icmp_seq=0 ttl=64 time=1.365 ms
64 bytes from 10.0.2.9: icmp_seq=1 ttl=64 time=1.813 ms
64 bytes from 10.0.2.9: icmp_seq=2 ttl=64 time=1.290 ms
^C
--- 10.0.2.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.290/1.489/1.813/0.231 ms
```

```
user@R2> ping 10.0.2.10
```

```
PING 10.0.2.10 (10.0.2.10): 56 data bytes
64 bytes from 10.0.2.10: icmp_seq=0 ttl=63 time=1.318 ms
64 bytes from 10.0.2.10: icmp_seq=1 ttl=63 time=1.394 ms
64 bytes from 10.0.2.10: icmp_seq=2 ttl=63 time=1.366 ms
64 bytes from 10.0.2.10: icmp_seq=3 ttl=63 time=1.305 ms
^C
--- 10.0.2.10 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.305/1.346/1.394/0.036 ms
```

```
user@R3> ping 10.0.1.10
```

```
PING 10.0.1.10 (10.0.1.10): 56 data bytes
64 bytes from 10.0.1.10: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.1.10: icmp_seq=1 ttl=63 time=1.418 ms
64 bytes from 10.0.1.10: icmp_seq=2 ttl=63 time=1.277 ms
^C
--- 10.0.1.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.277/1.337/1.418/0.059 ms
```

```
user@R3> ping 10.0.1.9
```

```
PING 10.0.1.9 (10.0.1.9): 56 data bytes
64 bytes from 10.0.1.9: icmp_seq=0 ttl=64 time=1.381 ms
64 bytes from 10.0.1.9: icmp_seq=1 ttl=64 time=1.499 ms
64 bytes from 10.0.1.9: icmp_seq=2 ttl=64 time=1.300 ms
64 bytes from 10.0.1.9: icmp_seq=3 ttl=64 time=1.397 ms
^C
--- 10.0.1.9 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.300/1.394/1.499/0.071 ms
```

Meaning Routers R1, R2, and R3 have a peer relationship with each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on Routers R1, R2, and R3, and that they are adjacent to each other.

Action Use the `show isis adjacency detail` command to check the adjacency between the routers.

Router R1

```
user@R1> show isis adjacency detail
```

R2

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R2

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:58 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R3

```
Interface: so-1/0/0, Level: 1, State: Up, Expires in 7 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

R3

```
Interface: so-1/0/0, Level: 2, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

Router R2

```
user@R2> show isis adjacency detail
```

R1

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 20 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9
```

R1

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 26 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
```

Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9

R3

Interface: so-1/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.10

R3

Interface: so-1/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.10

Router R3

user@R3> show isis adjacency detail

R2

Interface: so-0/0/0, Level: 1, State: Up, Expires in 18 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.9

R2

Interface: so-0/0/0, Level: 2, State: Up, Expires in 22 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.9

R1

Interface: so-1/0/0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.9

R1

Interface: so-1/0/0, Level: 2, State: Up, Expires in 19 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.9

Meaning IS-IS is configured on Routers R1, R2, and R3, and they are adjacent to each other.

Verifying the Configured Multicast Metric Values

Purpose Make sure that the SPF calculations are accurate as per the configured multicast metric values on Routers R1, R2, and R3.

Action Use the `show isis spf results` command to check the SPF calculations for the network.

Router R1

```
user@R1> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  28         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  18         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  17         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  0
      4 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  40         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  22         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  14         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R1.00  0
      4 nodes
```

Router R2

```
user@R2> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  29         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  18         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  12         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  12
R2.00  0
      5 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  45         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  26         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  23         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  23
R2.00  0
      5 nodes
```

Router R3

```
user@R3> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  26
R1.00  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.02  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.00  11         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R3.03  11
```

```

R3.00 0
      6 nodes

IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R2.02 34 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R2.00 21 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R3.03 21
R1.00 20 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R3.02 20
R3.00 0
      6 nodes

```

Meaning The configured multicast metric values are used in SPF calculations for the IS-IS network.

Verifying the Configuration of the Multicast Topology

Purpose Make sure that the multicast topology is configured on Routers R1, R2, and R3.

Action Use the **show isis database detail** command to verify the multicast topology configuration on the routers.

Router R1

```
user@R1> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```

R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 663 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up

```

```

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 883 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
  IPv4 Unicast IS neighbor: R3.03 Metric: 14
  IPv4 Multicast IS neighbor: R2.02 Metric: 12
  IPv4 Multicast IS neighbor: R3.03 Metric: 18
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 14 Internal Up

```

```

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 913 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0

```

```

R3.00-00 Sequence: 0x13c, Checksum: 0xc8de, Lifetime: 488 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 16
  IPv4 Unicast IS neighbor: R3.03 Metric: 19
  IPv4 Multicast IS neighbor: R3.02 Metric: 26
  IPv4 Multicast IS neighbor: R3.03 Metric: 11
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 16 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 19 Internal Up

```

```

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 625 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

```
R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 714 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

IS-IS level 2 link-state database:

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 816 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 20
  IPv4 Unicast IS neighbor: R3.02 Metric: 31
  IPv4 Multicast IS neighbor: R2.02 Metric: 14
  IPv4 Multicast IS neighbor: R3.02 Metric: 22
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 20 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 29
  IPv4 Unicast IS neighbor: R3.03 Metric: 32
  IPv4 Multicast IS neighbor: R2.02 Metric: 23
  IPv4 Multicast IS neighbor: R3.03 Metric: 26
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 29 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 28 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 32 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 805 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 30
  IPv4 Unicast IS neighbor: R3.03 Metric: 27
  IPv4 Multicast IS neighbor: R3.02 Metric: 20
  IPv4 Multicast IS neighbor: R3.03 Metric: 21
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 30 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 27 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

```
R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

Router R2

```
user@R2> show isis database detail
```

IS-IS level 1 link-state database:

```
R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 524 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 748 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
```

```

IPV4 Unicast IS neighbor: R3.03      Metric:      14
IPV4 Multicast IS neighbor: R2.02     Metric:      12
IPV4 Multicast IS neighbor: R3.03     Metric:      18
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      13 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      14 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 777 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1102 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      16
IPV4 Unicast IS neighbor: R3.03      Metric:      19
IPV4 Multicast IS neighbor: R3.02     Metric:      26
IPV4 Multicast IS neighbor: R3.03     Metric:      11
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      16 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 488 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 577 secs
IPV4 Unicast IS neighbor: R2.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

IS-IS level 2 link-state database:

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 676 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      20
IPV4 Unicast IS neighbor: R3.02      Metric:      31
IPV4 Multicast IS neighbor: R2.02     Metric:      14
IPV4 Multicast IS neighbor: R3.02     Metric:      22
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      20 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      29 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      29
IPV4 Unicast IS neighbor: R3.03      Metric:      32
IPV4 Multicast IS neighbor: R2.02     Metric:      23
IPV4 Multicast IS neighbor: R3.03     Metric:      26
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      29 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      28 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 667 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      30
IPV4 Unicast IS neighbor: R3.03      Metric:      27
IPV4 Multicast IS neighbor: R3.02     Metric:      20
IPV4 Multicast IS neighbor: R3.03     Metric:      21
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      30 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 707 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0

```

```
IPV4 Unicast IS neighbor: R3.00    Metric:      0
```

```
R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 707 secs
```

```
IPV4 Unicast IS neighbor: R2.00    Metric:      0
```

```
IPV4 Unicast IS neighbor: R3.00    Metric:      0
```

Router R3

```
user@R3> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R1.00-00 Sequence: 0x143, Checksum: 0xb08, Lifetime: 1155 secs
```

```
IPV4 Unicast IS neighbor: R2.02    Metric:      15
```

```
IPV4 Unicast IS neighbor: R3.02    Metric:      15
```

```
IPV4 Multicast IS neighbor: R2.02   Metric:      18
```

```
IPV4 Multicast IS neighbor: R3.02   Metric:      17
```

```
IP IPV4 Unicast prefix: 10.0.1.8/30 Metric:      15 Internal Up
```

```
IP IPV4 Unicast prefix: 10.0.2.8/30 Metric:      15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 687 secs
```

```
IPV4 Unicast IS neighbor: R2.02    Metric:      13
```

```
IPV4 Unicast IS neighbor: R3.03    Metric:      14
```

```
IPV4 Multicast IS neighbor: R2.02   Metric:      12
```

```
IPV4 Multicast IS neighbor: R3.03   Metric:      18
```

```
IP IPV4 Unicast prefix: 10.0.1.8/30 Metric:      13 Internal Up
```

```
IP IPV4 Unicast prefix: 10.0.3.8/30 Metric:      14 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 716 secs
```

```
IPV4 Unicast IS neighbor: R1.00    Metric:      0
```

```
IPV4 Unicast IS neighbor: R2.00    Metric:      0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1044 secs
```

```
IPV4 Unicast IS neighbor: R3.02    Metric:      16
```

```
IPV4 Unicast IS neighbor: R3.03    Metric:      19
```

```
IPV4 Multicast IS neighbor: R3.02   Metric:      26
```

```
IPV4 Multicast IS neighbor: R3.03   Metric:      11
```

```
IP IPV4 Unicast prefix: 10.0.2.8/30 Metric:      16 Internal Up
```

```
IP IPV4 Unicast prefix: 10.0.3.8/30 Metric:      19 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 430 secs
```

```
IPV4 Unicast IS neighbor: R1.00    Metric:      0
```

```
IPV4 Unicast IS neighbor: R3.00    Metric:      0
```

```
R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 519 secs
```

```
IPV4 Unicast IS neighbor: R2.00    Metric:      0
```

```
IPV4 Unicast IS neighbor: R3.00    Metric:      0
```

```
IS-IS level 2 link-state database:
```

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 617 secs
```

```
IPV4 Unicast IS neighbor: R2.02    Metric:      20
```

```
IPV4 Unicast IS neighbor: R3.02    Metric:      31
```

```
IPV4 Multicast IS neighbor: R2.02   Metric:      14
```

```
IPV4 Multicast IS neighbor: R3.02   Metric:      22
```

```
IP IPV4 Unicast prefix: 10.0.1.8/30 Metric:      20 Internal Up
```

```
IP IPV4 Unicast prefix: 10.0.2.8/30 Metric:      31 Internal Up
```

```
IP IPV4 Unicast prefix: 10.0.3.8/30 Metric:      29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 769 secs
```

```
IPV4 Unicast IS neighbor: R2.02    Metric:      29
```

```
IPV4 Unicast IS neighbor: R3.03    Metric:      32
```

```

IPv4 Multicast IS neighbor: R2.02    Metric:      23
IPv4 Multicast IS neighbor: R3.03    Metric:      26
IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      29 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      28 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 769 secs
IPv4 Unicast IS neighbor: R1.00      Metric:      0
IPv4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 610 secs
IPv4 Unicast IS neighbor: R3.02      Metric:      30
IPv4 Unicast IS neighbor: R3.03      Metric:      27
IPv4 Multicast IS neighbor: R3.02     Metric:      20
IPv4 Multicast IS neighbor: R3.03     Metric:      21
IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      31 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      30 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 649 secs
IPv4 Unicast IS neighbor: R1.00      Metric:      0
IPv4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 649 secs
IPv4 Unicast IS neighbor: R2.00      Metric:      0
IPv4 Unicast IS neighbor: R3.00      Metric:      0

```

Meaning Multicast topology is configured on Routers R1, R2, and R3.

Related Documentation

- [Example: Configuring Multitopology Routing Based on a Multicast Source](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies](#)

Example: Configuring IS-IS for CLNS

- [Understanding IS-IS for CLNS on page 3684](#)
- [Example: Configuring IS-IS for CLNS on page 3684](#)

Understanding IS-IS for CLNS

IS-IS extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable the IS-IS protocol on all interfaces.

- [Requirements on page 3685](#)
- [Overview on page 3685](#)

- [Configuration on page 3685](#)
- [Verification on page 3686](#)

Requirements

Before you begin, configure the network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing instance called `aaaa`, enable IS-IS on all interfaces, define the BGP export policy name (`dist-bgp`), family (`ISO`), and protocol (`BGP`), and apply the export policy to IS-IS.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS for CLNS:

1. Enable CLNS routing.

```
[edit routing-instances aaaa]
user@host# set protocols isis clns-routing
```
2. Enable IS-IS on all interfaces.

```
[edit routing-instances aaaa]
user@host# set protocols isis interface all
```
3. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network.

```
[edit routing-instances aaaa]
user@host# set protocols isis no-ipv4-routing no-ipv6-routing
```
4. Define the BGP export policy name, family, and protocol.

```
[edit policy-options]
user@host# set policy-statement dist-bgp from family iso protocol bgp
```
5. Define the action for the export policy.

```
[edit policy-options]
user@host# set policy-statement dist-bgp then accept
```

6. Apply the export policy to IS-IS.

```
[edit routing-instances aaaa]  
user@host# set protocols isis export dist-bgp
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-instances  
aaaa {  
  protocols {  
    isis {  
      export dist-bgp;  
      no-ipv4-routing;  
      no-ipv6-routing;  
      clns-routing;  
      interface all;  
    }  
  }  
}  
  
user@host# show policy-options  
policy-statement dist-bgp {  
  from {  
    family iso;  
    protocol bgp;  
  }  
  then accept;  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the ISO Routes on page 3686](#)
- [Checking the SPF Calculations on page 3686](#)

Verifying the ISO Routes

Purpose Verify that the expected ISO routes are displayed in the IS-IS routing table.

Action From operational mode, enter the [show isis route](#) command.

Checking the SPF Calculations

Purpose Display information about IS-IS shortest-path-first (SPF) calculations.

Action From operational mode, enter the **show isis spf** command.

Example: Configuring IS-IS Designated Routers

- [Understanding IS-IS Designated Routers on page 3687](#)
- [Example: Configuring Designated Router Election Priority for IS-IS on page 3687](#)

Understanding IS-IS Designated Routers

A router advertises its priority to become a designated router in its hello packets. On all multiaccess networks (physical networks that support the attachment of more than two routers, such as Ethernet networks), IS-IS uses the advertised priorities to elect a designated router for the network. This router is responsible for sending network link-state advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area. The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127, which you configure on the IS-IS interface. The router with the highest priority becomes the designated router for the area (Level 1, Level 2, or both), also configured on the IS-IS interface. If routers in the network have the same priority, then the router with the highest MAC address is elected as the designated router. By default, routers have a priority value of 64.

Example: Configuring Designated Router Election Priority for IS-IS

This example shows how to configure the designated router election priority for IS-IS.

Before you begin:

- Configure network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Enable IS-IS on the interfaces. See [“Example: Configuring IS-IS” on page 3633](#).

In this example, you configure the priority for logical interface ge-0/0/1.0 to be 100 and the level number to be 1. If this interface has the highest priority value, the router becomes the designated router for the Level 1 area.

To configure a designated router election priority for IS-IS:

```
[edit]
user@host# set protocols isis interface ge-0/0/1.0 level 1 priority 100
```

Related Documentation

- [Example: Configuring IS-IS](#)

Example: Enabling Packet Checksums on IS-IS Interfaces

This example shows how to enable packet checksums for IS-IS interfaces.

- [Requirements on page 3687](#)
- [Overview on page 3688](#)
- [Configuration on page 3688](#)
- [Verification on page 3689](#)

Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3633](#) for information about the sample IS-IS configuration.

Overview


Junos OS supports IS-IS checksums as documented in RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

IS-IS protocol data units (PDUs) include link-state PDUs, complete sequence number PDUs (CSNPs), partial sequence number PDUs (PSNPs), and IS-IS hello (IIH) packets. These PDUs can be corrupt due to faulty implementations of Layer 2 hardware or lack of checksums on a specific network technology. Corruption of length or type, length, and value (TLV) fields can lead to the generation of extensive numbers of empty link-state PDUs in the receiving node. Because authentication is not a replacement for a checksum mechanism, you might want to enable the optional checksum TLV on your IS-IS interfaces.

The checksum cannot be enabled with MD5 hello authentication on the same interface.

Figure 118 on page 3688 shows the topology used in this example.

Figure 118: IS-IS Checksum Topology



This example describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1  set protocols isis traceoptions file isis
            set protocols isis traceoptions flag all
            set protocols isis interface fe-1/2/0.1 checksum
  
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS checksums:

1. Enable checksums.


```

[edit protocols isis interface fe-1/2/0.1]
user@R1# set checksum
      
```
2. (Optional) Enable tracing for tracking checksum operations.


```

[edit protocols isis traceoptions]
user@R1# set file isis
user@R1# set flag all
      
```

Results From configuration mode, confirm your configuration by entering the **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
isis {
  traceoptions {
    file isis;
    flag all;
  }
  interface fe-1/2/0.1 {
    checksum;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying Checksums

Purpose Verify that checksums are performed.

Action From operational mode, enter the **show log isis | match checksum** command.

```
user@R1> show log isis | match checksum
```

```
May 31 16:47:39.513267      sequence 0x49 checksum 0x8e64
May 31 16:47:39.513394      sequence 0x4e checksum 0x34b3
May 31 16:47:39.513517      sequence 0x50 checksum 0x9dcb
May 31 16:47:46.563781      sequence 0x45 checksum 0x7e1a
May 31 16:47:46.563970      sequence 0x46 checksum 0x226d
May 31 16:47:46.564104      sequence 0x52 checksum 0x99cd
May 31 16:47:46.581087      sequence 0x49 checksum 0x8e64
May 31 16:47:46.581222      sequence 0x4e checksum 0x34b3
May 31 16:47:46.581353      sequence 0x50 checksum 0x9dcb
May 31 16:47:55.799090      sequence 0x45 checksum 0x7e1a
May 31 16:47:55.799223      sequence 0x46 checksum 0x226d
May 31 16:47:55.799347      sequence 0x52 checksum 0x99cd
May 31 16:47:55.818255      sequence 0x49 checksum 0x8e64
May 31 16:47:55.818473      sequence 0x4e checksum 0x34b3
May 31 16:47:55.818606      sequence 0x50 checksum 0x9dcb
May 31 16:48:03.455816      sequence 0x49 checksum 0x8e64
May 31 16:48:03.455973      sequence 0x4e checksum 0x34b3
```

Meaning The output shows that checksum information is captured in the IS-IS trace log file.

Related Documentation

- *Understanding Checksums on IS-IS Interfaces*

Configuration Tasks

- [Configuring IS-IS Authentication on page 3690](#)
- [Configuring Authentication Without Network-Wide Deployment on page 3691](#)

Configuring IS-IS Authentication

All IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in the autonomous system (AS) routing. By default, IS-IS authentication is disabled on the routing device.

To configure IS-IS authentication, you must define an authentication password and specify the authentication type.

You can configure one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet. Simple authentication is included for compatibility with existing IS-IS implementations. However, we recommend that you do *not* use this authentication method because it is insecure (the text can be “sniffed”).



CAUTION: A simple password that exceeds 254 characters is truncated.

- HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving routing device uses an authentication key (password) to verify the packet.

You can also configure more fine-grained interface-level authentication for hello packets.

To enable authentication and specify an authentication method, include the **authentication-type** statement, specifying the **simple** or **md5** authentication type:

authentication-type *authentication*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure a password, include the **authentication-key** statement. The authentication password for all routing devices in a domain must be the same.

authentication-key *key*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure hitless authentication key rollover, include the **authentication-key-chain (Protocols IS-IS)** statement.

The password can contain up to 255 characters. If you include spaces, enclose all characters in quotation marks (“ ”).

If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces that are shared with a Junos OS implementation.

Authentication of hello packets, partial sequence number PDU (PSNP), and complete sequence number PDU (CSNP) can be suppressed to enable interoperability with the routing software of different vendors. Different vendors handle authentication in various ways, and suppressing authentication for different PDU types might be the simplest way to allow compatibility within the same network.

To configure IS-IS to generate authenticated packets, but not to check the authentication on received packets, include the **no-authentication-check** statement:

```
no-authentication-check;
```

To suppress authentication of IS-IS hello packets, include the **no-hello-authentication** statement:

```
no-hello-authentication;
```

To suppress authentication of PSNPs, include the **no-psnp-authentication** statement:

```
no-psnp-authentication;
```

To suppress authentication of CSNPs, include the **no-csnp-authentication** statement:

```
no-csnp-authentication;
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.



NOTE: The **authentication** and the **no-authentication** statements must be configured at the same hierarchy level. Configuring authentication at the [edit protocols isis interface *interface-name*] hierarchy level and configuring **no-authentication** at the [edit protocols isis] hierarchy level has no effect.

Related Documentation

- [Configuring Authentication Without Network-Wide Deployment on page 3691](#)

Configuring Authentication Without Network-Wide Deployment

To allow the use of authentication without requiring network-wide deployment, include the **loose-authentication-check** statement:

```
loose-authentication-check;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Related Documentation

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS](#)

Configuration Statements

- [authentication-key \(Protocols IS-IS\) on page 3694](#)
- [authentication-key-chain \(Protocols IS-IS\) on page 3695](#)
- [authentication-type \(Protocols IS-IS\) on page 3696](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 3697](#)
- [checksum \(Protocols IS-IS\) on page 3699](#)
- [csnp-interval on page 3700](#)
- [disable \(Protocols IS-IS\) on page 3701](#)
- [export \(Protocols IS-IS\) on page 3702](#)
- [external-preference \(Protocols IS-IS\) on page 3703](#)
- [family \(Protocols IS-IS\) on page 3704](#)
- [hello-authentication-key on page 3705](#)
- [hello-authentication-key-chain on page 3706](#)
- [hello-authentication-type on page 3707](#)
- [hello-interval \(Protocols IS-IS\) on page 3708](#)
- [hello-padding on page 3709](#)
- [hold-time \(Protocols IS-IS\) on page 3711](#)
- [ignore-attached-bit on page 3712](#)
- [interface \(Protocols IS-IS\) on page 3713](#)
- [ipv4-multicast on page 3715](#)
- [ipv4-multicast-metric on page 3716](#)
- [ipv6-multicast on page 3716](#)
- [ipv6-multicast-metric on page 3717](#)
- [ipv6-unicast on page 3718](#)
- [ipv6-unicast-metric on page 3719](#)
- [isis on page 3720](#)
- [level \(Global IS-IS\) on page 3721](#)
- [loose-authentication-check on page 3722](#)
- [lsp-interval on page 3723](#)
- [lsp-lifetime on page 3724](#)
- [max-areas on page 3725](#)
- [mesh-group \(Protocols IS-IS\) on page 3726](#)
- [metric \(Protocols IS-IS\) on page 3727](#)
- [no-adjacency-holddown on page 3728](#)
- [no-authentication-check on page 3729](#)
- [no-csnp-authentication on page 3729](#)

- [no-hello-authentication](#) on page 3730
- [no-ipv4-multicast](#) on page 3730
- [no-ipv4-routing](#) on page 3731
- [no-ipv6-multicast](#) on page 3732
- [no-ipv6-routing](#) on page 3733
- [no-ipv6-unicast](#) on page 3734
- [no-psnp-authentication](#) on page 3734
- [no-unicast-topology](#) on page 3735
- [overload \(Protocols IS-IS\)](#) on page 3736
- [passive \(Protocols IS-IS\)](#) on page 3739
- [point-to-point](#) on page 3740
- [preference \(Protocols IS-IS\)](#) on page 3741
- [prefix-export-limit \(Protocols IS-IS\)](#) on page 3742
- [priority \(Protocols IS-IS\)](#) on page 3743
- [reference-bandwidth \(Protocols IS-IS\)](#) on page 3744
- [rib-group \(Protocols IS-IS\)](#) on page 3745
- [topologies \(Protocols IS-IS\)](#) on page 3746
- [traceoptions \(Protocols IS-IS\)](#) on page 3747
- [traffic-engineering \(Protocols IS-IS\)](#) on page 3750
- [wide-metrics-only](#) on page 3753

authentication-key (Protocols IS-IS)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis <i>level level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <i>level level-number</i>], [edit protocols isis <i>level level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis <i>level level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the authentication-type statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p>
Default	If you do not include this statement and the authentication-type statement, IS-IS authentication is disabled.
Options	key —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").



CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

authentication-key-chain (Protocols IS-IS)

Syntax	authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Apply and enable an authentication keychain to the routing device.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647 • Example: Configuring Route Authentication for BGP on page 3430 • Example: Configuring BFD Authentication for Static Routes on page 2848 • Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols • Understanding Hitless Authentication Key Rollover for IS-IS on page 3631

authentication-type (Protocols IS-IS)

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the authentication-key statement.
Default	If you do not include this statement and the authentication-key statement, IS-IS authentication is disabled.
Options	<i>authentication</i> —Authentication scheme: <ul style="list-style-type: none">• md5—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.• simple—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>• authentication-key on page 3694• no-authentication-check on page 3729

bfd-liveness-detection (Protocols IS-IS)

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. detection-time threshold and transmit-interval threshold options added in Junos OS Release 8.2. Support for logical systems introduced in Junos OS Release 8.3. no-adaptation statement introduced in Junos OS Release 9.0. authentication algorithm, authentication key-chain, and authentication loose-check options introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure bidirectional failure detection timers and authentication.
Options	<p>authentication algorithm <i>algorithm-name</i>—Configure the algorithm used to authenticate the specified BFD session: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>authentication key-chain <i>key-chain-name</i>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p> <p>authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p>

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000

multiplier *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BFD for IS-IS</i>• <i>Example: Configuring BFD Authentication for IS-IS</i>
------------------------------	--

checksum (Protocols IS-IS)

Syntax	checksum;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable checksums for packets on this interface. Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i> . The checksum cannot be enabled with MD5 hello authentication on the same interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i>


csnp-interval

Syntax	<code>csnp-interval (seconds disable);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds. To protect against link-state PDU flooding, we recommend modifying the default interval.</p> <p>To modify the CSNP interval, include the csnp-interval statement.</p> <p>To configure the interface not to send any CSNPs, specify the disable option.</p>
Default	By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds.
Options	<p>disable—Do not send CSNPs on this interface.</p> <p>seconds—Number of seconds between the sending of CSNPs.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i>

disable (Protocols IS-IS)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering],</p> <p>[edit protocols isis],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the [edit protocols isis traffic-engineering] hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances routing-instance-name protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Default	<p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which family iso is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Multi-Level IS-IS on page 3639 • IS-IS Overview on page 3625

export (Protocols IS-IS)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis],</code> <code>[edit protocols isis],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p>
<div> NOTE: For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</div>	
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Redistributing OSPF Routes into IS-IS</i>• <i>Example: Configuring an IS-IS Default Route Policy on Logical Systems</i>

external-preference (Protocols IS-IS)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the preference of external routes.
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Route Preferences Overview</i> • <i>Example: Redistributing OSPF Routes into IS-IS</i> • <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i> • preference on page 3741

family (Protocols IS-IS)

Syntax	<pre>family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering], [edit protocols isis traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3. Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.
Options	<p>inet—IPv4 address family</p> <p>inet6—IPv6 address family</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">•


hello-authentication-key

Syntax	<code>hello-authentication-key password;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the hello-authentication-type statement.
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>password—Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 3694 • authentication-type on page 3696 • hello-authentication-type on page 3707

hello-authentication-key-chain

Syntax	hello-authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Apply an authentication keychain to the IS-IS interface.
Options	<i>key-chain-name</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

hello-authentication-type

Syntax	hello-authentication-type (md5 simple);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the hello-authentication-key statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</p> </div> </div>	
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>md5—Specifies Message Digest 5 as the packet verification type.</p> <p>simple—Specifies simple authentication as the packet verification type.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 3694 • authentication-type on page 3696 • hello-authentication-key on page 3705

hello-interval (Protocols IS-IS)

Syntax	<code>hello-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the hello-interval value to 1.</p>
Options	<i>seconds</i> —Frequency of transmission for hello packets. Range: 1 through 20,000 seconds Default: 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>hold-time</i>

hello-padding

Syntax	hello-padding (adaptive disable loose strict);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none"> Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the adaptive option to configure enough padding to establish an adjacency to neighbors. Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the disable option to accommodate interfaces that support less than the default packet size of 1492 bytes. Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the loose option to configure enough padding to initialize an adjacency to neighbors.

- **Strict padding**—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

Options **adaptive**—Configure padding until the neighbor adjacency is established and active.

disable—Disable padding on all types of interfaces for all adjacency states.

loose—Configure padding until the state of the adjacency is initialized.

strict—Configure padding for all adjacency states.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • *Example: Configuring IS-IS*

hold-time (Protocols IS-IS)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p>
Options	<p>seconds—Hold-time value, in seconds.</p> <p>Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p>Default: 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS</i> • hello-interval on page 3708

ignore-attached-bit

Syntax	ignore-attached-bit;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The ignore-attached-bit statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p>
Default	The ignore-attached-bit statement is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">•

interface (Protocols IS-IS)

```

Syntax interface (all | interface-name) {
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
    }
    checksum;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-adjacency-holddown;
    no-ipv4-multicast;
    no-ipv6-multicast;
    no-ipv6-unicast;
    no-unicast-topology;
    passive;
    point-to-point;
    level level-number {
        disable;
        hello-authentication-key key;
        hello-authentication-key-chain key-chain-name;
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric metric;
        ipv6-multicast-metric metric;
        ipv6-unicast-metric metric;
        metric metric;
        passive;
        priority number;
        te-metric metric;
    }
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure interface-specific IS-IS properties. To configure more than one interface, include the interface statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Options	<p>all—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p>interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS</i>• <i>Example: Configuring Multi-Level IS-IS</i>

ipv4-multicast

Syntax	ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure alternate IPv4 multicast topologies.



NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.

Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology on page 3668

ipv4-multicast-metric

Syntax	ipv4-multicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the multicast topology metric value for the level.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 3668


ipv6-multicast

Syntax	ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 3668

ipv6-multicast-metric

Syntax	<code>ipv6-multicast-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Specify the IPv6 alternate multicast topology metric value for the level.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 0 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology on page 3668

ipv6-unicast

Syntax	ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure alternate IPv6 unicast topologies.</p> <p>This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.</p>
	<div> NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</div>
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

ipv6-unicast-metric

Syntax	<code>ipv6-unicast-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 0 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

isis

Syntax	isis { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable IS-IS routing on the routing device or for a routing instance. The isis statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.
Default	IS-IS is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS on page 3633• Example: Configuring Multi-Level IS-IS on page 3639

level (Global IS-IS)

Syntax	<pre> level <i>level-number</i> { authentication-key <i>key</i>; authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>; authentication-type <i>type</i>; disable; external-preference <i>preference</i>; no-csnp-authentication; no-hello-authentication; no-psnp-authentication; preference <i>preference</i>; wide-metrics-only; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p>
Options	<p><i>level-number</i>—IS-IS level number.</p> <p>Values: 1 or 2</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Configuring IS-IS*
 - *Example: Configuring Multi-Level IS-IS*

loose-authentication-check

Syntax	loose-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Allow the use of MD5 authentication without requiring network-wide deployment.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

lsp-interval

Syntax	<code>lsp-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the lsp-interval statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p>
Default	By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.
Options	<p>milliseconds—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.</p> <p>Range: 0 through 1000 milliseconds</p> <p>Default: 100 milliseconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i>

lsp-lifetime

Syntax	<code>lsp-lifetime seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis],</code> <code>[edit protocols isis],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a router needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The show isis overview command displays the link-state PDU lifetime.</p>
Default	By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.
Options	seconds —link-state PDU lifetime, in seconds. Range: 350 through 65,535 seconds Default: 1200 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i>• http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

max-areas

Syntax	<code>max-areas <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis]</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p>
Options	<p><i>number</i>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p>Range: 3 through 36</p> <p>Default: 3</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring Multi-Level IS-IS</i>

mesh-group (Protocols IS-IS)

Syntax	mesh-group (blocked <i>value</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an interface to be part of a mesh group, which is a set of fully connected nodes.</p> <p>A <i>mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.</p> <p>To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.</p>
Options	<p>blocked—Configure the interface so that it does not flood link-state PDUs.</p> <p>value—Number that identifies the mesh group.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$; 32 bits are allocated to identify a mesh group)</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i>

metric (Protocols IS-IS)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the metric value for the level.

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ($2^{24} - 1$) if you are using wide metrics.

Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.

[Table 294 on page 3727](#) depicts IS-IS route export metric default values.

Table 294: Default Metric Values for Routes Exported into IS-IS

Protocol Used for Learning the Route	Default Metric Value
Direct	10
Static	Same as reported by the protocol used for exporting the route
Aggregate	10
Generate	10
RIP	Same as reported by the protocol used for exporting the route
OSPF	Same as reported by the protocol used for exporting the route
BGP	10

The default metric values behavior can be customized by using routing policies.

Options	<i>metric</i> —Metric value. Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)
----------------	--

Default: 10 (for all interfaces except lo0), 0 (for the lo0 interface)

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i>• <i>te-metric</i>• wide-metrics-only on page 3753

no-adjacency-holddown

Syntax	no-adjacency-holddown;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Disable the hold-down timer for IS-IS adjacencies.</p> <p>A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• hold-time on page 3711

no-authentication-check

Syntax	no-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hello-authentication-type on page 3707

no-csnp-authentication

Syntax	no-csnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number PDU (CSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • csnp-interval on page 3700


no-hello-authentication

Syntax	no-hello-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• hello-authentication-type on page 3707

no-ipv4-multicast

Syntax	no-ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from IPv4 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 3668

no-ipv4-routing

Syntax	no-ipv4-routing;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets. • The routing device does not advertise any IPv4 addresses in Junos OS hello packets. • The routing device does not calculate any IPv4 routes.
	<p> NOTE: Note: Even when no-ipv4-routing is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

no-ipv6-multicast

Syntax	no-ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology on page 3668

no-ipv6-routing

Syntax	no-ipv6-routing;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Disable IP version 6 (IPv6) routing. Disabling IPv6 routing has the following results: <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets. • The routing device does not advertise any IPv6 addresses in Junos OS hello packets. • The routing device does not calculate any IPv6 routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

no-ipv6-unicast

Syntax	no-ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

no-psnp-authentication

Syntax	no-psnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on partial sequence number PDU (PSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IS-IS Authentication on page 3690

no-unicast-topology

Syntax	no-unicast-topology;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from the IPv4 unicast topologies.
Default	IPv4 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology on page 3668

overload (Protocols IS-IS)

Syntax	<pre>overload { advertise-high-metrics; allow-route-leaking; timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none">1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared.2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared.3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value. <p>In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are</p>

advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.

Options **advertise-high-metrics**—Advertise maximum link metrics in NLRIs instead of setting the overload bit.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode.

When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

allow-route-leaking—Enable leaking of route information into the network even if the overload bit is set.



NOTE: The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.

timeout seconds—Number of seconds at which the overloading is reset.


Range: 60 through 1800 seconds

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • *Example: Configuring IS-IS*

passive (Protocols IS-IS)

Syntax	<code>passive;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the disable statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p>
<div>  <p>NOTE: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the passive statement was used in the configuration of the interface.</p> </div>	
<p>If neither passive mode nor the family iso option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the family iso option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p>	
Default	By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Multi-Level IS-IS on page 3639 • <code>disable</code>

point-to-point

Syntax	point-to-point;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the point-to-point statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.</p> <p>The point-to-point statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• IS-IS Overview on page 3625• Understanding IS-IS Designated Routers on page 3687• <i>Example: Configuring Synchronization Between IS-IS and LDP</i>

preference (Protocols IS-IS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the preference statement (for internal routes) or the external-preference statement.</p>
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Route Preferences Overview</i> • <i>Example: Redistributing OSPF Routes into IS-IS</i> • <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i> • external-preference on page 3703

prefix-export-limit (Protocols IS-IS)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</code> <code>[edit protocols isis level level-number],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the prefix-export-limit statement. The prefix-export-limit statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the clear isis overload command.</p> <p>The show isis overview command displays the prefix export limit when it is configured.</p>
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i>• <i>Example: Redistributing OSPF Routes into IS-IS</i>

priority (Protocols IS-IS)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p>
Options	<p><i>number</i>—Priority value.</p> <p>Range: 0 through 127</p> <p>Default: 64</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Designated Routers on page 3686

reference-bandwidth (Protocols IS-IS)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</code> <code>[edit protocols <i>isis</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p>
Options	<p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bps</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS</i>• http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

rib-group (Protocols IS-IS)

Syntax	<pre> rib-group { inet <i>group-name</i>; inet6 <i>group-name</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p>
Options	<p><i>group-name</i>—Name of the routing table group.</p> <p>inet—Install IPv4 IS-IS routes.</p> <p>inet6—Install IPv6 IS-IS routes.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Understanding Multiprotocol BGP</i>

topologies (Protocols IS-IS)

Syntax	<pre>topologies { ipv4-multicast; ipv6-multicast; ipv6-unicast; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure alternate IS-IS topologies. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>• Example: Configuring IS-IS Multicast Topology on page 3668

traceoptions (Protocols IS-IS)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default IS-IS protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory /var/log. We recommend that you place IS-IS tracing output in the file isis-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p>

IS-IS Protocol-Specific Tracing Flags

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i> • <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i> • <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i>

traffic-engineering (Protocols IS-IS)

Syntax	<pre>traffic-engineering { disable; credibility-protocol-preference; family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; } multipath { lsp-equal-cost; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit protocols isis],
Release Information	Statement introduced before Junos OS Release 7.4. Support for the family statement introduced in Junos OS Release 9.3. Support for the credibility-protocol-preference statement introduced in Junos OS Release 9.4. Support for the multipath statement introduced in Junos OS Release 9.6. Support for the lsp-equal-cost statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure traffic engineering properties for IS-IS.</p> <p>IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS installs the prefix for IPv4 routes in the inet.3 routing table as well. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.</p> <p>If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the LSP as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses an LSP along the path to reach the egress router.</p> <p>In Junos OS Release 9.3 and later, IS-IS traffic engineering shortcuts support IPv6 routes. LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is for only BGP to use LSPs in its calculations. If you configure MPLS so that both BGP and interior gateway protocols use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table. IS-IS ensures that the IPv6 routes running over the IPv4 MPLS LSP are correctly de-encapsulated at the</p>

tunnel egress by pushing an extra IPv6 explicit null label between the IPv6 payload and the IPv4 transport label.

RSVP LSPs with a higher preference than IS-IS routes are not considered during the computation of traffic engineering shortcuts.

To configure IS-IS so that it uses LSPs as shortcuts when installing information in the inet.3 or inet6.3 routing table, include the following statements:

```
family inet {
  shortcuts {
    multicast-rpf-routes;
  }
}
family inet6 {
  shortcuts;
}
```

For IPv4 traffic, include the **inet** statement. For IPv6 traffic, include the **inet6** statement.

To configure load balancing across multiple LSPs, include the **multipath** statement.

When traffic engineering shortcuts are used, RSVP first looks at the **metric2** value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load-balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the **lsp-equal-cost** statement to retain the equal cost path information in the routing table.

```
multipath {
  lsp-equal-cost;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

Default IS-IS traffic engineering support is enabled.

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database. To disable this support, and to disable IS-IS shortcuts if they are configured, include the **disable** statement.

Options **credibility-protocol-preference**—Specify that IS-IS should use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value. By default, the traffic engineering database prefers IS-IS routes even when the routes of another IGP are configured with a lower, that is, more preferred value. Use this statement to override this default behavior.

The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.



NOTE: This feature is also supported for OSPFv2.

lsp-equal-cost—Configure LSPs to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation. When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default IS-IS behavior, include the **lsp-equal-cost** statement for load balancing, so that the equal cost path information is retained in the routing table.

multipath—Enable load balancing for multiple LSPs.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Enabling OSPF Traffic Engineering Support on page 3910](#)
- [Example: Enabling IS-IS Traffic Engineering Support](#)
- [traffic-engineering \(OSPF\) on page 4026](#)

wide-metrics-only

Syntax	wide-metrics-only;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis. Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ($2^{24} - 1$). To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the wide-metrics-only statement.
Default	By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i> • <i>te-metric</i>

CHAPTER 41

Administration

- [Operational Commands on page 3755](#)

Operational Commands

- [clear isis adjacency](#)
- [clear isis database](#)
- [clear isis overload](#)
- [clear isis statistics](#)
- [show isis adjacency](#)
- [show isis authentication](#)
- [show isis database](#)
- [show isis hostname](#)
- [show isis interface](#)
- [show isis overview](#)
- [show isis route](#)
- [show isis statistics](#)

clear isis adjacency

List of Syntax	Syntax on page 3756 Syntax (EX Series Switches and QFX Series) on page 3756
Syntax	<pre>clear isis adjacency <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor></pre>
Syntax (EX Series Switches and QFX Series)	<pre>clear isis adjacency <instance <i>instance-name</i>> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove entries from the IS-IS adjacency database.
Options	<p>none—Remove all entries from the adjacency database.</p> <p>instance <i>instance-name</i>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Clear all adjacencies for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) Clear adjacencies for the specified neighbor only.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis adjacency on page 3764
List of Sample Output	clear isis adjacency on page 3756
Output Fields	See show isis adjacency for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
```


so-1/0/0.0	karakul	3	Up	26
so-1/1/3.0	1921.6800.5080	3	Up	23
so-5/0/0.0	1921.6800.5080	3	Up	19

user@host> clear isis adjacency karakul

user@host> show isis adjacency

IS-IS adjacency database:

Interface	System	L State	Hold (secs)	SNPA
so-1/0/0.0	karakul	3 Initializing	26	
so-1/1/3.0	1921.6800.5080	3 Up	24	
so-5/0/0.0	1921.6800.5080	3 Up	21	

clear isis database

List of Syntax	Syntax on page 3758 Syntax (EX Series Switches and QFX Series) on page 3758
Syntax	clear isis database <entries> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <purge>
Syntax (EX Series Switches and QFX Series)	clear isis database <entries> <instance <i>instance-name</i> > <purge>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. purge option introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information. You can also use purge with any of the options to initiate a network-wide purge of link-state PDUs rather than the local deletion of entries from the IS-IS link-state database.
	<div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <div style="display: flex; align-items: center;"> <div> <p>CAUTION: In a production network, the purge command option might cause short-term network-wide traffic disruptions.</p> </div> </div> </div>
Options	<p>none—Remove all entries from the IS-IS link-state database for all routing instances.</p> <p>entries—(Optional) Name of the database entry.</p> <p>instance <i>instance-name</i>—(Optional) Clear all entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>purge—(Optional) Discard all entries in the IS-IS link-state database.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> show isis database on page 3770
List of Sample Output	clear isis database on page 3759
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd             1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x19   0xe92c             1134
badlands.00-00        0x16   0x1454             985
carlsbad.00-00        0x33   0x220b            1015
ranier.00-00          0x2e   0xfc31             1007
1921.6800.5066.00-00  0x11   0x7313             566
1921.6800.5067.00-00  0x14   0xd9d4             939
  6 LSPs
```

```
user@host> clear isis database
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
```

clear isis overload

List of Syntax	Syntax on page 3760 Syntax (EX Series Switches and QFX Series) on page 3760
Syntax	<code>clear isis overload</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>clear isis overload</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p>
Options	<p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 3770
List of Sample Output	clear isis overload on page 3760
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                               Sequence Checksum Lifetime Attributes
```

```

pro3-c.00-00          0x4   0x10db    1185 L1 L2 Overload

```

```

  1 LSPs

```

```

IS-IS level 2 link-state database:

```

```

LSP ID          Sequence Checksum Lifetime Attributes
pro3-c.00-00    0x5   0x429f    1185 L1 L2 Overload

```

```

pro2-a.00-00    0x91e  0x2589     874 L1 L2

```

```

pro2-a.02-00    0x1   0xcabc     874 L1 L2

```

```

  3 LSPs

```

```

user@host> clear isis overload

```

```

user@host> show isis database

```

```

IS-IS level 1 link-state database:

```

```

LSP ID          Sequence Checksum Lifetime Attributes
pro3-c.00-00    0xa   0x429e    1183 L1 L2

```

```

  1 LSPs

```

```

IS-IS level 2 link-state database:

```

```

LSP ID          Sequence Checksum Lifetime Attributes
pro3-c.00-00    0xc   0x9c39    1183 L1 L2
pro2-a.00-00    0x91e  0x2589     783 L1 L2
pro2-a.02-00    0x1   0xcabc     783 L1 L2

```

```

  3 LSPs

```

clear isis statistics

List of Syntax	Syntax on page 3762 Syntax (EX Series Switches and QFX Series) on page 3762
Syntax	<code>clear isis statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>clear isis statistics</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set statistics about IS-IS traffic to zero.
Options	none —Set IS-IS traffic statistics to zero for all routing instances. instance <i>instance-name</i> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show isis statistics on page 3791
List of Sample Output	clear isis statistics on page 3762
Output Fields	See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12793	12793	0	8666	719
IIH	116751	116751	0	118834	0
CSNP	203956	203956	0	204080	0
PSNP	7356	7350	6	8635	0
Unknown	0	0	0	0	0
Totals	340856	340850	6	340215	719

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 1064
Fragments rebuilt: 1087
LSP regenerations: 436
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	0	0	0	0	0
IIH	3	3	0	3	0
CSNP	2	2	0	4	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	5	5	0	7	0

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 0
Fragments rebuilt: 0
LSP regenerations: 0
Purges initiated: 0

show isis adjacency

List of Syntax	Syntax on page 3764 Syntax (EX Series Switches and QFX Series) on page 3764
Syntax	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about IS-IS neighbors.
Options	<p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p><i>system id</i>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p><i>brief detail extensive</i>—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p><i>instance instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis adjacency on page 3756
List of Sample Output	show isis adjacency on page 3766 show isis adjacency brief on page 3766 show isis adjacency detail on page 3767 show isis adjacency extensive on page 3767
Output Fields	Table 295 on page 3765 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.

Table 295: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels
System	System identifier (sysid), displayed as a name, if possible.	brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected .	All levels
Hold (secs)	Remaining hold time of the adjacency.	brief
SNPA	Subnetwork point of attachment (MAC address of the next hop).	brief
Expires in	How long until the adjacency expires, in seconds.	detail
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up .	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No .	detail extensive
Adjacency advertisement: Advertise	This router has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the router's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 295: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition log	<p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. 	extensive

Sample Output

show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up         23

```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 3766](#).

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
Transition log:
When           State      Event      Down reason
Wed Nov  8 21:24:25  Up        Seenself
```

show isis authentication

List of Syntax	Syntax on page 3768 Syntax (EX Series Switches and QFX Series) on page 3768
Syntax	<pre>show isis authentication <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis authentication <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for hitless authentication key rollover introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display information about IS-IS authentication.
Options	<p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis authentication on page 3769 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 3769
Output Fields	<p>Table 296 on page 3768 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.</p>

Table 296: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.
IIH Auth	<p>IS-IS Hello (IIH) packet authentication type.</p> <p>Displays the name of the active keychain if hitless authentication key rollover is configured.</p>
CSN Auth	Complete sequence number authentication type.

Table 296: show isis authentication Output Fields (*continued*)

Field Name	Field Description
PSN Auth	Partial sequence number authentication type.
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple    Simple    Simple
                   2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

show isis database

List of Syntax	Syntax on page 3770 Syntax (EX Series Switches and QFX Series) on page 3770
Syntax	<pre>show isis database <system-id> <brief detail extensive> <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis database <system-id> <brief detail extensive> <level (1 2)> <instance <i>instance-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the entries in the IS-IS link-state database, which contains data about PDU packets.
Options	<p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p>level (1 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis database on page 3758
List of Sample Output	show isis database on page 3772 show isis database brief on page 3773 show isis database detail on page 3773

[show isis database extensive on page 3773](#)

Output Fields [Table 297 on page 3771](#) describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 297: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always IS-IS for this command.	All levels
level	Level of intermediate system: <ul style="list-style-type: none"> • 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • 2—Intermediate system routes between areas and toward other ASs. 	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only).	none brief
# LSPs	Total number of link-state PDUs in the specified link-state database.	none brief
IP prefix	Prefix advertised by this link-state PDU.	detail extensive
IS neighbor	IS-IS neighbor of the advertising system.	detail extensive
ES neighbor	(J Series routers only) An ES-IS neighbor of the advertising system.	detail extensive
IP prefix	IPv4 prefix advertised by this link-state PDU.	detail extensive
V6 prefix	IPv6 prefix advertised by this link-state PDU.	detail extensive
Metric	Metric of the prefix or neighbor.	detail extensive
Header	<ul style="list-style-type: none"> • LSP ID—Link state PDU identifier of the header. • Length—Header length. • Allocated Length—Amount of length available for the header. • Router ID—Address of the local routing device. • Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. 	extensive

Table 297: show isis database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet	<ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. 	extensive
TLVs	<ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, and IP address. 	extensive

Sample Output

show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3    0x3167    1057 L1 L2
camaro.00-00          0x5    0x770e    1091 L1 L2
ranier.00-00          0x4    0xaa95    1091 L1 L2
glacier.00-00         0x4    0x206f    1089 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x3    0x87a2    1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6    0x8d6b    1096 L1 L2
camaro.00-00          0x9    0x877b    1101 L1 L2
ranier.00-00          0x8    0x855d    1103 L1 L2
glacier.00-00         0x7    0xf892    1098 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x6    0x562     1105 L1 L2
  6 LSPs

```


show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 3772](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02          Metric:      10
  ES neighbor: 0015.0015.0015          Metric:      10 Down
  ES neighbor: 0025.0025.0025          Metric:      10 Down
  ES neighbor: 0030.0030.0030          Metric:      10 Down
  ES neighbor: 0040.0040.0040          Metric:      10 Down
  ES neighbor: sisira                    Metric:       0
  IP prefix: 1.0.0.0/24                 Metric:      10 External Down
  IP prefix: 3.0.0.0/24                 Metric:      10 External Down
  IP prefix: 4.0.0.0/24                 Metric:      10 External Down
  IP prefix: 5.0.0.0/24                 Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32             Metric:      10 External Down
  IP prefix: 25.25.25.25/32             Metric:      10 External Down
  IP prefix: 30.30.30.30/32             Metric:      10 External Down
  IP prefix: 40.40.40.40/32             Metric:      10 External Down
  IP prefix: 60.60.60.60/32             Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02          Metric:      10
  IP prefix: 1.0.0.0/24                 Metric:      10 External Down
  IP prefix: 3.0.0.0/24                 Metric:      10 External Down
  IP prefix: 4.0.0.0/24                 Metric:      10 External Down
  IP prefix: 5.0.0.0/24                 Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32             Metric:      10 External Down
  IP prefix: 25.25.25.25/32             Metric:      10 External Down
  IP prefix: 30.30.30.30/32             Metric:      10 External Down
  IP prefix: 40.40.40.40/32             Metric:      10 External Down
  IP prefix: 50.50.50.50/32             Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32             Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

show isis database extensive

```
user@host> show isis database logical-system CE3 sisira.00-00 extensive
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
```

```

IS neighbor: hemantha-CE3.02           Metric:      10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
ES neighbor: 0015.0015.0015           Metric:      10 Down
ES neighbor: 0025.0025.0025           Metric:      10 Down
ES neighbor: 0030.0030.0030           Metric:      10 Down
ES neighbor: 0040.0040.0040           Metric:      10 Down
ES neighbor: sisira                    Metric:      0
IP prefix: 1.0.0.0/24                  Metric:      10 External Down
IP prefix: 3.0.0.0/24                  Metric:      10 External Down
IP prefix: 4.0.0.0/24                  Metric:      10 External Down
IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
IP prefix: 15.15.15.15/32              Metric:      10 External Down
IP prefix: 25.25.25.25/32              Metric:      10 External Down
IP prefix: 30.30.30.30/32              Metric:      10 External Down
IP prefix: 40.40.40.40/32              Metric:      10 External Down
IP prefix: 60.60.60.60/32              Metric:      0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
Allocated length: 336 bytes, Router ID: 0.0.0.0
Remaining lifetime: 970 secs, Level: 1, Interface: 333
Estimated free bytes: 144, Actual free bytes: 0
Aging timer expires in: 970 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
ES neighbor TLV: Internal, Metric: default 0, Up
  ES: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0040.0040.0040
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0025.0025.0025
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0015.0015.0015
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0030.0030.0030
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down

```

```

IP extended prefix: 30.30.30.30/32 metric 10 down
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
No queued transmissions

```

IS-IS level 2 link-state database:

```

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 988 secs
IS neighbor: hemantha-CE3.02 Metric: 10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
IP prefix: 1.0.0.0/24 Metric: 10 External Down
IP prefix: 3.0.0.0/24 Metric: 10 External Down
IP prefix: 4.0.0.0/24 Metric: 10 External Down
IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
IP prefix: 15.15.15.15/32 Metric: 10 External Down
IP prefix: 25.25.25.25/32 Metric: 10 External Down
IP prefix: 30.30.30.30/32 Metric: 10 External Down
IP prefix: 40.40.40.40/32 Metric: 10 External Down
IP prefix: 50.50.50.50/32 Metric: 10 Internal Up
IP prefix: 60.60.60.60/32 Metric: 0 Internal Up
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
Metric: 0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 427 bytes
Allocated length: 427 bytes, Router ID: 0.0.0.0
Remaining lifetime: 988 secs, Level: 2, Interface: 333
Estimated free bytes: 130, Actual free bytes: 0
Aging timer expires in: 988 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 427 bytes, Lifetime : 1198 secs
Checksum: 0x69ac, Sequence: 0x13, Attributes: 0x3 L1 L2
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down

```

```
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down
IP extended prefix: 30.30.30.30/32 metric 10 down
ISO prefix-neighbor TLV: Internal, Metric: default 0, Up
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP prefix: 50.50.50.50/32, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
IP extended prefix: 50.50.50.50/32 metric 10 up
No queued transmissions
```

show isis hostname

List of Syntax	Syntax on page 3777 Syntax (EX Series Switches and QFX Series) on page 3777
Syntax	show isis hostname <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis hostname
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS hostname database information.
Options	none —Display IS-IS hostname database information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show isis hostname on page 3777
Output Fields	Table 298 on page 3777 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear.

Table 298: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. <ul style="list-style-type: none"> Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>. Static—Hostname mapping configured by user.

Sample Output

show isis hostname

```


user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
Type
Dynamic

```

1921.6800.4202 isis2
1921.6800.4203 isis3

Static
Dynamic

show isis interface

List of Syntax	Syntax on page 3779 Syntax (EX Series Switches and QFX Series) on page 3779
Syntax	<pre>show isis interface <brief detail extensive> <interface-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis interface <brief detail extensive> <interface-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Display status information about IS-IS-enabled interfaces.</p>
<div>  <p>NOTE: If the configured metric for an IS-IS level is above 63, and the wide-metrics-only statement is not configured, the show isis interface detail command and the show isis interface extensive command display 63 as the metric value for that level. Configure the wide-metrics-only statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The show isis interface command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p> </div>	
Options	<p>none—Display standard information about all IS-IS-enabled interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i>
List of Sample Output	show isis interface on page 3781 show isis interface brief on page 3782 show isis interface detail on page 3782 show isis interface extensive on page 3782 show isis interface extensive (With LDP) on page 3782

Output Fields Table 299 on page 3780 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

Table 299: show isis interface Output Fields

Field Name	Field Description	Level of Output
<i>interface-name</i>	Name of the interface.	detail
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	detail
Index	Interface index assigned by the Junos OS kernel.	detail
State	Internal implementation information.	detail
Circuit id	Circuit identifier.	detail
Circuit type	Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	detail
LSP interval	Interval between link-state PDUs sent from the interface.	detail
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	detail extensive
Sysid	System identifier.	detail
Interface	Interface through which the adjacency is made.	none brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	All levels
CirID	Circuit identifier.	none brief
Level 1 DR	Level 1 designated intermediate system.	none brief
Level 2 DR	Level 2 designated intermediate system.	none brief
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.	none brief
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	detail extensive

Table 299: show isis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	detail extensive
Adjacencies	Number of adjacencies established on this interface.	detail
Priority	Priority value for this interface.	detail
Metric	Metric value for this interface.	detail
Hello(s) / Hello Interval	Interface's hello interval.	detail extensive
Hold(s) / Hold Time	Interface's hold time.	detail extensive
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routers attached to the network.	detail
Hello padding	Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. 	extensive
LDP sync state	Current LDP synchronization state: in sync , in holddown , or not supported .	extensive
reason	Reason for being in the LDP sync state.	extensive
config holdtime	Configured value of the hold timer.	extensive
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	extensive

Sample Output

show isis interface

```

user@host> show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0     3   0x1 Point to Point    Point to Point    10/10
1o0.0          0   0x1 Passive          Passive          0/0

```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 3781](#).

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             1         64      10     9.000      27
    2             1         64      10     9.000      27
1o0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0         64       0  Passive
    2             0         64       0  Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s, Loose Hello padding
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
1o0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
  Level 2
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
```

show isis interface extensive (With LDP)

```
user@host> show isis interface extensive
IS-IS interface database:
so-1/1/2.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:01:28, reason: LDP up during config
  config holddtime: 20 seconds
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 11
    Hello Interval: 9.000 s, Hold Time: 27 s
    IPV4 MulticastMetric: 10
    IPV6 UnicastMetric: 10
```


show isis overview

Syntax	show isis overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis overview <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS overview information.
Options	none —Display standard overview information about IS-IS for all routing instances. instance <i>instance-name</i> —(Optional) Display overview information for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show isis overview on page 3786
Output Fields	Table 300 on page 3784 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 300: show isis overview Output Fields

Field Name	Field Description
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: enabled or disabled .
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Attached bit evaluation	Attached bit capability: enabled or disabled .
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.

Table 300: show isis overview Output Fields (*continued*)

Field Name	Field Description
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: enabled or disabled .
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: enabled or disabled .
Restart	Graceful restart capability: enabled or disabled .
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: enabled or disabled .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
CLNS is enabled	(J Series routers only) OSI CLNP capability is enabled.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the prefix-export-limit statement.
Prefix export count	Number of prefixes exported.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.

Sample Output

show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Adjacency holddown: disabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
```

show isis route

List of Syntax	Syntax on page 3787 Syntax (EX Series Switches and QFX Series) on page 3787
Syntax	<pre>show isis route <destination> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis route <destination> <inet inet6> <instance <i>instance-name</i>> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display the routes in the IS-IS routing table.
Options	<p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p><i>destination</i>—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p>
Required Privilege Level	view
List of Sample Output	show isis route logical-system on page 3788 show isis route (CLNS) on page 3788 show isis route on page 3789
Output Fields	<p>Table 301 on page 3788 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p>

Table 301: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table           Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2    11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2    11    60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table   Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2    11    20 int  gr-0/2/0.0  h

```

show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table           Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0   1    10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa

```



```

0/0
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
  1      10      20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                   IPV4 camaro
10.255.71.238/32            2    13        20   int  so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
10.255.71.239/32            2    13        20   int  so-6/0/0.0             IPV4 olympic
                               ae0.0                   IPV4 camaro
10.255.71.242/32            2    13        10   int  as0.0                   IPV4 glacier
10.255.71.243/32            2    13        10   int  so-6/0/0.0             IPV4 olympic
12.13.0.0/30                 2    13        20   int  so-6/0/0.0             IPV4 olympic
12.15.0.0/30                 2    13        20   int  so-6/0/0.0             IPV4 olympic
13.15.0.0/30                 2    13        30   int  ae0.0                   IPV4 camaro
                               so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
13.16.0.0/30                 2    13        25   int  as0.0                   IPV4 glacier
14.15.0.0/30                 2    13        20   int  ae0.0                   IPV4 camaro
192.2.1.0/30                 2    13        30   int  so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
1eee::/64                    2    13        30   int  so-6/0/0.0             IPV6 olympic
                               as0.0                   IPV6 glacier
abcd::10:255:71:52/128      2    13        10   int  ae0.0                   IPV6 camaro
abcd::10:255:71:238/128     2    13        20   int  so-6/0/0.0             IPV6 olympic

```

					as0.0	IPV6 glacier
abcd::10:255:71:239/128	2	13	20	int	so-6/0/0.0	IPV6 olympic
					ae0.0	IPV6 camaro
abcd::10:255:71:242/128	2	13	10	int	as0.0	IPV6 glacier
abcd::10:255:71:243/128	2	13	10	int	so-6/0/0.0	IPV6 olympic

show isis statistics

List of Syntax	Syntax on page 3791 Syntax (EX Series Switches and QFX Series) on page 3791
Syntax	<pre>show isis statistics <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis statistics <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display statistics about IS-IS traffic.
Options	<p>none—Display IS-IS traffic statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear isis statistics on page 3762
List of Sample Output	show isis statistics on page 3793
Output Fields	<p>Table 302 on page 3792 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 302: show isis statistics Output Fields

Field Name	Field Description
PDU type	<p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12227	12227	0	8184	683
IIH	113808	113808	0	115817	0
CSNP	198868	198868	0	198934	0
PSNP	6985	6979	6	8274	0
Unknown	0	0	0	0	0
Totals	331888	331882	6	331209	683

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:          0 Drops:          0  
LSP queue length:          0 Drops:          0
```

```
SPF runs:                  1014  
Fragments rebuilt:        1038  
LSP regenerations:        425  
Purges initiated:         0
```


PART 14

Open Shortest Path First

- [Overview on page 3797](#)
- [Configuration on page 3809](#)
- [Administration on page 4031](#)

CHAPTER 42

Overview

- [OSPF Overview on page 3797](#)

OSPF Overview

- [OSPF Overview on page 3798](#)
- [OSPF Areas and Router Functionality Overview on page 3803](#)
- [Packets Overview on page 3805](#)
- [OSPF External Metrics Overview on page 3808](#)

OSPF Overview

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.



NOTE: On SRX Series devices, when only one link-protection is configured under the OSPF interface, the device does not install an alternative route in the forwarding table. When the per-packet load-balancing is enabled as a workaround, the device does not observe both the OSPF metric and sending the traffic through both the interfaces.

An OSPF AS can consist of a single area, or it can be subdivided into multiple areas. In a single-area OSPF network topology, each router maintains a database that describes the topology of the AS. Link-state information for each router is flooded throughout the AS. In a multiarea OSPF topology, each router maintains a database that describes the topology of its area, and link-state information for each router is flooded throughout that area. All routers maintain summarized topologies of other areas within an AS. Within each area, OSPF routers have identical topological databases. When the AS or area topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPFv2 protocol exchanges can be authenticated. OSPFv3 relies on IPsec to provide this functionality. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. A single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.



NOTE: By default, Junos OS is compatible with RFC 1583, *OSPF Version 2*. In Junos OS Release 8.5 and later, you can disable compatibility with RFC 1583 by including the `no-rfc-1583` statement. For more information, see [“Example: Disabling OSPFv2 Compatibility with RFC 1583” on page 3836](#).

This topic describes the following information:

- [OSPF Default Route Preference Values on page 3800](#)
- [OSPF Routing Algorithm on page 3800](#)
- [OSPF Three-Way Handshake on page 3801](#)
- [OSPF Version 3 on page 3802](#)

OSPF Default Route Preference Values

The Junos OS routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference value is from 0 through 4,294,967,295 ($2^{32} - 1$), with a lower value indicating a more preferred route. [Table 303 on page 3800](#) lists the default preference values for OSPF.

Table 303: Default Route Preference Values for OSPF

How Route Is Learned	Default Preference	Statement to Modify Default Preference
OSPF internal route	10	OSPF <code>preference</code>
OSPF AS external routes	150	OSPF <code>external-preference</code>

OSPF Routing Algorithm

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated

router form adjacencies with other routing devices.) Adjacencies determine the distribution of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.


The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

OSPF Three-Way Handshake

OSPF creates a topology map by flooding LSAs across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in [Figure 119 on page 3801](#).

Figure 119: OSPF Three-Way Handshake



In [Figure 119 on page 3801](#), Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

OSPF Version 3

OSPFv3 is a modified version of OSPF that supports IP version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.
- The protocol runs per link rather than per subnet.
- Router and network link-state advertisements (LSAs) do not carry prefix information.
- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.
- Flooding scopes are as follows:
 - Link-local
 - Area
 - AS
- Link-local addresses are used for all neighbor exchanges except virtual links.
- Authentication is removed. The IPv6 authentication header relies on the IP layer.
- The packet format has changed as follows:
 - Version number 2 is now version number 3.
 - The **db** option field has been expanded to 24 bits.
 - Authentication information has been removed.
 - Hello messages do not have address information.
 - Two new option bits are included: **R** and **V6**.
- Type 3 summary LSAs have been renamed *inter-area-prefix-LSAs*.
- Type 4 summary LSAs have been renamed *inter-area-router-LSAs*.

Related Documentation

- [Understanding OSPF Areas and Backbone Areas on page 3814](#)
- [OSPF Configuration Overview](#)
- [OSPF Version 3 for IPv6](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 3836](#)

OSPF Areas and Router Functionality Overview

In OSPF, a single autonomous system (AS) can be divided into smaller groups called *areas*. This reduces the number of link-state advertisements (LSAs) and other OSPF overhead traffic sent on the network, and it reduces the size of the topology database that each router must maintain. The routing devices that participate in OSPF routing perform one or more functions based on their location in the network.

This topic describes the following OSPF area types and routing device functions:

- [Areas on page 3803](#)
- [Area Border Routers on page 3803](#)
- [Backbone Areas on page 3803](#)
- [AS Boundary Routers on page 3804](#)
- [Backbone Router on page 3804](#)
- [Internal Router on page 3804](#)
- [Stub Areas on page 3804](#)
- [Not-So-Stubby Areas on page 3804](#)
- [Transit Areas on page 3805](#)

Areas

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Also, routing within the area is determined only by the area's topology, providing the area with some protection from bad routing data.

All routing devices within an area have identical topology databases.

Area Border Routers

Routing devices that belong to more than one area and connect one or more OSPF areas to the backbone area are called *area border routers* (ABRs). At least one interface is within the backbone while another interface is in another area. ABRs also maintain a separate topological database for each area to which they are connected.

Backbone Areas

An OSPF *backbone area* consists of all networks in area ID 0.0.0.0, their attached routing devices, and all ABRs. The backbone itself does not have any ABRs. The backbone distributes routing information between areas. The backbone is simply another area, so the terminology and rules of areas apply: a routing device that is directly connected to the backbone is an internal router on the backbone, and the backbone's topology is hidden from the other areas in the AS.

The routing devices that make up the backbone must be physically contiguous. If they are not, you must configure *virtual links* to create the appearance of backbone connectivity. You can create virtual links between any two ABRs that have an interface to a common nonbackbone area. OSPF treats two routing devices joined by a virtual link as if they were connected to an unnumbered point-to-point network.

AS Boundary Routers

Routing devices that exchange routing information with routing devices in non-OSPF networks are called *AS boundary routers*. They advertise externally learned routes throughout the OSPF AS. Depending on the location of the AS boundary router in the network, it can be an ABR, a backbone router, or an internal router (with the exception of stub areas). Internal routers within a stub area cannot be an AS boundary router because stub areas cannot contain any Type 5 LSAs.

Routing devices within the area where the AS boundary router resides know the path to that AS boundary router. Any routing device outside the area only knows the path to the nearest ABR that is in the same area where the AS boundary router resides.

Backbone Router

Backbone routers are routing devices that have one or more interfaces connected to the OSPF backbone area (area ID 0.0.0.0).

Internal Router

Routing devices that connect to only one OSPF area are called *internal routers*. All interfaces on internal routers are directly connected to networks within a single area.

Stub Areas

Stub areas are areas through which or into which AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological databases and therefore the amount of memory required on the internal routers in the stub area.

Routing devices within a stub area rely on the default routes originated by the area's ABR to reach external AS destinations. You must configure the **default-metric** option on the ABR before it advertises a default route. Once configured, the ABR advertises a default route in place of the external routes that are not being advertised within the stub area, so that routing devices in the stub area can reach destinations outside the area.

The following restrictions apply to stub areas: you cannot create a virtual link through a stub area, a stub area cannot contain an AS boundary router, the backbone cannot be a stub area, and you cannot configure an area as both a stub area and a not-so-stubby area.

Not-So-Stubby Areas

An OSPF stub area has no external routes in it, so you cannot redistribute from another protocol into a stub area. A *not-so-stubby area* (NSSA) allows external routes to be

flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

The following restriction applies to NSSAs: you cannot configure an area as both a stub area and an NSSA.

Transit Areas

Transit areas are used to pass traffic from one adjacent area to the backbone (or to another area if the backbone is more than two hops away from an area). The traffic does not originate in, nor is it destined for, the transit area.

Related Documentation

- [OSPF Overview on page 3798](#)
- [Packets Overview on page 3805](#)
- [OSPF Configuration Overview](#)
- [Understanding OSPF Areas and Backbone Areas on page 3814](#)
- [Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 3821](#)

Packets Overview

There are several types of link-state advertisement (LSA) packets.

This topic describes the following information:

- [OSPF Packet Header on page 3805](#)
- [Hello Packets on page 3806](#)
- [Database Description Packets on page 3806](#)
- [Link-State Request Packets on page 3806](#)
- [Link-State Update Packets on page 3806](#)
- [Link-State Acknowledgment Packets on page 3807](#)
- [Link-State Advertisement Packet Types on page 3807](#)

OSPF Packet Header

All OSPFv2 packets have a common 24-byte header, and OSPFv3 packets have a common 16-byte header, that contains all information necessary to determine whether OSPF should accept the packet. The header consists of the following fields:

- Version number—The current OSPF version number. This can be either **2** or **3**.
- Type—Type of OSPF packet.
- Packet length—Length of the packet, in bytes, including the header.
- Router ID—IP address of the router from which the packet originated.
- Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0.

- Checksum—Fletcher checksum.
- Authentication—(OSPFv2 only) Authentication scheme and authentication information.
- Instance ID—(OSPFv3 only) Identifier used when there are multiple OSPFv3 realms configured on a link.

Hello Packets

Routers periodically send hello packets on all interfaces, including virtual links, to establish and maintain neighbor relationships. Hello packets are multicast on physical networks that have a multicast or broadcast capability, which enables dynamic discovery of neighboring routers. (On nonbroadcast networks, dynamic neighbor discovery is not possible, so you must configure all neighbors statically as described in [“Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network” on page 3841.](#))

Hello packets consist of the OSPF header plus the following fields:

- Network mask—(OSPFv2 only) Network mask associated with the interface.
- Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval.
- Options—Optional capabilities of the router.
- Router priority—The router’s priority to become the designated router.
- Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be down. All routers on a shared network must use the same router dead interval.
- Designated router—IP address of the designated router.
- Backup designated router—IP address of the backup designated router.
- Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

Database Description Packets

When initializing an adjacency, OSPF exchanges database description packets, which describe the contents of the topological database. These packets consist of the OSPF header, packet sequence number, and the link-state advertisement’s header.

Link-State Request Packets

When a router detects that portions of its topological database are out of date, it sends a link-state request packet to a neighbor requesting a precise instance of the database. These packets consist of the OSPF header plus fields that uniquely identify the database information that the router is seeking.

Link-State Update Packets

Link-state update packets carry one or more link-state advertisements one hop farther from their origin. The router multicasts (floods) these packets on physical networks that support multicast or broadcast mode. The router acknowledges all link-state update

packets and, if retransmission is necessary, sends the retransmitted advertisements unicast.

Link-state update packets consist of the OSPF header plus the following fields:

- Number of advertisements—Number of link-state advertisements included in this packet.
- Link-state advertisements—The link-state advertisements themselves.

Link-State Acknowledgment Packets

The router sends link-state acknowledgment packets in response to link-state update packets to verify that the update packets have been received successfully. A single acknowledgment packet can include responses to multiple update packets.

Link-state acknowledgment packets consist of the OSPF header plus the link-state advertisement header.

Link-State Advertisement Packet Types

Link-state request, link-state update, and link-state acknowledgment packets are used to reliably flood link-state advertisement packets. OSPF sends the following types of link-state advertisements:

- Router link advertisements—Are sent by all routers to describe the state and cost of the router's links to the area. These link-state advertisements are flooded throughout a single area only.
- Network link advertisements—Are sent by designated routers to describe all the routers attached to the network. These link-state advertisements are flooded throughout a single area only.
- Summary link advertisements—Are sent by area border routers to describe the routes that they know about in other areas. There are two types of summary link advertisements: those used when the destination is an IP network, and those used when the destination is an AS boundary router. Summary link advertisements describe interarea routes, that is, routes to destinations outside the area but within the AS. These link-state advertisements are flooded throughout the advertisement's associated areas.
- AS external link advertisement—Are sent by AS boundary routers to describe external routes that they know about. These link-state advertisements are flooded throughout the AS (except for stub areas).

Each link-state advertisement type describes a portion of the OSPF routing domain. All link-state advertisements are flooded throughout the AS.

Each link-state advertisement packet begins with a common 20-byte header.

Related Documentation

- [OSPF Overview on page 3798](#)
- [OSPF Areas and Router Functionality Overview on page 3803](#)
- [OSPF Configuration Overview](#)

- [OSPF Designated Router Overview on page 3809](#)
- [Understanding OSPFv2 Authentication](#)
- [OSPF Timers Overview on page 3875](#)

OSPF External Metrics Overview

When OSPF exports route information from external autonomous systems (ASs), it includes a cost, or *external metric*, in the route. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router. Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router. By default, OSPF uses the Type 2 external metric.

CHAPTER 43

Configuration

- [Basic OSPF Area Configuration on page 3809](#)
- [Advanced OSPF Area Configuration on page 3820](#)
- [OSPF Interface Configuration on page 3837](#)
- [OSPF Route Control Configuration on page 3852](#)
- [OSPF Fault Detection Configuration on page 3875](#)
- [OSPF Redundancy Features Configuration on page 3892](#)
- [OSPF Traffic Engineering Configuration on page 3907](#)
- [OSPF Database Protection Configuration on page 3919](#)
- [OSPF Policy Configuration on page 3921](#)
- [OSPF Monitoring Configuration on page 3954](#)
- [Configuration Statements on page 3961](#)

Basic OSPF Area Configuration

- [Examples: Configuring OSPF Designated Routers on page 3809](#)
- [Examples: Configuring OSPF Areas on page 3814](#)

Examples: Configuring OSPF Designated Routers

- [OSPF Designated Router Overview on page 3809](#)
- [Example: Configuring an OSPF Router Identifier on page 3810](#)
- [Example: Controlling OSPF Designated Router Election on page 3812](#)

OSPF Designated Router Overview

Large LANs that have many routing devices and therefore many OSPF adjacencies can produce heavy control-packet traffic as link-state advertisements (LSAs) are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers on all multiaccess networks (broadcast and nonbroadcast multiaccess [NBMA] networks types). Rather than broadcasting LSAs to all their OSPF neighbors, the routing devices send their LSAs to the designated router. Each multiaccess network has a designated router, which performs two main functions:

- Originate network link advertisements on behalf of the network.

- Establish adjacencies with all routing devices on the network, thus participating in the synchronizing of the link-state databases.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the routing device with the highest router identifier (defined by the **router-id** configuration value, which is typically the IP address of the routing device, or the loopback address) is elected the designated router. The routing device with the second highest router identifier is elected the backup designated router. If the designated router fails or loses connectivity, the backup designated router assumes its role and a new backup designated router election takes place between all the routers in the OSPF network.

OSPF uses the router identifier for two main purposes: to elect a designated router, unless you manually specify a priority value, and to identify the routing device from which a packet is originated. At designated router election, the router priorities are evaluated first, and the routing device with the highest priority is elected designated router. If router priorities tie, the routing device with the highest router identifier, which is typically the routing device's IP address, is chosen as the designated router. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

At least one routing device on each logical IP network or subnet must be eligible to be the designated router for OSPFv2. At least one routing device on each logical link must be eligible to be the designated router for OSPFv3.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router. A priority of 255 means the routing device is always the designated router.

Example: Configuring an OSPF Router Identifier

This example shows how to configure an OSPF router identifier.

- [Requirements on page 3810](#)
- [Overview on page 3811](#)
- [Configuration on page 3811](#)
- [Verification on page 3812](#)

Requirements

Before you begin:

- Identify the interfaces on the routing device that will participate in OSPF. You must enable OSPF on all interfaces within the network on which OSPF traffic is to travel.
- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

The router identifier is used by OSPF to identify the routing device from which a packet originated. Junos OS selects a router identifier according to the following set of rules:

1. By default, Junos OS selects the lowest configured physical IP address of an interface as the router identifier.
2. If a loopback interface is configured, the IP address of the loopback interface becomes the router identifier.
3. If multiple loopback interfaces are configured, the lowest loopback address becomes the router identifier.
4. If a router identifier is explicitly configured using the **router-id address** statement under the **[edit routing-options]** hierarchy level, the above three rules are ignored.



NOTE: If the router identifier is modified in a network, the link-state advertisements (LSAs) advertised by the previous router identifier are retained in the OSPF database until the LSA retransmit interval has timed out.

If the router identifier is not configured explicitly and an interface IP address is used as the router identifier, the established OSPF adjacency flaps when the interface goes down, or when it is brought back into the network. When the interface is brought back into the network, or a new interface is introduced into the network, the router identifier is selected again based on the rules stated above. Hence, it is strongly recommended that you explicitly configure the router identifier under the **[edit routing-options]** hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.



NOTE: The router identifier behavior described here holds good even when configured under **[edit routing-instances routing-instance-name routing-options]** and **[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options]** hierarchy levels.

In this example, you configure the OSPF router identifier by setting its router ID value to the IP address of the device, which is 177.162.4.24.

Configuration

CLI Quick Configuration

To quickly configure an OSPF router identifier, copy the following command and paste it into the CLI.

```
[edit]
set routing-options router-id 177.162.4.24
```

Step-by-Step Procedure

To configure an OSPF router identifier:

1. Configure the OSPF router identifier by entering the **[router-id]** configuration value.
- ```
[edit]
```

```
user@host# set routing-options router-id 177.162.4.24
```

2. If you are done configuring the device, commit the configuration.

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

### Results

Confirm your configuration by entering the **show routing-options router-id** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options router-id
router-id 177.162.4.24;
```

### Verification

After you configure the router ID and activate OSPF on the routing device, the router ID is referenced by multiple OSPF operational mode commands that you can use to monitor and troubleshoot the OSPF protocol. The router ID fields are clearly marked in the output.

### Example: Controlling OSPF Designated Router Election

---

This example shows how to control OSPF designated router election.

- [Requirements on page 3812](#)
- [Overview on page 3812](#)
- [Configuration on page 3812](#)
- [Verification on page 3813](#)

### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.

### Overview

This example shows how to control OSPF designated router election. Within the example, you set the OSPF interface to **ge-/0/0/1** and the device priority to 200. The higher the priority value, the greater likelihood the routing device will become the designated router.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPF designated router election, copy the following command and paste it into the CLI.



```
[edit]
set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

**Step-by-Step Procedure** To control OSPF designated router election:

1. Configure an OSPF interface and specify the device priority.



**NOTE:** To specify an OSPFv3 interface, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

2. If you are done configuring the device, commit the configuration.

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

### Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.3 {
 interface ge-0/0/1.0 {
 priority 200;
 }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

### Verification

Confirm that the configuration is working properly.

- [Verifying the Designated Router Election on page 3813](#)

### Verifying the Designated Router Election

**Purpose** Based on the priority you configured for a specific OSPF interface, you can confirm the address of the area's designated router. The DR ID, DR, or DR-ID field displays the address of the area's designated router. The BDR ID, BDR, or BDR-ID field displays the address of the backup designated router.

**Action** From operational mode, enter the `show ospf interface` and the `show ospf neighbor` commands for OSPFv2, and enter the `show ospf3 interface` and the `show ospf3 neighbor` commands for OSPFv3.

**Related Documentation** • [OSPF Areas and Router Functionality Overview on page 3803](#)

- [OSPF Configuration Overview](#)

## Examples: Configuring OSPF Areas


- [Understanding OSPF Areas and Backbone Areas on page 3814](#)
- [Example: Configuring a Single-Area OSPF Network on page 3815](#)
- [Example: Configuring a Multiarea OSPF Network on page 3817](#)

### Understanding OSPF Areas and Backbone Areas

OSPF networks in an autonomous system (AS) are administratively grouped into *areas*. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions similar to a network address. Within an area, the topology database contains only information about the area, link-state advertisements (LSAs) are flooded only to nodes within the area, and routes are computed only within the area. The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Subnetworks are divided into other areas, which are connected to form the whole of the main network. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The central area of an AS, called the *backbone area*, has a special function and is always assigned the area ID 0.0.0.0. (Within a simple, single-area network, this is also the ID of the area.) Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a routing device that has interfaces in more than one area. These connecting routing devices are called *area border routers* (ABRs). [Figure 120 on page 3814](#) shows an OSPF topology of three areas connected by two ABRs.

Figure 120: Multiarea OSPF Topology



Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for transmitting the traffic through the appropriate ABR to the destination area. The ABRs summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate ABR. For example, in the OSPF areas shown in [Figure 120 on page 3814](#), packets sent from Router A to Router C are automatically routed through ABR B.

Junos OS supports active backbone detection. Active backbone detection is implemented to verify that ABRs are connected to the backbone. If the connection to the backbone area is lost, then the routing device's default metric is not advertised, effectively rerouting traffic through another ABR with a valid connection to the backbone. Active backbone detection enables transit through an ABR with no active backbone connection. An ABR advertises to other routing devices that it is an ABR even if the connection to the backbone is down, so that the neighbors can consider it for interarea routes.

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate ABR and on to the remote host within the destination area.

### Example: Configuring a Single-Area OSPF Network

This example shows how to configure a single-area OSPF network.

- [Requirements on page 3815](#)
- [Overview on page 3815](#)
- [Configuration on page 3816](#)
- [Verification on page 3817](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 3810.

#### Overview


To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

In an autonomous system (AS), the backbone area is always assigned area ID 0.0.0.0 (within a simple, single-area network, this is also the ID of the area). Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an

AS. All other networks or areas in the AS must be directly connected to the backbone area by area border routers that have interfaces in more than one area. You must also create a backbone area if your network consists of multiple areas. In this example, you create the backbone area and add interfaces, such as **ge-0/0/0**, as needed to the OSPF area.

To use OSPF on the device, you must configure at least one OSPF area, such as the one shown in [Figure 121 on page 3816](#).

**Figure 121: Typical Single-Area OSPF Network Topology**



#### Configuration

##### CLI Quick Configuration

To quickly configure a single-area OSPF network, copy the following command and paste it into the CLI. You repeat this configuration for all interfaces that are part of the OSPF area.

```
[edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

##### Step-by-Step Procedure

To configure a single-area OSPF network:

1. Configure the single-area OSPF network by specifying the area ID and associated interface.



**NOTE:** For a single-area OSPFv3 network, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. If you are done configuring the device, commit the configuration.

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

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface ge-0/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Interfaces in the Area

- |                |                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured. |
| <b>Action</b>  | From operational mode, enter the <b>show ospf interface</b> command for OSPFv2, and enter the <b>show ospf3 interface</b> command for OSPFv3.                  |

### Example: Configuring a Multiarea OSPF Network

This example shows how to configure a multiarea OSPF network. To reduce traffic and topology maintenance for the devices in an OSPF autonomous system (AS), you can group the OSPF-enabled routing devices into multiple areas.

- [Requirements on page 3817](#)
- [Overview on page 3818](#)
- [Configuration on page 3818](#)
- [Verification on page 3820](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.

### Overview


To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

Each OSPF area consists of routing devices configured with the same area number. The backbone area is always assigned area ID 0.0.0.0. (All area identifiers (IDs) must be unique within an AS.) All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. In

[Figure 122 on page 3818](#), Devices A and C are in the backbone area (area 0), and Devices D and E are in area 2. Device B has a special role. This is the area border router that connects area 0 and area 2. The area border router maintains a separate topological database for each area to which it is connected.

To reduce traffic and topology maintenance for the devices in an OSPF AS, you can group them into multiple areas as shown in [Figure 122 on page 3818](#). In this example, you create the backbone area, create an additional area (area 2) and assign it unique area ID 0.0.0.2, and you configure Device B as the area border router, where interface **ge-0/0/0** participates in OSPF area 0 and interface **ge-0/0/2** participates in OSPF area 2.

**Figure 122: Typical Multiarea OSPF Network Topology**



### Configuration

**CLI Quick Configuration** To quickly configure a multiarea OSPF network, copy the following commands and paste them into the CLI. You repeat this configuration for all interfaces that are part of the OSPF area.

**Device A** [edit]  
 set protocols ospf area 0.0.0.0 interface ge-0/0/0  
 set protocols ospf area 0.0.0.0 interface ge-0/0/1

**Device C** [edit]  
 set protocols ospf area 0.0.0.0 interface ge-0/0/0

**Device B** [edit]

```
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

**Device D**      [edit]  
 set protocols ospf area 0.0.0.2 interface ge-0/0/0  
 set protocols ospf area 0.0.0.2 interface ge-0/0/2

**Device E**      [edit]  
 set protocols ospf area 0.0.0.2 interface ge-0/0/2

**Step-by-Step Procedure**      To configure a multiarea OSPF network:

1.      Configure the backbone area.



**NOTE:** For an OSPFv3 network, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

```
[edit]
user@C# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

```
[edit]
user@B# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2.      Configure an additional area for your OSPF network.

```
[edit]
user@B# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/0
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@E# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

3.      If you are done configuring the device, commit the configuration.

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

### Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show protocols ospf
area 0.0.0.0 {
 interface ge-0/0/0.0;
 interface ge-0/0/1.0;
}
```

```
user@C# show protocols ospf
area 0.0.0.0 {
 interface ge-0/0/0.0;
}

user@B# show protocols ospf
area 0.0.0.0 {
 interface ge-0/0/0.0;
}
area 0.0.0.2 {
 interface ge-0/0/2.0;
}

user@D# show protocols ospf
area 0.0.0.2 {
 interface ge-0/0/0.0;
 interface ge-0/0/2.0;
}

user@E# show protocols ospf
area 0.0.0.2 {
 interface ge-0/0/2.0;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 3820](#)

### **Verifying the Interfaces in the Area**

**Purpose** Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.

**Action** From operational mode, enter the **show ospf interface** command for OSPFv2, and enter the **show ospf3 interface** command for OSPFv3.

**Related Documentation**

- [OSPF Areas and Router Functionality Overview on page 3803](#)
- [OSPF Configuration Overview](#)

---

## **Advanced OSPF Area Configuration**

- [Examples: Configuring OSPF Stub and Not-So-Stubby Areas on page 3821](#)
- [Example: Configuring OSPF Multiarea Adjacency on page 3831](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 3835](#)




## Examples: Configuring OSPF Stub and Not-So-Stubby Areas

- [Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 3821](#)
- [Example: Configuring OSPF Stub and Totally Stubby Areas on page 3822](#)
- [Example: Configuring OSPF Not-So-Stubby Areas on page 3826](#)

### Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas

[Figure 123 on page 3821](#) shows an autonomous system (AS) across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

**Figure 123: OSPF AS Network with Stub Areas and NSSAs**



To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router (ABR) interface to the area as a stub interface, you suppress external route advertisements through the ABR. Instead, the ABR advertises a default route (through itself) in place of the external routes and generates network summary (Type 3) link-state advertisements (LSAs). Packets destined for external routes are automatically sent to the ABR, which acts as a gateway for outbound traffic and routes the traffic appropriately.



**NOTE:** You must explicitly configure the ABR to generate a default route when attached to a stub or not-so-stubby-area (NSSA). To inject a default route with a specified metric value into the area, you must configure the `default-metric` option and specify a metric value.

For example, area 0.0.0.3 in [Figure 123 on page 3821](#) is not directly connected to the outside network. All outbound traffic is routed through the ABR to the backbone and then to the

destination addresses. By designating area 0.0.0.3 as a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

A stub area that only allows routes internal to the area and restricts Type 3 LSAs from entering the stub area is often called a *totally stubby area*. You can convert area 0.0.0.3 to a totally stubby area by configuring the ABR to only advertise and allow the default route to enter into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area.



**NOTE:** If you incorrectly configure a totally stubby area, you might encounter network connectivity issues. You should have advanced knowledge of OSPF and understand your network environment before configuring totally stubby areas.

Similar to area 0.0.0.3 in [Figure 123 on page 3821](#), area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route advertisements to the area and advertise the static customer routes by designating the area an NSSA. In an NSSA, the AS boundary router generates NSSA external (Type 7) LSAs and floods them into the NSSA, where they are contained. Type 7 LSAs allow an NSSA to support the presence of AS boundary routers and their corresponding external routing information. The ABR converts Type 7 LSAs into AS external (Type 5) LSAs and leaks them to the other areas, but external routes from other areas are not advertised within the NSSA.

---

### Example: Configuring OSPF Stub and Totally Stubby Areas

---

This example shows how to configure an OSPF stub area and a totally stubby area to control the advertisement of external routes into an area.

- [Requirements on page 3822](#)
- [Overview on page 3823](#)
- [Configuration on page 3824](#)
- [Verification on page 3825](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

### Overview

The backbone area, which is 0 in [Figure 124 on page 3824](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an autonomous system (AS). All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by area border routers (ABRs) that have interfaces in more than one area.

Stub areas are areas through which or into which OSPF does not flood AS external link-state advertisements (Type 5 LSAs). You might create stub areas when much of the topology database consists of AS external advertisements and you want to minimize the size of the topology databases on the internal routers in the stub area.

The following restrictions apply to stub areas:

- You cannot create a virtual link through a stub area.
- A stub area cannot contain an AS boundary router.
- You cannot configure the backbone as a stub area.
- You cannot configure an area as both a stub area and a not-so-stubby area (NSSA).

In this example, you configure each routing device in area 7 (area ID 0.0.0.7) as a stub router and some additional settings on the ABR:


- **stub**—Specifies that this area become a stub area and not be flooded with Type 5 LSAs. You must include the **stub** statement on all routing devices that are in area 7 because this area has no external connections.
- **default-metric**—Configures the ABR to generate a default route with a specified metric into the stub area. This default route enables packet forwarding from the stub area to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to a stub. You must explicitly configure this option to generate a default route.
- **no-summaries**—(Optional) Prevents the ABR from advertising summary routes into the stub area by converting the stub area into a totally stubby area. If configured in combination with the **default-metric** statement, a totally stubby area only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area. Only the ABR requires this additional configuration because it is the only routing device within the totally stubby area that creates Type 3 LSAs used to receive and send traffic from outside of the area.

**NOTE:**

In Junos OS Release 8.5 and later, the following applies:

- A router-identifier interface that is not configured to run OSPF is no longer advertised as a stub network in OSPF LSAs.
- OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also advertises the direct route with the configured mask length, as in earlier releases.

Figure 124: OSPF Network Topology with Stub Areas and NSSAs

**Configuration****CLI Quick Configuration**

- To quickly configure an OSPF stub area, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the stub area.

[edit]

```
set protocols ospf area 0.0.0.7 stub
```

- To quickly configure the ABR to inject a default route into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

[edit]

```
set protocols ospf area 0.0.0.7 stub default-metric 10
```

- (Optional) To quickly configure the ABR to restrict all summary advertisements and allow only internal routes and default route advertisements into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

[edit]

```
set protocols ospf area 0.0.0.7 stub no-summaries
```

**Step-by-Step  
Procedure**

To configure OSPF stub areas:

1. On all routing devices in the area, configure an OSPF stub area.



**NOTE:** To specify an OSPFv3 stub area, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub
```

2. On the ABR, inject a default route into the area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub default-metric 10
```

3. (Optional) On the ABR, restrict summary LSAs from entering the area. This step converts the stub area into a totally stubby area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub no-summaries
```

4. If you are done configuring the devices, commit the configuration.

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

**Results**

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices:

```
user@host# show protocols ospf
area 0.0.0.7 {
 stub;
}
```

Configuration on the ABR (the output also includes the optional setting):

```
user@host# show protocols ospf
area 0.0.0.7 {
 stub default-metric 10 no-summaries;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 3826](#)
- [Verifying the Type of OSPF Area on page 3826](#)

### ***Verifying the Interfaces in the Area***

**Purpose** Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub as the type of OSPF area.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

### ***Verifying the Type of OSPF Area***

**Purpose** Verify that the OSPF area is a stub area. Confirm that the output displays Normal Stub as the Stub type.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

---

### **Example: Configuring OSPF Not-So-Stubby Areas**

This example shows how to configure an OSPF not-so-stubby area (NSSA) to control the advertisement of external routes into an area.

- [Requirements on page 3826](#)
- [Overview on page 3826](#)
- [Configuration on page 3828](#)
- [Verification on page 3830](#)

#### ***Requirements***

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### ***Overview***

The backbone area, which is 0 in [Figure 125 on page 3828](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by ABRs that have interfaces in more than one area.

An OSPF stub area has no external routes, so you cannot redistribute routes from another protocol into a stub area. OSPF NSSAs allow external routes to be flooded within the area.

In addition, you might have a situation when exporting Type 7 LSAs into the NSSA is unnecessary. When an AS boundary router is also an ABR with an NSSA attached, Type 7 LSAs are exported into the NSSA by default. If the ABR is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. During route redistribution, this routing device generates both Type 5 LSAs and Type 7 LSAs. You can disable exporting Type 7 LSAs into the NSSA.



**NOTE:** The following restriction applies to NSSAs: You cannot configure an area as both a stub area and an NSSA.

You configure each routing device in area 9 (area ID 0.0.0.9) with the following setting:

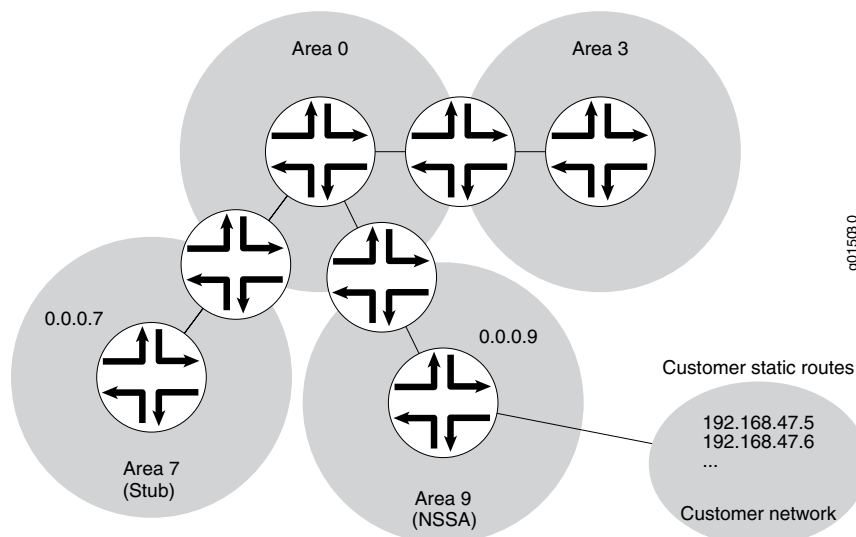
- **nssa**—Specifies an OSPF NSSA. You must include the **nssa** statement on all routing devices in area 9 because this area only has external connections to static routes.

You also configure the ABR in area 9 with the following additional settings:

- **no-summaries**—Prevents the ABR from advertising summary routes into the NSSA. If configured in combination with the **default-metric** statement, the NSSA only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into the NSSA. Only the ABR requires this additional configuration because it is the only routing device within the NSSA that creates Type 3 LSAs used to receive and send traffic from outside the area.
- **default-lsa**—Configures the ABR to generate a default route into the NSSA. In this example, you configure the following:
  - **default-metric**—Specifies that the ABR generate a default route with a specified metric into the NSSA. This default route enables packet forwarding from the NSSA to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to an NSSA. You must explicitly configure this option for the ABR to generate a default route.
  - **metric-type**—(Optional) Specifies the external metric type for the default LSA, which can be either Type 1 or Type 2. When OSPF exports route information from external ASs, it includes a cost, or external metric, in the route. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. Type 2 external metrics use only the external cost assigned by the AS boundary router. By default, OSPF uses the Type 2 external metric.
  - **type-7**—(Optional) Floods Type 7 default LSAs into the NSSA if the **no-summaries** statement is configured. By default, when the **no-summaries** statement is configured, a Type 3 LSA is injected into NSSAs for Junos OS release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the **type-7** statement.

The second example also shows the optional configuration required to disable exporting Type 7 LSAs into the NSSA by including the **no-nssa-abr** statement on the routing device that performs the functions of both an ABR and an AS boundary router.

Figure 125: OSPF Network Topology with Stub Areas and NSSAs



### Configuration

- [Configuring Routing Devices to Participate in a Not-So-Stubby-Area on page 3828](#)
- [Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas on page 3830](#)

### Configuring Routing Devices to Participate in a Not-So-Stubby-Area

**CLI Quick Configuration** To quickly configure an OSPF NSSA, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the NSSA.

```
[edit]
set protocols ospf area 0.0.0.9 nssa
```

To quickly configure an ABR that participates in an OSPF NSSA, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.9 nssa default-lsa default-metric 10
set protocols ospf area 0.0.0.9 nssa default-lsa metric-type 1
set protocols ospf area 0.0.0.9 nssa default-lsa type-7
set protocols ospf area 0.0.0.9 nssa no-summaries
```

**Step-by-Step Procedure** To configure OSPF NSSAs:

1. On all routing devices in the area, configure an OSPF NSSA.



**NOTE:** To specify an OSPFv3 NSSA area, include the **ospf3** statement at the **[edit protocols]** hierarchy level.



```
[edit]
user@host# set protocols ospf area 0.0.0.9 nssa
```

2. On the ABR, enter OSPF configuration mode and specify the NSSA area 0.0.0.9 that you already created.

```
[edit]
user@host# edit protocols ospf area 0.0.0.9 nssa
```

3. On the ABR, inject a default route into the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa default-metric 10
```

4. (Optional) On the ABR, specify the external metric type for the default route.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa metric-type 1
```

5. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa type-7
```

6. On the ABR, restrict summary LSAs from entering the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set no-summaries
```

7. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices in the area:

```
user@host# show protocols ospf
area 0.0.0.9 {
 nssa;
}
```

Configuration on the ABR. The output also includes the optional **metric-type** and **type-7** statements.

```
user@host# show protocols ospf
area 0.0.0.9 {
 nssa {
 default-lsa {
 default-metric 10;
 metric-type 1;
 type-7;
 }
 no-summaries;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

#### *Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas*

**CLI Quick Configuration** To quickly disable exporting Type 7 LSAs into the NSSA, copy the following command and paste it into the CLI. You configure this setting on an AS boundary router that is also an ABR with an NSSA area attached.

```
[edit]
set protocols ospf no-nssa-abr
```

**Step-by-Step Procedure** You can configure this setting if you have an AS boundary router that is also an ABR with an NSSA area attached.

1. Disable exporting Type 7 LSAs into the NSSA.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf no-nssa-abr
```

2. If you are done configuring the device, commit the configuration.

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

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-nssa-abr;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

#### *Verification*

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 3830](#)
- [Verifying the Type of OSPF Area on page 3831](#)
- [Verifying the Type of LSAs on page 3831](#)

#### *Verifying the Interfaces in the Area*

**Purpose** Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub NSSA as the type of OSPF area.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Verifying the Type of OSPF Area**

**Purpose** Verify that the OSPF area is a stub area. Confirm that the output displays Not so Stubby Stub as the Stub type.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

**Verifying the Type of LSAs**

**Purpose** Verify the type of LSAs that are in the area. If you disabled exporting Type 7 LSAs into an NSSA, confirm that the Type field does not include NSSA as a type of LSA.

**Action** From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

- Related Documentation**
- *Example: Configuring OSPFv3 Stub and Totally Stubby Areas*
  - [OSPF Areas and Router Functionality Overview on page 3803](#)
  - *OSPF Configuration Overview*

**Example: Configuring OSPF Multiarea Adjacency**

- [Multiarea Adjacency for OSPF on page 3831](#)
- [Example: Configuring Multiarea Adjacency for OSPF on page 3832](#)

**Multiarea Adjacency for OSPF**

An area is a set of networks and hosts within an autonomous system (AS) that have been administratively grouped together. By default, a single interface can belong to only one OSPF area. However, in some situations, you might want to configure an interface to belong to more than one area. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. For example, you can configure an interface to belong to multiple areas with a high-speed backbone link between two area border routers (ABRs) so you can create multiarea adjacencies that belong to different areas.

In Junos OS Release 9.2 and later, you can configure a logical interface to belong to more than one OSPFv2 area. Support for OSPFv3 was introduced in Junos OS Release 9.4. As defined in RFC 5185, *OSPF Multi-Area Adjacency*, the ABRs establish multiple adjacencies belonging to different areas over the same logical interface. Each multiarea adjacency is announced as a point-to-point unnumbered link in the configured area by the routers connected to the link. For each area, one of the logical interfaces is treated as primary, and the remaining interfaces that are configured for the area are designated as secondary.

Any logical interface not configured as a secondary interface for an area is treated as the primary interface for that area. A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.

### Example: Configuring Multiarea Adjacency for OSPF

---

This example shows how to configure multiarea adjacency for OSPF.

- [Requirements on page 3832](#)
- [Overview on page 3832](#)
- [Configuration on page 3833](#)
- [Verification on page 3835](#)

#### Requirements

Before you begin, plan your multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).

#### Overview

By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:

- For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.
- Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.
- Since the purpose of a secondary interface is to advertise a topological path through an OSPF area, you cannot configure a secondary interface or a primary interface with one or more secondary interfaces to be passive. Passive interfaces advertise their address, but do not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- Any logical interface not configured as a secondary interface for an area is treated as a primary interface for that area. A logical interface can be configured as the primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.
- You cannot configure the **secondary** statement with the **interface all** statement.
- You cannot configure a secondary interface by its IP address.

In this example, you configure an interface to be in two areas, creating a multiarea adjacency with a link between two ABRs: ABR R1 and ABR R2. On each ABR, area 0.0.0.1 contains the primary interface and is the primary link between the ABRs, and area 0.0.0.2 contains the secondary logical interface, which you configure by including the **secondary** statement. You configure interface **so-0/0/0** on ABR R1 and interface **so-1/0/0** on ABR R2.

**Configuration**

**CLI Quick Configuration** To quickly configure a secondary logical interface for an OSPF area, copy the following commands and paste them into the CLI.

Configuration on ABR R1:

```
[edit]
set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
set routing-options router-id 10.255.0.1
set protocols ospf area 0.0.0.1 interface so-0/0/0
set protocols ospf area 0.0.0.2 interface so-0/0/0 secondary
```

Configuration on ABR R2:

```
[edit]
set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
set routing-options router-id 10.255.0.2
set protocols ospf area 0.0.0.1 interface so-1/0/0
set protocols ospf area 0.0.0.2 interface so-1/0/0 secondary
```

**Step-by-Step Procedure** To configure a secondary logical interface:

1. Configure the device interfaces.



**NOTE:** For OSPFv3, on each interface specify the inet6 address family and include the IPv6 address.

```
[edit]
user@R1# set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
```

```
[edit]
user@R2# set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
```

2. Configure the router identifier.

```
[edit]
user@R1# set routing-options router-id 10.255.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.255.0.2
```

3. On each ABR, configure the primary interface for the OSPF area.



**NOTE:** For OSPFv3, include the ospf3 statement at the [edit protocols] hierarchy level.

```
[edit]
user@R1# set protocols ospf 0.0.0.1 interface so-0/0/0
```

```
[edit]
user@R2# set protocols ospf 0.0.0.2 interface so-1/0/0
```

4. On each ABR, configure the secondary interface for the OSPF area.

```
[edit]
user@R1# set protocols ospf area 0.0.0.1 so-0/0/0 secondary
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.2 so-1/0/0 secondary
```

5. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

### **Results**

Confirm your configuration by entering the **show interfaces**, **show routing-options**, and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R1:

```
user@R1# show interfaces
so-0/0/0 {
 unit 0 {
 family inet {
 address 192.168.8.45/30;
 }
 }
}

user@R1# show routing-options
router-id 10.255.0.1;

user@R1# show protocols ospf
area 0.0.0.1 {
 interface so-0/0/0.0;
}
area 0.0.0.2 {
 interface so-0/0/0.0 {
 secondary;
 }
}
```

Configuration on ABR R2:

```
user@R2# show interfaces
so-0/0/0 {
 unit 0 {
 family inet {
 address 192.168.8.37/30;
 }
 }
}

user@R2# show routing-options
router-id 10.255.0.2;

user@R2# show protocols ospf
area 0.0.0.1 {
```

```

interface so-1/0/0.0;
}
area 0.0.0.2 {
 interface so-1/0/0.0 {
 secondary;
 }
}

```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Secondary Interface on page 3835](#)
- [Verifying the Interfaces in the Area on page 3835](#)
- [Verifying Neighbor Adjacencies on page 3835](#)

### **Verifying the Secondary Interface**

**Purpose** Verify that the secondary interface appears for the configured area. The Secondary field displays if the interface is configured as a secondary interface. The output might also show the same interface listed in multiple areas.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

### **Verifying the Interfaces in the Area**

**Purpose** Verify the interfaces configured for the specified area.

**Action** From operational mode, enter the **show ospf interface area *area-id*** command for OSPFv2, and enter the **show ospf3 interface area *area-id*** command for OSPFv3..

### **Verifying Neighbor Adjacencies**

**Purpose** Verify the primary and secondary neighbor adjacencies. The Secondary field displays if the neighbor is on a secondary interface.

**Action** From operational mode, enter the **show ospf neighbor detail** command for OSPFv2, and enter the **show ospf3 neighbor detail** command for OSPFv3.

**Related Documentation**

- [OSPF Areas and Router Functionality Overview on page 3803](#)
- [Understanding OSPF Areas and Backbone Areas on page 3814](#)
- [OSPF Configuration Overview](#)

## **Example: Disabling OSPFv2 Compatibility with RFC 1583**

- [OSPFv2 Compatibility with RFC 1583 Overview on page 3836](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 3836](#)

## OSPFv2 Compatibility with RFC 1583 Overview

---

In the first implementation of OSPF (RFC1583, *OSPF Version 2*), the summary route assumes the cost of the granular route with the lowest cost. OSPF RFC 2328, *OSPF Version 2* changes the behavior so that the summary route assumes the cost of the granular route with the highest cost. OSPF readvertises the summary route whenever the cost of the summary changes. When using the default RFC 1583 behavior, this happens when the granular route with the lowest metric is changed or lost. When RFC 2328 is used, this happens when the granular route with the highest cost is changed or lost.

By default, the Junos OS implementation of OSPF is compatible with RFC 1583. This means that Junos OS maintains a single best route to an autonomous system (AS) boundary router in the OSPF routing table, rather than multiple intra-AS paths, if they are available. You can disable compatibility with RFC 1583. It is preferable to do so when the same external destination is advertised by AS boundary routers that belong to different OSPF areas. When you disable compatibility with RFC 1583, the OSPF routing table maintains the multiple intra-AS paths that are available, which the router uses to calculate AS external routes as defined in RFC 2328. Being able to use multiple available paths to calculate an AS external route can prevent routing loops.

## Example: Disabling OSPFv2 Compatibility with RFC 1583

---

This example shows how to disable OSPFv2 compatibility with RFC 1583 on the routing device.

- [Requirements on page 3836](#)
- [Overview on page 3836](#)
- [Configuration on page 3836](#)
- [Verification on page 3837](#)

### Requirements

No special configuration beyond device initialization is required before disabling OSPFv2 compatibility with RFC 1583.

### Overview

The introduction of RFC 2328 changed the method used to calculate the routes in an OSPF network. By default, the Junos OS implementation of OSPFv2 is compatible with RFC 1583, so OSPF uses the minimum cost to determine the route to any of the networks within the specified range. When you disable RFC 1583 compatibility, OSPF uses the maximum cost to determine the route to any of the networks within the specified range. To minimize the potential for routing loops, configure the same RFC compatibility on all OSPF devices in an OSPF domain.

### Configuration

#### CLI Quick Configuration

To quickly disable OSPFv2 compatibility with RFC 1583, copy the following command and paste it into the CLI. You configure this setting on all devices that are part of the OSPF domain.

[edit]



```
set protocols ospf no-rfc-1583
```

**Step-by-Step Procedure** To disable OSPFv2 compatibility with RFC 1583:

1. Disable RFC 1583.  

```
[edit]
user@host# set protocols ospf no-rfc-1583
```
2. If you are done configuring the device, commit the configuration.  

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



**NOTE:** Repeat this configuration on each routing device that participates in an OSPF routing domain.

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-rfc-1583;
```

### Verification

Confirm that the configuration is working properly.

#### Verifying the OSPF Routes

**Purpose** Verify that the OSPF routing table maintains the intra-AS paths with the largest metric, which the router uses to calculate AS external routes.

**Action** From operational mode, enter the **show ospf route detail** command.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

## OSPF Interface Configuration

- [Examples: Configuring OSPF Interfaces on page 3837](#)

### Examples: Configuring OSPF Interfaces

- [About OSPF Interfaces on page 3838](#)
- [Example: Configuring an Interface on a Broadcast or Point-to-Point Network on page 3839](#)

- [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network on page 3841](#)
- [Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network on page 3844](#)
- [Example: Configuring OSPF Demand Circuits on page 3846](#)
- [Example: Configuring a Passive OSPF Interface on page 3848](#)
- [Example: Configuring OSPFv2 Peer interfaces on page 3850](#)

---

### About OSPF Interfaces

To activate OSPF on a network, you must enable the OSPF protocol on one or more interfaces on each device within the network on which traffic is to travel. How you configure the interface depends on whether the interface is connected to a broadcast or point-to-point network, a point-to-multipoint network, a nonbroadcast multiaccess (NBMA) network, or across a demand circuit.

- A broadcast interface behaves as if the routing device is connected to a LAN.
- A point-to-point interface provides a connection between a single source and a single destination (there is only one OSPF adjacency).
- A point-to-multipoint interface provides a connection between a single source and multiple destinations.
- An NBMA interface behaves in a similar fashion to a point-to-multipoint interface, but you might configure an NBMA interface to interoperate with other equipment.
- A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You can also configure an OSPF interface to be passive, to operate in passive traffic engineering mode, or to be a peer interface.

- A passive interface advertises its address, but does not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- An interface operating in OSPF passive traffic engineering mode floods link address information within the autonomous system (AS) and makes it available for traffic engineering calculations.
- A peer interface can be configured for OSPFv2 routing devices. A peer interface is required for Generalized MPLS (GMPLS) to transport traffic engineering information through a link separate from the control channel. You establish this separate link by configuring a peer interface. The peer interface name must match the Link Management Protocol (LMP) peer name. A peer interface is optional for a hierarchy of RSVP label-switched paths (LSPs). After you configure the forwarding adjacency, you can configure OSPFv2 to advertise the traffic engineering properties of a forwarding adjacency to a specific peer.

Point-to-point interfaces differ from multipoint in that only one OSPF adjacency is possible. (A LAN, for instance, can have multiple addresses and can run OSPF on each subnet simultaneously.) As such, when you configure a numbered point-to-point interface

to OSPF by name, multiple OSPF interfaces are created. One, which is unnumbered, is the interface on which the protocol is run. An additional OSPF interface is created for each address configured on the interface, if any, which is automatically marked as passive.

For OSPFv3, one OSPF-specific interface must be created per interface name configured under OSPFv3. OSPFv3 does not allow interfaces to be configured by IP address.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



**NOTE:** When you configure OSPFv2 on an interface, you must also include the **family inet** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. When you configure OSPFv3 on an interface, you must also include the **family inet6** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support address families other than unicast IPv6.

### Example: Configuring an Interface on a Broadcast or Point-to-Point Network

This example shows how to configure an OSPF interface on a broadcast or point-to-point network.

- [Requirements on page 3839](#)
- [Overview on page 3839](#)
- [Configuration on page 3840](#)
- [Verification on page 3841](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### Overview

If the interface on which you are configuring OSPF supports broadcast mode (such as a LAN), or if the interface supports point-to-point mode (such as a PPP interface or a point-to-point logical interface on Frame Relay), you specify the interface by including the IP address or the interface name for OSPFv2, or only the interface name for OSPFv3. In Junos OS Release 9.3 and later, an OSPF point-to-point interface can be an Ethernet

interface without a subnet. If you configure an interface on a broadcast network, designated router and backup designated router election is performed.



**NOTE:** Using both the interface name and the IP address of the same interface produces an invalid configuration.

In this example, you configure interface **ge-0/2/0** as an OSPFv2 interface in OSPF area 0.0.0.1.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF interface on a broadcast or point-to-point network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
set protocols ospf area 0.0.0.1 interface ge-0/2/0
```

#### Step-by-Step Procedure

To configure an OSPF interface on a broadcast or point-to-point network:

1. Configure the interface.



**NOTE:** For an OSPFv3 interface, specify an IPv6 address.

```
[edit]
user@host# set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
```

2. Create an OSPF area.



**NOTE:** For an OSPFv3 interface, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface ge-0/2/0
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

### Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-0/2/0 {
 unit 0 {
 family inet {
 address 10.0.0.1/32;
 }
 }
}

user@host# show protocols ospf
area 0.0.0.1 {
 interface ge-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

### Verification

Confirm that the configuration is working properly.

#### Verifying the OSPF Interface

- |                |                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify the interface configuration. Depending on your deployment, the Type field might display LAN or P2P.                                                  |
| <b>Action</b>  | From operational mode, enter the <b>show ospf interface detail</b> command for OSPFv2, and enter the <b>show ospf3 interface detail</b> command for OSPFv3. |

### Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

This example shows how to configure an OSPFv2 interface on a nonbroadcast multiaccess (NBMA) network.

- [Requirements on page 3841](#)
- [Overview on page 3842](#)
- [Configuration on page 3843](#)
- [Verification on page 3844](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).

- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).

### Overview

When you configure OSPFv2 on an NBMA network, you can use nonbroadcast mode rather than point-to-multipoint mode. Using this mode offers no advantages over point-to-multipoint mode, but it has more disadvantages than point-to-multipoint mode. Nevertheless, you might occasionally find it necessary to configure nonbroadcast mode to interoperate with other equipment. Because there is no autodiscovery mechanism, you must configure each neighbor.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routing devices must have a direct connection to both the designated and backup designated routers, or unpredictable results occur.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration. For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the interface name.

In this example, you configure the Asynchronous Transfer Mode (ATM) interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify the following settings:

- **interface-type nbma**—Sets the interface to run in NBMA mode. You must explicitly configure the interface to run in NBMA mode.
- **neighbor address <eligible>**—Specifies the IP address of the neighboring device. OSPF routing devices normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the device cannot discover its neighbors dynamically, so you must configure all the neighbors statically. To configure multiple neighbors, include multiple **neighbor** statements. If you want the neighbor to be a designated router, include the **eligible** keyword.
- **poll-interval**—Specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before it establishes adjacency with a neighbor. Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The range is from 1 through 255 seconds. By default, the device sends hello packets out the interface every 120 seconds before it establishes adjacency with a neighbor.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the **poll-interval** statement to the time specified in the **hello-interval** statement.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPFv2 interface on an NBMA network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 interface-type nbma
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 neighbor 192.0.2.2 eligible
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 poll-interval 130
```

**Step-by-Step Procedure** To configure an OSPFv2 interface on an NBMA network:

1. Configure the interface.  

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
```
2. Create an OSPF area.  

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```
3. Assign the interface to the area.  
 In this example, include the **eligible** keyword to allow the neighbor to be a designated router.  

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 interface-type nbma neighbor 192.0.2.2 eligible
```
4. Configure the poll interval.  

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 poll-interval 130
```
5. If you are done configuring the device, commit the configuration.  

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

### Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
 unit 0 {
 family inet {
 address 192.0.2.1/32;
 }
 }
}

user@host# show protocols ospf
area 0.0.0.1 {
 interface at-0/1/0.0 {
```

```
 interface-type nbma;
 neighbor 192.0.2.2 eligible;
 poll-interval 130;
 }
}
```

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the OSPF Interface**

**Purpose** Verify the interface configuration. Confirm that the Type field displays NBMA.

**Action** From operational mode, enter the **show ospf interface detail** command.

---

### **Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network**

This example shows how to configure an OSPFv2 interface on a point-to-multipoint network.

- [Requirements on page 3844](#)
- [Overview on page 3844](#)
- [Configuration on page 3845](#)
- [Verification on page 3845](#)

### **Requirements**

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

### **Overview**

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint Asynchronous Transfer Mode (ATM) or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, you must configure each neighbor.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration.

In this example, you configure ATM interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify 192.0.2.1 as the neighbor's IP address.



**Configuration**

**CLI Quick Configuration** To quickly configure an OSPFv2 interface on a point-to-multipoint network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
set protocols ospf area 0.0.0.1 interface at-0/1/0 neighbor 192.0.2.1
```

**Step-by-Step Procedure** To configure an OSPFv2 interface on a point-to-multipoint network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area and specify the neighbor.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 neighbor 192.0.2.1
```

To configure multiple neighbors, include a **neighbor** statement for each neighbor.

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

**Results**

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
 unit 0 {
 family inet {
 address 192.0.2.2/32;
 }
 }
}

user@host# show protocols ospf
area 0.0.0.1 {
 interface at-0/1/0.0 {
 neighbor 192.0.2.1;
 }
}
```

**Verification**

Confirm that the configuration is working properly.

### *Verifying the OSPF Interface*

**Purpose** Verify the interface configuration. Confirm that the Type field displays P2MP.

**Action** From operational mode, enter the **show ospf interface detail** command.

---

### **Example: Configuring OSPF Demand Circuits**

This example shows how to configure an OSPF demand circuit interface.

- [Requirements on page 3846](#)
- [Overview on page 3846](#)
- [Configuration on page 3847](#)
- [Verification on page 3848](#)

#### **Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.



**NOTE:** If you are using OSPF demand circuits over an ISDN link, you must configure an ISDN interface and enable dial-on-demand routing. See the *Junos® OS Network Interfaces*.

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### **Overview**

OSPF sends periodic hello packets to establish and maintain neighbor adjacencies and uses link-state advertisements (LSAs) to make routing calculations and decisions. OSPF support for demand circuits is defined in RFC 1793, *Extending OSPF to Support Demand Circuits*, and suppresses the periodic hello packets and LSAs. A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You configure demand circuits on an OSPF interface. When the interface becomes a demand circuit, all hello packets and LSAs are suppressed as soon as OSPF synchronization is achieved. LSAs have a DoNotAge bit that stops the LSA from aging and prevents periodic updates from being sent. Hello packets and LSAs are sent and received on a demand-circuit interface only when there is a change in the network topology. This reduces the amount of traffic through the OSPF interface.

Consider the following when configuring OSPF demand circuits:

- Periodic hellos are only suppressed on point-to-point and point-to-multipoint interfaces. If you configure demand circuits on an OSPF broadcast network or on an OSPF nonbroadcast multiaccess (NBMA) network, periodic hello packets are still sent.
- Demand circuit support on an OSPF point-to-multipoint interface resembles that for point-to-point interfaces. If you configure a point-to-multipoint interface as a demand circuit, the device negotiates hello suppression separately on each interface that is part of the point-to-multipoint network.

This example assumes that you have a point-to-point connection between two devices using SONET/SDH interfaces. A demand-circuit interface automatically negotiates the demand-circuit connection with its OSPF neighbor. If the neighbor does not support demand circuits, then no demand circuit connection is established.

In this example, you configure OSPF interface **so-0/1/0** in OSPF area 0.0.0.1 as a demand circuit.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF demand circuit interface, copy the following command and paste it into the CLI. You must configure both neighboring interfaces for OSPF demand circuits for the connection to be established.

[edit]

```
set protocols ospf area 0.0.0.1 interface so-0/1/0 demand-circuit
```

#### Step-by-Step Procedure

To configure an OSPF demand circuit interface on one neighboring interface:

1. Create an OSPF area.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

[edit ]

```
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the neighboring interface as a demand circuit.

[edit protocols ospf area 0.0.0.1]

```
user@host# set interface so-0/1/0 demand-circuit
```

3. If you are done configuring the device, commit the configuration.

[edit protocols ospf area 0.0.0.1]

```
user@host# commit
```



**NOTE:** Repeat this entire configuration on the other neighboring interface.

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
ospf {
 area 0.0.0.1 {
 interface so-0/1/0.0 {
 demand-circuit;
 }
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Status of Neighboring Interfaces

- |                |                                                                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify information about the neighboring interface. When the neighbor is configured for demand circuits, a DC flag displays.                              |
| <b>Action</b>  | From operational mode, enter the <b>show ospf neighbor detail</b> command for OSPFv2, and enter the <b>show ospf3 neighbor detail</b> command for OSPFv3. |

---

### Example: Configuring a Passive OSPF Interface

This example shows how to configure a passive OSPF interface. A passive OSPF interface advertises its address but does not run the OSPF protocol.

- [Requirements on page 3848](#)
- [Overview on page 3849](#)
- [Configuration on page 3849](#)
- [Verification on page 3850](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

### Overview

By default, OSPF must be configured on an interface for direct interface addresses to be advertised as interior routes. To advertise the direct interface addresses without actually running OSPF on that interface (adjacencies are not formed and hello packets are not generated), you configure that interface as a passive interface.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



**NOTE:** If you do not want to see notifications for state changes in a passive OSPF interface, you can disable the OSPF traps for the interface by including the **no-interface-state-traps** statement. The **no-interface-state-traps** statement is supported only for OSPFv2.

In this example, you configure interface **ge-0/2/0** as a passive OSPF interface in area 0.0.0.1 by including the **passive** statement.

### Configuration

#### CLI Quick Configuration

To quickly configure a passive OSPF interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface ge-0/2/0 passive
```

#### Step-by-Step Procedure

To configure a passive OSPF interface:

1. Create an OSPF area.



**NOTE:** For an OSPFv3 interface, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the passive interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface ge-0/2/0 passive
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
 area 0.0.0.1 {
 interface ge-0/2/0.0 {
 passive;
 }
 }
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Status of OSPF Interfaces

**Purpose** Verify the status of the OSPF interface. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

---

### Example: Configuring OSPFv2 Peer interfaces

This example shows how to configure an OSPFv2 peer interface.

- [Requirements on page 3850](#)
- [Overview on page 3851](#)
- [Configuration on page 3851](#)
- [Verification on page 3851](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.
- Configure Generalized MPLS per your network requirements. See *LMP Configuration Overview* in the *Junos OS MPLS Applications Configuration Guide*.

### Overview

You can configure an OSPFv2 peer interface for many reasons, including when you configure Generalized MPLS (GMPLS). This example configures a peer interface for GMPLS. GMPLS requires traffic engineering information to be transported through a link separate from the control channel. You establish this separate link by configuring a peer interface. The OSPFv2 peer interface name must match the Link Management Protocol (LMP) peer name. You configure GMPLS and the LMP settings separately from OSPF.

This example assumes that GMPLS and the LMP peer named **oxc1** are already configured, and you need to configure the OSPFv2 peer interface in area 0.0.0.0.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPFv2 peer interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 peer-interface oxc1
```

**Step-by-Step Procedure** To configure a peer OSPFv2 interface used by the LMP:

1. Create an OSPF area.  

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```
2. Configure the peer interface.  

```
[edit protocols ospf area 0.0.0.0]
user@host# set peer-interface oxc1
```
3. If you are done configuring the device, commit the configuration.  

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
 area 0.0.0.0 {
 peer-interface oxc1;
 }
```

### Verification

Confirm that the configuration is working properly.

#### Verifying the Configured OSPFv2 Peer

**Purpose** Verify the status of the OSPFv2 peer. When an OSPFv2 peer is configured for GMPLS, the Peer Name field displays the name of the LMP peer that you created for GMPLS, which is also the configured OSPFv2 peer.

**Action** From operational mode, enter the **show link-management** command.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

---

## OSPF Route Control Configuration

- [Examples: Configuring OSPF Route Summarization on page 3852](#)
- [Examples: Configuring OSPF Traffic Control on page 3861](#)
- [Example: Configuring OSPF Overload Mode on page 3871](#)

### Examples: Configuring OSPF Route Summarization

- [Understanding OSPF Route Summarization on page 3852](#)
- [Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements on page 3852](#)
- [Example: Limiting the Number of Prefixes Exported to OSPF on page 3858](#)
- [Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 3860](#)

---

#### Understanding OSPF Route Summarization

Area border routers (ABRs) send summary link advertisements to describe the routes to other areas. Depending on the number of destinations, an area can get flooded with a large number of link-state records, which can utilize routing device resources. To minimize the number of advertisements that are flooded into an area, you can configure the ABR to coalesce, or summarize, a range of IP addresses and send reachability information about these addresses in a single link-state advertisement (LSA). You can summarize one or more ranges of IP addresses, where all routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

For an OSPF area, you can summarize and filter intra-area prefixes. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place. For an OSPF not-so-stubby area (NSSA), you can only coalesce or filter NSSA external (Type 7) LSAs before they are translated into AS external (Type 5) LSAs and enter the backbone area. All external routes learned within the area that do not fall into the range of one of the prefixes are advertised individually to other areas.

In addition, you can also limit the number of prefixes (routes) that are exported into OSPF. By setting a user-defined maximum number of prefixes, you prevent the routing device from flooding an excessive number of routes into an area.

---

#### Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements

This example shows how to summarize routes sent into the backbone area.

- [Requirements on page 3853](#)
- [Overview on page 3853](#)



- [Configuration on page 3854](#)
- [Verification on page 3857](#)

### Requirements

Before you begin:


- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#)
- Configure a static route. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Configuration Guide*.

### Overview

You can summarize a range of IP addresses to minimize the size of the backbone router's link-state database. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

[Figure 126 on page 3853](#) shows the topology used in this example. R5 is the ABR between area 0.0.0.4 and the backbone. The networks in area 0.0.0.4 are 10.0.8.4/30, 10.0.8.0/30, and 10.0.8.8/30, which can be summarized as 10.0.8.0/28. R3 is the ABR between NSSA area 0.0.0.3 and the backbone. The networks in area 0.0.0.3 are 10.0.4.4/30, 10.0.4.0/30, and 10.0.4.12/30, which can be summarized as 10.0.4.0/28. Area 0.0.0.3 also contains external static route 3.0.0.8 that you will prevent from flooding throughout the network.

**Figure 126: Summarizing Ranges of Routes in OSPF**



In this example, you configure the ABRs for route summarization by including the following settings:

- **area-range**—For an area, summarizes a range of IP addresses when sending summary intra-area link advertisements. For an NSSA, summarizes a range of IP addresses when sending NSSA link-state advertisements (Type 7 LSAs). The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas.
- **network/mask-length**—Indicates the summarized IP address range and the number of significant bits in the network mask.
- **restrict**—On the NSSA ABR, prevents the configured summary from being advertised. In this example, we do not want to flood the external route outside of area 0.0.0.3.

### Configuration

#### CLI Quick Configuration

- To quickly configure route summarization for an OSPF area, copy the following commands and paste them into the CLI. The following is the configuration on ABR R5:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3
set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5
set protocols ospf area 0.0.0.4 stub
set protocols ospf area 0.0.0.4 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28
```

- To quickly configure route summarization for an OSPF NSSA, copy the following commands and paste them into the CLI. The following is the configuration on ABR R3:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10
set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7
set protocols ospf area 0.0.0.3 interface fe-0/0/1
set protocols ospf area 0.0.0.3 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
set protocols ospf area 0.0.0.3 nssa
set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict
```

#### Step-by-Step Procedure

To summarize routes sent to the backbone area:

1. Configure the interfaces.



**NOTE:** For OSPFv3, include IPv6 addresses.

```
[edit]
user@R5# set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3
user@R5# set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4
user@R5# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3
user@R5# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5

[edit]
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10
user@R3# set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1
user@R3# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1
user@R3# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7
```

2. Configure the type of OSPF area.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 stub

[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa
```

3. Assign the interfaces to the OSPF areas.

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/1
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/2
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/4

user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/1
user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/2
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/4
```

4. Summarize the routes that are flooded into the backbone.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28

[edit]
user@R3# set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
```

5. On ABR R3, restrict the external static route from leaving area 0.0.0.3.

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict
```

6. If you are done configuring the devices, commit the configuration.

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

### Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R5:

```
user@R5# show interfaces
fe-0/0/0 {
 unit 0 {
 family inet {
 address 10.0.2.3/32;
 }
 }
}
fe-0/0/1 {
 unit 0 {
 family inet {
 address 10.0.8.3/32;
 }
 }
}
fe-0/0/2 {
 unit 0 {
 family inet {
 address 10.0.8.4/32;
 }
 }
}
fe-0/0/4 {
 unit 0 {
 family inet {
 address 10.0.2.5/32;
 }
 }
}

user@R5# show protocols ospf
area 0.0.0.0 {
 interface fe-0/0/0.0;
 interface fe-0/0/4.0;
}
area 0.0.0.4 {
 stub;
 area-range 10.0.8.0/28;
 interface fe-0/0/1.0;
 interface fe-0/0/2.0;
}
```

Configuration on ABR R3:

```
user@R3# show interfaces
fe-0/0/0 {
 unit 0 {
 family inet {
 address 10.0.2.1/32;
 }
 }
}
fe-0/0/1 {
 unit 0 {
 family inet {
```

```

 address 10.0.4.10/32;
 }
}
fe-0/0/2 {
 unit 0 {
 family inet {
 address 10.0.4.1/32;
 }
 }
}
fe-0/0/4 {
 unit 0 {
 family inet {
 address 10.0.2.7/32;
 }
 }
}

user@R3t# show protocols ospf
area 0.0.0.0 {
 interface fe-0/0/0.0;
 interface fe-0/0/4.0;
}
area 0.0.0.3 {
 nssa {
 area-range 3.0.0.0/8 restrict;
 }
 area-range 10.0.4.0/28;
 interface fe-0/0/1.0;
 interface fe-0/0/2.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces** and **show protocols ospf3** commands.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the Summarized Route**

**Purpose** Verify that the routes you configured for route summarization are being aggregated by the ABRs before the routes enter the backbone area. Confirm route summarization by checking the entries of the OSPF link-state database for the routing devices in the backbone.

**Action** From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

### Example: Limiting the Number of Prefixes Exported to OSPF

---

This example shows how to limit the number of prefixes exported to OSPF.

- [Requirements on page 3858](#)
- [Overview on page 3858](#)
- [Configuration on page 3858](#)
- [Verification on page 3859](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### Overview

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area.

You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem. If the routing device exceeds the configured prefix export value, the routing device purges the external prefixes and enters into an overload state. This state ensures that the routing device is not overwhelmed as it attempts to process routing information. The prefix export limit number can be a value from 0 through 4,294,967,295.

In this example, you configure a prefix export limit of 100,000 by including the **prefix-export-limit** statement.

#### Configuration

##### CLI Quick Configuration

To quickly limit the number of prefixes exported to OSPF, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf prefix-export-limit 100000
```

**Step-by-Step Procedure** To limit the number of prefixes exported to OSPF:

1. Configure the prefix export limit value.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf prefix-export-limit 100000
```

2. If you are done configuring the device, commit the configuration.

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

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
prefix-export-limit 100000;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Prefix Export Limit

**Purpose** Verify the prefix export counter that displays the number of routes exported into OSPF.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

### Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

---

The OSPF standard requires that every link-state advertisement (LSA) be refreshed every 30 minutes. The Juniper Networks implementation refreshes LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. This requirement can result in traffic overhead that makes it difficult to scale OSPF networks. You can override the default behavior by specifying that the DoNotAge bit be set in self-originated LSAs when they are initially sent by the router or switch. Any LSA with the DoNotAge bit set is reflooded only when a change occurs in the LSA. This feature thus reduces protocol traffic overhead while permitting any changed LSAs to be flooded immediately. Routers or switches enabled for flood reduction continue to send hello packets to their neighbors and to age self-originated LSAs in their databases.

The Juniper implementation of OSPF refresh and flooding reduction is based on RFC 4136, *OSPF Refresh and Flooding Reduction in Stable Topologies*. However, the Juniper implementation does not include the forced-flooding interval defined in the RFC. Not implementing the forced-flooding interval ensures that LSAs with the DoNotAge bit set are reflooded only when a change occurs.

This feature is supported for the following:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 realms
- OSPFv2 and OSPFv3 virtual links
- OSPFv2 sham links
- OSPFv2 peer interfaces
- All routing instances supported by OSPF
- Logical systems

To configure flooding reduction for an OSPF interface, include the **flood-reduction** statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-id]` hierarchy level.



**NOTE:** If you configure flooding reduction for an interface configured as a demand circuit, the LSAs are not initially flooded, but sent only when their content has changed. Hello packets and LSAs are sent and received on a demand-circuit interface only when a change occurs in the network topology.

---

In the following example, the OSPF interface **so-0/0/1.0** is configured for flooding reduction. As a result, all the LSAs generated by the routes that traverse the specified interface have the DoNotAge bit set when they are initially flooded, and LSAs are refreshed only when a change occurs.

```
[edit]
protocols ospf {
 area 0.0.0.0 {
```



```

interface so-0/0/1.0 {
 flood-reduction;
}
interface lo0.0;
interface so-0/0/0.0;
}

```



**NOTE:** Beginning with Junos OS Release 12.2, you can configure a global default link-state advertisement (LSA) flooding interval in OSPF for self-generated LSAs by including the `lsa-refresh-interval minutes` statement at the `[edit protocols (ospf | ospf3)]` hierarchy level. The Juniper Networks implementation refreshes LSAs every 50 minutes. The range is 25 through 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes.

If you have both the global LSA refresh interval configured for OSPF and OSPF flooding reduction configured for a specific interface in an OSPF area, the OSPF flood reduction configuration takes precedence for that specific interface.

- Related Documentation**
- [OSPF Overview on page 3798](#)
  - [OSPF Configuration Overview](#)

## Examples: Configuring OSPF Traffic Control

- [Understanding OSPF Traffic Control on page 3861](#)
- [Example: Controlling the Cost of Individual OSPF Network Segments on page 3863](#)
- [Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 3867](#)
- [Example: Controlling OSPF Route Preferences on page 3869](#)

### Understanding OSPF Traffic Control

Once a topology is shared across the network, OSPF uses the topology to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest-path-first (SPF) algorithm. Routes with lower total path metrics are preferred over those with higher path metrics.

You can use the following methods to control OSPF traffic:

- Control the cost of individual OSPF network segments
- Dynamically adjust OSPF interface metrics based on bandwidth
- Control OSPF route selection

### *Controlling the Cost of Individual OSPF Network Segments*

OSPF uses the following formula to determine the cost of a route:

$\text{cost} = \text{reference-bandwidth} / \text{interface bandwidth}$

You can modify the reference-bandwidth value, which is used to calculate the default interface cost. The interface bandwidth value is not user-configurable and refers to the actual bandwidth of the physical interface.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface.

To control the flow of packets across the network, OSPF allows you to manually assign a cost (or metric) to a particular path segment. When you specify a metric for a specific OSPF interface, that value is used to determine the cost of routes advertised from that interface. For example, if all routers in the OSPF network use default metric values, and you increase the metric on one interface to 5, all paths through that interface have a calculated metric higher than the default and are not preferred.



**NOTE:** Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface.

When there are multiple equal-cost routes to the same destination in a routing table, an equal-cost multipath (ECMP) set is formed. If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose one of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. Define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the **[edit policy-options]** hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table.

#### ***Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth***

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, the Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. Junos OS uses the smallest configured bandwidth threshold value that is equal to or greater than the actual interface bandwidth to determine the metric value. If the interface bandwidth is greater than any of the configured bandwidth threshold values, the metric value configured for the interface is used instead of any of the bandwidth-based metric values configured. The ability to recalculate the metric for an interface when its bandwidth changes is especially useful for aggregate interfaces.



**NOTE:** You must also configure a metric for the interface when you enable bandwidth-based metrics.

### Controlling OSPF Route Preferences

You can control the flow of packets through the network using route preferences. Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. Although the default settings are appropriate for most environments, you might want to modify the default settings if all of the routing devices in your OSPF network use the default preference values, or if you are planning to migrate from OSPF to a different interior gateway protocol (IGP). If all of the devices use the default route preference values, you can change the route preferences to ensure that the path through a particular device is selected for the forwarding table any time multiple equal-cost paths to a destination exist. When migrating from OSPF to a different IGP, modifying the route preferences allows you to perform the migration in a controlled manner.

### Example: Controlling the Cost of Individual OSPF Network Segments

This example shows how to control the cost of individual OSPF network segments.

- [Requirements on page 3863](#)
- [Overview on page 3863](#)
- [Configuration on page 3864](#)
- [Verification on page 3866](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.

#### Overview

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred to those with higher path metrics. In this example, we explore how to control the cost of OSPF network segments.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface. This means that all interfaces faster than 100 Mbps have the same default cost metric of 1. If multiple equal-cost paths exist between a source

and destination address, OSPF routes packets along each path alternately, in round-robin fashion.


Having the same default metric might not be a problem if all of the interfaces are running at the same speed. If the interfaces operate at different speeds, you might notice that traffic is not routed over the fastest interface because OSPF equally routes packets across the different interfaces. For example, if your routing device has Fast Ethernet and Gigabit Ethernet interfaces running OSPF, each of these interfaces have a default cost metric of 1.

In the first example, you set the reference bandwidth to 10g (10 Gbps, as denoted by 10,000,000,000 bits) by including the **reference-bandwidth** statement. With this configuration, OSPF assigns the Fast Ethernet interface a default metric of 100, and the Gigabit Ethernet interface a metric of 10. Since the Gigabit Ethernet interface has the lowest metric, OSPF selects it when routing packets. The range is 9600 through 1,000,000,000,000 bits.

Figure 127 on page 3864 shows three routing devices in area 0.0.0.0 and assumes that the link between Device R2 and Device R3 is congested with other traffic. You can also control the flow of packets across the network by manually assigning a metric to a particular path segment. Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface. To prevent the traffic from Device R3 going directly to Device R2, you adjust the metric on the interface on Device R3 that connects with Device R1 so that all traffic goes through Device R1.

In the second example, you set the metric to 5 on interface **fe-1/0/1** on Device R3 that connects with Device R1 by including the **metric** statement. The range is 1 through 65,535.

**Figure 127: OSPF Metric Configuration**



#### Configuration

- [Configuring the Reference Bandwidth on page 3865](#)
- [Configuring a Metric for a Specific OSPF Interface on page 3865](#)

**Configuring the Reference Bandwidth**

**CLI Quick Configuration** To quickly configure the reference bandwidth, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf reference-bandwidth 10g
```

**Step-by-Step Procedure** To configure the reference bandwidth:

1. Configure the reference bandwidth to calculate the default interface cost.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf reference-bandwidth 10g
```



**TIP:** As a shortcut in this example, you enter `10g` to specify 10 Gbps reference bandwidth. Whether you enter `10g` or `10000000000`, the output of `show protocols ospf` command displays 10 Gbps as `10g`, not `10000000000`.

2. If you are done configuring the device, commit the configuration.

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



**NOTE:** Repeat this entire configuration on all routing devices in a shared network.

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
reference-bandwidth 10g;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

**Configuring a Metric for a Specific OSPF Interface**

**CLI Quick Configuration** To quickly configure a metric for a specific OSPF interface, copy the following command and paste it into the CLI.

```
[edit]
```

```
set protocols ospf area 0.0.0.0 interface fe-1/0/1 metric 5
```

**Step-by-Step  
Procedure**

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
```

```
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host# set interface fe-1/0/1 metric 5
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host# commit
```

**Results**

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface fe-1/0/1.0 {
 metric 5;
 }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Configured Metric on page 3866](#)
- [Verifying the Route on page 3867](#)

**Verifying the Configured Metric****Purpose**

Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

**Action**

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

**Verifying the Route**

**Purpose** When choosing paths to a destination, OSPF uses the path with the lowest total cost. Confirm that OSPF is using the appropriate path.

**Action** From operational mode, enter the **show route** command.

**Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth**

This example shows how to dynamically adjust OSPF interface metrics based on bandwidth.

- [Requirements on page 3867](#)
- [Overview on page 3867](#)
- [Configuration on page 3868](#)
- [Verification on page 3869](#)

**Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.

**Overview**

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface. When the bandwidth of an interface changes, the Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. When you configure bandwidth-based metric values, you typically configure multiple bandwidth and metric values.

In this example, you configure OSPF interface **ae0** for bandwidth-based metrics by including the **bandwidth-based-metrics** statement and the following settings:

- **bandwidth**—Specifies the bandwidth threshold in bits per second. The range is 9600 through 1,000,000,000,000,000.
- **metric**—Specifies the metric value to associate with a specific bandwidth value. The range is 1 through 65,535.

### Configuration

**CLI Quick Configuration** To quickly configure bandwidth threshold values and associated metric values for an OSPF interface, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface ae0.0 metric 5
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 1g
metric 60
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 10g
metric 50
```

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ae0 metric 5
```

3. Configure the bandwidth threshold values and associated metric values.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 1g metric 60
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 10g metric 50
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface ae0.0 {
 bandwidth-based-metrics {
 bandwidth 1g metric 60;
 bandwidth 10g metric 50;
 }
 metric 5;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.



**Verification**

Confirm that the configuration is working properly.

**Verifying the Configured Metric**

- Purpose** Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.
- Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Example: Controlling OSPF Route Preferences**

This example shows how to control OSPF route selection in the forwarding table. This example also shows how you might control route selection if you are migrating from OSPF to another IGP.

- [Requirements on page 3869](#)
- [Overview on page 3869](#)
- [Configuration on page 3870](#)
- [Verification on page 3871](#)

**Requirements**

This example assumes that OSPF is properly configured and running in your network, and you want to control route selection because you are planning to migrate from OSPF to a different IGP.

- Configure the device interfaces. See the *Junos® OS Network Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the IGP that you want to migrate to. See the *Junos OS Routing Protocols Configuration Guide*.

**Overview**

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. You might want to modify this setting if you are planning to migrate from OSPF to a different IGP. Modifying the route preferences enables you to perform the migration in a controlled manner.

This example makes the following assumptions:

- OSPF is already running in your network.
- You want to migrate from OSPF to IS-IS.

- You configured IS-IS per your network requirements and confirmed it is working properly.

In this example, you increase the OSPF route preference values to make them less preferred than IS-IS routes by specifying 168 for internal OSPF routes and 169 for external OSPF routes. IS-IS internal routes have a preference of either 15 (for Level 1) or 18 (for Level 2), and external routes have a preference of 160 (for Level 1) or 165 (for Level 2). In general, it is preferred to leave the new protocol at its default settings to minimize complexities and simplify any future addition of routing devices to the network. To modify the OSPF route preference values, configure the following settings:

- **preference**—Specifies the route preference for internal OSPF routes. By default, internal OSPF routes have a value of 10. The range is from 0 through 4,294,967,295 ( $2^{32} - 1$ ).
- **external-preference**—Specifies the route preference for external OSPF routes. By default, external OSPF routes have a value of 150. The range is from 0 through 4,294,967,295 ( $2^{32} - 1$ ).

### Configuration

#### CLI Quick Configuration

To quickly configure the OSPF route preference values, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf preference 168 external-preference 169
```

To configure route selection:

1. Enter OSPF configuration mode and set the external and internal routing preferences.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf preference 168 external-preference 169
```

2. If you are done configuring the device, commit the configuration.

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

#### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
preference 168;
external-preference 169;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Route on page 3871](#)

**Verifying the Route**

**Purpose** Verify that the IGP is using the appropriate route. After the new IGP becomes the preferred protocol (in this example, IS-IS), you should monitor the network for any issues. After you confirm that the new IGP is working properly, you can remove the OSPF configuration from the routing device by entering the **delete ospf** command at the **[edit protocols]** hierarchy level.

**Action** From operational mode, enter the **show route** command.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

**Example: Configuring OSPF Overload Mode**

- [OSPF Overload Function Overview on page 3871](#)
- [Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 3872](#)

**OSPF Overload Function Overview**

If the time elapsed after the OSPF instance is enabled is less than the specified timeout, overload mode is set.

You can configure the local routing device so that it appears to be overloaded. An overloaded routing device determines it is unable to handle any more OSPF transit traffic, which results in sending OSPF transit traffic to other routing devices. OSPF traffic to directly attached interfaces continues to reach the routing device. You might configure overload mode for many reasons, including:

- If you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic. This could include a routing device that is connected to the network for analysis purposes, but is not considered part of the production network, such as network management routing devices.
- If you are performing maintenance on a routing device in a production network. You can move traffic off that routing device so network services are not interrupted during your maintenance window.

You configure or disable overload mode in OSPF with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the OSPF instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. When the timer expires, overload mode is cleared. In overload mode,

the router link-state advertisement (LSA) is originated with all the transit router links (except stub) set to a metric of 0xFFFF. The stub router links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and to take paths around the routing device. However, the overloaded routing device's own links are still accessible.



**NOTE:** The routing device can also dynamically enter the overload state, regardless of configuring the device to appear overloaded. For example, if the routing device exceeds the configured OSPF prefix limit, the routing device purges the external prefixes and enters into an overload state. You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem.

---

### Example: Configuring OSPF to Make Routing Devices Appear Overloaded

---

This example shows how to configure a routing device running OSPF to appear to be overloaded.

- [Requirements on page 3872](#)
- [Overview on page 3872](#)
- [Configuration on page 3873](#)
- [Verification on page 3874](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### Overview

You can configure a local routing device running OSPF to appear to be overloaded, which allows the local routing device to participate in OSPF routing, but not for transit traffic. When configured, the transit interface metrics are set to the maximum value of 65535.

This example includes the following settings:

- **overload**—Configures the local routing device so it appears to be overloaded. You might configure this if you want the routing device to participate in OSPF routing, but do not

want it to be used for transit traffic, or you are performing maintenance on a routing device in a production network.

- **timeout seconds**—(Optional) Specifies the number of seconds at which the overload is reset. If no timeout interval is specified, the routing device remains in the overload state until the overload statement is deleted or a timeout is set. In this example, you configure 60 seconds as the amount of time the routing device remains in the overload state. By default, the timeout interval is 0 seconds (this value is not configured). The range is from 60 through 1800 seconds.

### Configuration

**CLI Quick Configuration** To quickly configure a local routing device to appear as overloaded, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf overload timeout 60
```

**Step-by-Step Procedure** To configure a local routing device to appear overloaded:

1. Enter OSPF configuration mode.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host edit protocols ospf
```

2. Configure the local routing device to be overloaded.

```
[edit protocols ospf]
user@host set overload
```

3. (Optional) Configure the number of seconds at which overload is reset.

```
[edit protocols ospf]
user@host set overload timeout 60
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration. The output includes the optional **timeout** statement.

```
user@host# show protocols ospf
overload timeout 60;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying Traffic Has Moved Off Devices on page 3874](#)
- [Verifying Transit Interface Metrics on page 3874](#)
- [Verifying the Overload Configuration on page 3874](#)
- [Verifying the Viable Next Hop on page 3874](#)

**Verifying Traffic Has Moved Off Devices**

**Purpose** Verify that the traffic has moved off the upstream devices.

**Action** From operational mode, enter the **show interfaces detail** command.

**Verifying Transit Interface Metrics**

**Purpose** Verify that the transit interface metrics are set to the maximum value of 65535 on the downstream neighboring device.

**Action** From operational mode, enter the **show ospf database router detail advertising-router address** command for OSPFv2, and enter the **show ospf3 database router detail advertising-router address** command for OSPFv3.

**Verifying the Overload Configuration**

**Purpose** Verify that overload is configured by reviewing the Configured overload field. If the overload timer is also configured, this field also displays the time that remains before it is set to expire.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and the **show ospf3 overview** command for OSPFv3.

**Verifying the Viable Next Hop**

**Purpose** Verify the viable next hop configuration on the upstream neighboring device. If the neighboring device is overloaded, it is not used for transit traffic and is not displayed in the output.

**Action** From operational mode, enter the **show route address** command.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

---

## OSPF Fault Detection Configuration

---

- [Example: Configuring OSPF Timers on page 3875](#)
- [Example: Configuring BFD for OSPF on page 3881](#)
- [Example: Configuring BFD Authentication for OSPF on page 3887](#)

### Example: Configuring OSPF Timers

- [OSPF Timers Overview on page 3875](#)
- [Example: Configuring OSPF Timers on page 3876](#)

---

#### OSPF Timers Overview

OSPF routing devices constantly track the status of their neighbors, sending and receiving hello packets that indicate whether each neighbor still is functioning, and sending and receiving link-state advertisement (LSA) and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You configure OSPF timers on the interface of the routing device participating in OSPF. Depending on the timer, the configured interval must be the same on all routing devices on a shared network (area).

You can configure the following OSPF timers:

- **Hello interval**—Routing devices send hello packets at a fixed interval on all interfaces, including virtual links, to establish and maintain neighbor relationships. The hello interval specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. This interval must be the same on all routing devices on a shared network. By default, the routing device sends hello packets every 10 seconds (broadcast and point-to-point networks) and 30 seconds (nonbroadcast multiple access (NBMA) networks).
- **Poll interval**—(OSPFv2, Nonbroadcast networks only) Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The poll interval specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before establishing adjacency with a neighbor. By default, the routing device sends hello packets every 120 seconds until active neighbors are detected.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval.

- **LSA retransmission interval**—When a routing device sends LSAs to its neighbors, the routing device expects to receive an acknowledgment packet from each neighbor within a certain amount of time. The LSA retransmission interval specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting the LSA to an interface's neighbors. By default, the routing device waits 5 seconds for an acknowledgment before retransmitting the LSA.
- **Dead interval**—If a routing device does not receive a hello packet from a neighbor within a fixed amount of time, the routing device modifies its topology database to indicate

that the neighbor is nonoperational. The dead interval specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. This interval must be the same on all routing devices on a shared network. By default, this interval is four times the default hello interval, which is 40 seconds (broadcast and point-to-point networks) and 120 seconds (NBMA networks).

- Transit delay—Before a link-state update packet is propagated out of an interface, the routing device must increase the age of the packet. The transit delay sets the estimated time required to transmit a link-state update on the interface. By default, the transit delay is 1 second. You should never have to modify the transit delay time.

---

### Example: Configuring OSPF Timers

This example shows how to configure the OSPF timers.

- [Requirements on page 3876](#)
- [Overview on page 3876](#)
- [Configuration on page 3878](#)
- [Verification on page 3881](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### Overview

The default OSPF timer settings are optimal for most networks. However, depending on your network requirements, you might need to modify the timer settings. This example explains why you might need to modify the following timers:

- Hello interval
- Dead interval
- LSA retransmission interval
- Transit delay

#### Hello Interval and Dead Interval



The hello interval and the dead interval optimize convergence times by efficiently tracking neighbor status. By lowering the values of the hello interval and the dead interval, you can increase the convergence of OSPF routes if a path fails. These intervals must be the same on all routing devices on a shared network. Otherwise, OSPF cannot establish the appropriate adjacencies.

In the first example, you lower the hello interval to 2 seconds and the dead interval to 8 seconds on point-to-point OSPF interfaces **fe-0/0/1** and **fe-1/0/1** in area 0.0.0.0 by configuring the following settings:

- **hello-interval**—Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds. The range is from 1 through 255 seconds.
- **dead-interval**—Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. By default, the routing device waits 40 seconds (four times the hello interval). The range is 1 through 65,535 seconds.

#### LSA Retransmission Interval

The link-state advertisement (LSA) retransmission interval optimizes the sending and receiving of LSA and acknowledgement packets. You must configure the LSA retransmission interval to be equal to or greater than 3 seconds to avoid triggering a retransmit trap because the Junos OS delays LSA acknowledgments by up to 2 seconds. If you have a virtual link, you might find increased performance by increasing the value of the LSA retransmission interval.

In the second example, you increase the LSA retransmission timer to 8 seconds on OSPF interface **fe-0/0/1** in area 0.0.0.1 by configuring the following setting:

- **retransmit-interval**—Specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting LSA to an interface's neighbors. By default, the routing device retransmits LSAs to its neighbors every 5 seconds. The range is from 1 through 65,535 seconds.

#### Transit Delay

The transit delay sets the time the routing device uses to age a link-state update packet. If you have a slow link (for example, one with an average propagation delay of multiple seconds), you should increase the age of the packet by a similar amount. Doing this ensures that you do not receive a packet back that is younger than the original copy.

In the final example, you increase the transit delay to 2 seconds on OSPF interface **fe-1/0/1** in area 0.0.0.1. By configuring the following setting, this causes the routing device to age the link-state update packet by 2 seconds:

- **transit-delay**—Sets the estimated time required to transmit a link-state update on the interface. You should never have to modify the transit delay time. By default, the routing device ages the packet by 1 second. The range is from 1 through 65,535 seconds.

### Configuration

- [Configuring the Hello Interval and the Dead Interval on page 3878](#)
- [Controlling the LSA Retransmission Interval on page 3879](#)
- [Specifying the Transit Delay on page 3880](#)

### Configuring the Hello Interval and the Dead Interval

#### CLI Quick Configuration

To quickly configure the hello and dead intervals, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-0/0/1 dead-interval 8
set protocols ospf area 0.0.0.0 interface fe-1/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-1/0/1 dead-interval 8
```

#### Step-by-Step Procedure

To configure the hello and dead intervals:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
user@host# set interface fe-1/0/1
```

3. Configure the hello interval.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1 hello-interval 2
user@host# set interface fe-1/0/1 hello-interval 2
```

4. Configure the dead interval.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1 dead-interval 8
user@host# set interface fe-1/0/1 dead-interval 8
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```



**NOTE:** Repeat this entire configuration on all routing devices in a shared network.

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface fe-0/0/1.0 {
 hello-interval 2;
 dead-interval 8;
 }
 interface fe-1/0/1.0 {
 hello-interval 2;
 dead-interval 8;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

#### *Controlling the LSA Retransmission Interval*

**CLI Quick Configuration** To quickly configure the LSA retransmission interval, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-0/0/1 retransmit-interval 8
```

**Step-by-Step Procedure** To configure the LSA retransmission interval:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1
```

3. Configure the LSA retransmission interval.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1 retransmit-interval 8
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
 interface fe-0/0/1.0 {
 retransmit-interval 8;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### *Specifying the Transit Delay*

**CLI Quick Configuration** To quickly configure the transit delay, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-1/0/1 transit-delay 2
```

**Step-by-Step Procedure** To configure the transit delay:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-1/0/1
```

3. Configure the transit delay.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-1/0/1 transit-delay 2
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show protocols ospf
area 0.0.0.1 {
 interface fe-1/0/1.0 {
 transit-delay 2;
 }
}

```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

### Verifying the Timer Configuration

**Purpose** Verify that the interface for OSPF or OSPFv3 has been configured with the applicable timer values. Confirm that the Hello field, the Dead field, and the ReXmit field display the values that you configured.

**Action** From operational mode, enter the **show ospf interface detail** for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

## Example: Configuring BFD for OSPF

- [BFD for OSPF Overview on page 3881](#)
- [Example: Configuring BFD for OSPF on page 3883](#)

### BFD for OSPF Overview

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



**NOTE:** BFD is supported for OSPFv3 in Junos OS Release 9.3 and later.

You can configure the following BFD protocol settings:

- **detection-time threshold**—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- **full-neighbors-only**—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors. This setting is available in Junos OS Release 9.5 and later.
- **minimum-interval**—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
  - For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
  - For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms. In OSPFv3, BFD is always based in the Routing Engine, meaning that BFD is not distributed. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- **minimum-receive-interval**—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device expects to receive a hello packet from a neighbor with which it has established a BFD

session. You can also specify the minimum receive interval using the **minimum-interval** statement.

- **multiplier**—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- **no-adaptation**—Disables BFD adaptation. This setting disables BFD sessions from adapting to changing network conditions. This setting is available in Junos OS Release 9.0 and later.



**NOTE:** We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- **transmit-interval minimum-interval**—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the **minimum-interval** statement.
- **transmit-interval threshold**—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- **version**—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

### Example: Configuring BFD for OSPF

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

- [Requirements on page 3883](#)
- [Overview on page 3884](#)
- [Configuration on page 3885](#)
- [Verification on page 3886](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 3815](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).

### Overview

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.

BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the **bfd-liveness-detection** statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.





**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
  - For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
  - For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

### Configuration

#### CLI Quick Configuration

To quickly configure the BFD protocol for OSPF, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

#### Step-by-Step Procedure

To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```



**NOTE:** Repeat this entire configuration on the other neighboring interface.

---

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface fe-0/0/1.0 {
 bfd-liveness-detection {
 minimum-interval 300;
 multiplier 4;
 full-neighbors-only;
 }
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the BFD Sessions**

**Purpose** Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

**Action** From operational mode, enter the **show bfd session detail** command.

**Meaning** The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.
- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

**Related Documentation**

- [OSPF Configuration Overview](#)
- [BFD Authentication for OSPF Overview on page 3887](#)

## Example: Configuring BFD Authentication for OSPF

- [BFD Authentication for OSPF Overview on page 3887](#)
- [Configuring BFD Authentication for OSPF on page 3889](#)

### BFD Authentication for OSPF Overview

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over OSPFv2. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3888](#)
- [Security Authentication Keychains on page 3888](#)
- [Strict Versus Loose Authentication on page 3889](#)

### **BFD Authentication Algorithms**

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

### **Security Authentication Keychains**

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### ***Strict Versus Loose Authentication***

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Configuring BFD Authentication for OSPF**

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over OSPFv2. Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on OSPF:

- [Configuring BFD Authentication Parameters on page 3889](#)
- [Viewing Authentication Information for BFD Sessions on page 3890](#)

### ***Configuring BFD Authentication Parameters***

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the OSPFv2 protocol.
2. Associate the authentication keychain with the OSPFv2 protocol.
3. Configure the related security authentication keychain.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on an OSPF route or routing instance.

[edit]

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection
authentication algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified OSPF route or routing instance with the unique security authentication keychain attributes.

This keychain should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection
authentication keychain bfd-ospf
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
  - The matching keychain name as specified in Step 2.
  - At least one key, a unique integer between 0 and 63. Creating multiple keys enables multiple clients to use the BFD session.
  - The secret data used to allow access to the session.
  - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# authentication-key-chains key-chain bfd-ospf key 53 secret
9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set protocols ospf interface if2-ospf bfd-liveness-detection authentication
loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat the steps in this procedure to configure the other end of the BFD session.



**NOTE:** BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

### *Viewing Authentication Information for BFD Sessions*

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if2-ospf** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-ospf**. The authentication keychain is configured with two keys. Key 1 contains the secret data "**\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data "**\$9\$a5jiKW9L.reP38ny.TszF2/9**" and a start time of June 1, 2009, at 3:29:20 PM PST.

```

[edit protocols ospf]
area 0.0.0.1 {
 interface if2-ospf {
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-ospf;
 }
 }
 }
}
[edit security]
authentication key-chains {
 key-chain bfd-ospf {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
 start-time "2009-6-1.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9l.reP38ny.TszF2/9";
 start-time "2009-6-1.15:29:20 -0700";
 }
 }
}

```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured.

#### show bfd session detail

```
user@host# show bfd session detail
```

| Address   | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-----------|-------|------------|-------------|-------------------|------------|
| 10.9.1.33 | Up    | so-7/1/0.0 | 0.600       | 0.200             | 3          |

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**  
 Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Replicated

1 sessions, 1 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address   | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-----------|-------|------------|-------------|-------------------|------------|
| 10.9.1.33 | Up    | so-7/1/0.0 | 0.600       | 0.200             | 3          |

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**  
 keychain bfd-ospf, algo keyed-md5, mode loose

```
Session up time 3d 00:34
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated
Min async interval 0.200, min slow interval 1.000
Adaptive async tx interval 0.200, rx interval 0.200
Local min tx interval 0.200, min rx interval 0.200, multiplier 3
Remote min tx interval 0.100, min rx interval 0.100, multiplier 3
Threshold transmission interval 0.000, Threshold for detection time 0.000
Local discriminator 11, remote discriminator 80
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-ospf, algo keyed-sha-1, mode strict
1 sessions, 1 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

- Related Documentation**
- [OSPF Configuration Overview](#)
  - [BFD for OSPF Overview on page 3881](#)

---

## OSPF Redundancy Features Configuration

---

- [Examples: Configuring Graceful Restart for OSPF on page 3892](#)

### Examples: Configuring Graceful Restart for OSPF

- [Graceful Restart for OSPF Overview on page 3892](#)
- [Example: Configuring Graceful Restart for OSPF on page 3894](#)
- [Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 3898](#)
- [Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 3901](#)
- [Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 3904](#)

---

### Graceful Restart for OSPF Overview

---

Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting device and its neighbors continue forwarding packets without disrupting network performance. Because neighboring devices assist in the restart (these neighbors are called *helper routers*), the restarting device can quickly resume full operation without recalculating algorithms.



**NOTE:** On a broadcast link with a single neighbor, when the neighbor initiates an OSPFv3 graceful restart operation, the restart might be terminated at the point when the local routing device assumes the role of a helper. A change in the LSA is considered a topology change, which terminates the neighbor's restart operation.

Graceful restart is disabled by default. You can globally enable graceful restart for all routing protocols by including the **graceful-restart** statement at the **[edit routing-options]**



hierarchy level. To enable graceful restart specifically for OSPF, first you need to globally enable graceful restart for all routing protocols.

This topic describes the following information:

- [Helper Mode for Graceful Restart on page 3893](#)
- [Planned and Unplanned Graceful Restart on page 3894](#)

### ***Helper Mode for Graceful Restart***

When a device enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The device does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This device continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting device must send a grace LSA to all neighbors. In response, the helper routers enter helper mode (the ability to assist a neighboring device attempting a graceful restart) and send an acknowledgment back to the restarting device. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting device had remained in continuous OSPF operation.



**NOTE:** Helper mode is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode specifically for OSPF.

When the restarting device receives replies from all the helper routers, the restarting device selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting device receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting device or when the topology of the network changes, the helper routers also resume normal operation.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The Junos OS implementation is based on RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*. In restart signaling-based helper mode implementations, the restarting device informs its restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting device sends hello messages to its helper routers with the restart signal (RS) bit set in the hello packet header. When a helper router receives a hello packet with the RS bit set in the header, the helper router returns a hello message to the restarting device. The reply hello message from the helper router contains the ResyncState flag and the ResyncTimeout timer that enable the restarting device to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting device exits the restart mode.



**NOTE:** Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

### ***Planned and Unplanned Graceful Restart***

OSPF supports two types of graceful restart: planned and unplanned. During a planned restart, the restarting routing device informs the neighbors before restarting. The neighbors act as if the routing device is still within the network topology, and continue forwarding traffic to the restarting routing device. A grace period is set to specify when the neighbors should consider the restarting routing device as part of the topology. During an unplanned restart, the routing device restarts without warning.

### **Example: Configuring Graceful Restart for OSPF**

---

This example shows how to configure graceful restart specifically for OSPF.

- [Requirements on page 3894](#)
- [Overview on page 3894](#)
- [Configuration on page 3895](#)
- [Verification on page 3897](#)

#### ***Requirements***

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### ***Overview***

Graceful restart enables a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting routing device and its neighbors continue forwarding packets without disrupting network performance. By default, graceful restart is disabled. You can globally enable graceful restart for all routing protocols by including the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To enable graceful restart specifically for OSPF, first you need to globally enable graceful restart for all routing protocols.

The first example shows how to enable graceful restart and configure the optional settings for the grace period interval. In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. The grace period interval for OSPF graceful restart is determined as equal to or less than the sum of the **notify-duration** time interval and the **restart-duration** time interval. The grace period is the number of seconds that the routing device's neighbors continue to advertise the routing device as fully adjacent, regardless of the connection state between the routing device and its neighbors.

The **notify-duration** statement configures how long (in seconds) the routing device notifies helper routers that it has completed graceful restart by sending purged grace link-state advertisements (LSAs) over all interfaces. By default, the routing device sends grace LSAs for 30 seconds. The range is from 1 through 3600 seconds.

The **restart-duration** statement configures the amount of time the routing device waits (in seconds) to complete reacquisition of OSPF neighbors from each area. By default, the routing device allows 180 seconds. The range is from 1 through 3600 seconds.

The second example shows how to disable graceful restart for OSPF by including the **disable** statement.

### Configuration

- [Enabling Graceful Restart for OSPF on page 3895](#)
- [Disabling Graceful Restart for OSPF on page 3897](#)

### Enabling Graceful Restart for OSPF

#### CLI Quick Configuration

To quickly enable graceful restart for OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set routing-options graceful-restart
set protocols ospf graceful-restart restart-duration 190
set protocols ospf graceful-restart notify-duration 40
```

#### Step-by-Step Procedure

To enable graceful restart for OSPF:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Configure graceful restart globally

```
[edit]
user@host#edit routing-options graceful-restart
```

4. Configure OSPF graceful restart.

```
[edit]
user@host# edit protocols ospf graceful-restart
```

5. (Optional) Configure the restart duration time.

```
[edit protocols ospf graceful-restart]
user@host# set restart-duration 190
```

6. (Optional) Configure the notify duration time.

```
[edit protocols ospf graceful-restart]
user@host# set notify-duration 40
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf graceful-restart]
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
 unit 0 {
 family inet {
 address 10.0.0.4/32;
 }
 }
}
fe-1/1/2 {
 unit 0 {
 family inet {
 address 10.0.0.5/32;
 }
 }
}
user@host# show protocols ospf
graceful-restart {
 restart-duration 190;
 notify-duration 40;
}
area 0.0.0.0 {
 interface fe-1/1/1.0;
 interface fe-1/1/2.0;
}
```

To confirm an OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

**Disabling Graceful Restart for OSPF**

**CLI Quick Configuration** To quickly disable graceful restart for OSPF, copy the following command and paste it into the CLI.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

**Step-by-Step Procedure** To disable graceful restart for OSPF:

1. Disable graceful restart for the OSPF protocol only.  
This command does not affect the global graceful restart configuration setting.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

2. If you are done configuring the device, commit the configuration.

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

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
graceful-restart disable;
```

To confirm an OSPFv3 configuration, enter the `show protocols ospf3` command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 3897](#)
- [Verifying Graceful Restart Status on page 3898](#)

**Verifying the OSPF Graceful Restart Configuration**

**Purpose** Verify information about your OSPF graceful restart configuration.

**Action** From operational mode, enter the `show ospf overview` command for OSPFv2. Enter the `show ospf3 overview` command for OSPFv3.

**Meaning** The Restart field displays the status of graceful restart as either enabled or disabled. The Restart duration field displays how much time the restarted routing device requires to complete reacquisition of OSPF neighbors. The Restart grace period field displays how

much time the neighbors should consider the restarted routing device as part of the topology.

### ***Verifying Graceful Restart Status***

**Purpose** Verify the status of graceful restart.

**Action** From operational mode, enter the **show route instance detail** command.

**Meaning** The Restart State field displays Pending if the restart has not been completed or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have or have not yet completed graceful restart for the specified routing table.

### ***Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart***

---

This example shows how to disable and reenabling the helper mode capability for OSPFv2 graceful restart.

- [Requirements on page 3898](#)
- [Overview on page 3898](#)
- [Configuration on page 3899](#)
- [Verification on page 3901](#)

#### ***Requirements***

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### ***Overview***

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv2 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the standard OSPFv2 graceful restart helper capability by including the **helper-disable standard** statement. This configuration is useful if you have an environment that contains other vendor equipment that is configured for restart signaling-based graceful restart.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenabling the standard OSPFv2 restart helper capability that you disabled in the first example.

### Configuration

- [Disabling Helper Mode for OSPFv2 on page 3899](#)
- [Reenabling Helper Mode for OSPFv2 on page 3900](#)

### Disabling Helper Mode for OSPFv2

#### CLI Quick Configuration

To quickly enable graceful restart for OSPFv2 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart helper-disable standard
```

#### Step-by-Step Procedure

To enable graceful restart for OSPFv2 with helper mode disabled:

1. Configure the interfaces.
 

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```
2. Configure OSPFv2 on the interfaces
 

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```
3. Disable the OSPFv2 graceful restart helper capability.
 

If you disable the OSPFv2 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf graceful-restart helper-disable standard
```
4. If you are done configuring the device, commit the configuration.
 

```
[edit]
```

```
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
 unit 0 {
 family inet {
 address 10.0.0.4/32;
 }
 }
}
fe-1/1/2 {
 unit 0 {
 family inet {
 address 10.0.0.5/32;
 }
 }
}
user@host# show protocols ospf
graceful-restart {
 helper-disable {
 standard;
 }
}
area 0.0.0.0 {
 interface fe-1/1/1.0;
 interface fe-1/1/2.0;
}
```

#### *Reenabling Helper Mode for OSPFv2*

**CLI Quick Configuration** To quickly reenable standard helper-mode for OSPFv2, copy the following command and paste it into the CLI.

```
[edit]
delete protocols ospf graceful-restart helper-disable standard
```



**NOTE:** To reenable restart signaling-based helper mode, include the **restart-signaling** statement. To reenable both standard and restart signaling-based helper mode, include the **both** statement.

**Step-by-Step Procedure** To reenable standard helper mode for OSPFv2:

1. Delete the standard helper-mode statement from the OSPFv2 configuration.  

```
[edit]
user@host# delete protocols ospf graceful-restart helper-disable standard
```
2. If you are done configuring the device, commit the configuration.



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

**Results** After you reenable standard helper mode, the **show protocols ospf** command no longer displays the graceful restart configuration.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPFv2 Graceful Restart Configuration on page 3901](#)
- [Verifying Graceful Restart Status on page 3901](#)

### **Verifying the OSPFv2 Graceful Restart Configuration**

**Purpose** Verify information about your OSPFv2 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, the Graceful restart helper mode field displays the status of the standard helper mode capability as enabled or disabled, and the Restart-signaling helper mode field displays the status of the restart signaling-based helper mode as enabled or disabled. By default, both standard and restart signaling-based helper modes are enabled.

**Action** From operational mode, enter the **show ospf overview** command.

### **Verifying Graceful Restart Status**

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

### **Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart**

This example shows how to disable and reenable the helper mode capability for OSPFv3 graceful restart.

- [Requirements on page 3901](#)
- [Overview on page 3902](#)
- [Configuration on page 3902](#)
- [Verification on page 3904](#)

### **Requirements**

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#)
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 3815](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).

### Overview

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv3 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the OSPFv3 graceful restart helper capability by including the **helper-disable** statement.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenabling the OSPFv3 restart helper capability that you disabled in the first example.

### Configuration

- [Disabling Helper Mode for OSPFv3 on page 3902](#)
- [Reenabling Helper Mode for OSPFv3 on page 3903](#)

#### Disabling Helper Mode for OSPFv3

#### CLI Quick Configuration

To quickly enable graceful restart for OSPFv3 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
set interfaces fe-1/1/2 unit 0 family inet6 address 2002:0a00:0005::
set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
set protocols ospf3 graceful-restart helper-disable
```

#### Step-by-Step Procedure

To enable graceful restart for OSPFv3 with helper mode disabled:

1. Configure the interfaces.

```
[edit]
```

```

user@host# set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
user@host# set interfaces fe-1/1/1 unit 0 family inet address 2002:0a00:0005::

```

2. Configure OSPFv3 on the interfaces

```

[edit]
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/2

```

3. Disable the OSPFv3 graceful restart helper capability.  
If you disable the OSPFv3 graceful restart helper capability, you cannot disable strict LSA checking.

```

[edit]
user@host# set protocols ospf3 graceful-restart helper-disable

```

4. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

**Results** Confirm your configuration by entering the **show interfaces** and the **show protocols ospf3** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show interfaces
fe-1/1/1 {
 unit 0 {
 family inet6 {
 address 2002:0a00:0004::/128;
 }
 }
}
fe-1/1/2 {
 unit 0 {
 family inet6 {
 address 2002:0a00:0005::/128;
 }
 }
}
user@host# show protocols ospf3
graceful-restart {
 helper-disable;
}
area 0.0.0.0 {
 interface fe-1/1/1.0;
 interface fe-1/1/2.0;
}

```

#### *Reenabling Helper Mode for OSPFv3*

**CLI Quick Configuration** To quickly reenable helper-mode for OSPFv3, copy the following command and paste it into the CLI.

```

[edit]
delete protocols ospf3 graceful-restart helper-disable

```

**Step-by-Step Procedure**

To reenable helper mode for OSPFv3:

1. Delete the standard helper-mode statement from the OSPFv3 configuration.

[edit]

user@host# **delete protocols ospf3 graceful-restart helper-disable**

2. If you are done configuring the device, commit the configuration.

[edit]

user@host# **commit**

**Results** After you reenable standard helper mode, the **show protocols ospfs** command no longer displays the graceful restart configuration.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPFv3 Graceful Restart Configuration on page 3904](#)
- [Verifying Graceful Restart Status on page 3904](#)

**Verifying the OSPFv3 Graceful Restart Configuration**

**Purpose** Verify information about your OSPFv3 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, and the Helper mode field displays the status of the helper mode capability as either enabled or disabled.

**Action** From operational mode, enter the **show ospf3 overview** command.

**Verifying Graceful Restart Status**

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

---

**Example: Disabling Strict LSA Checking for OSPF Graceful Restart**

This example shows how to disable strict link-state advertisement (LSA) checking for OSPF graceful restart.

- [Requirements on page 3905](#)
- [Overview on page 3905](#)
- [Configuration on page 3905](#)
- [Verification on page 3907](#)

### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 3815](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 3817](#).

### Overview

You can disable strict LSA checking to prevent the termination of graceful restart by a helping router. You might configure this option for interoperability with other vendor devices. The OSPF graceful restart helper capability must be enabled if you disable strict LSA checking. By default, LSA checking is enabled.

In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable strict LSA checking by including the **no-strict-lsa-checking** statement.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

### Configuration

#### CLI Quick Configuration

To quickly enable graceful restart for OSPF with strict LSA checking disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart no-strict-lsa-checking
```

#### Step-by-Step Procedure

To enable graceful restart for OSPF with strict LSA checking disabled:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable strict LSA checking.  
If you disable the strict LSA checking, OSPF graceful restart helper capability must be enabled (which is the default behavior).

```
[edit]
user@host# set protocols ospf graceful-restart no-strict-lsa-checking
```

4. If you are done configuring the device, commit the configuration.

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

**Results** Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
 unit 0 {
 family inet {
 address 10.0.0.4/32;
 }
 }
}
fe-1/1/2 {
 unit 0 {
 family inet {
 address 10.0.0.5/32;
 }
 }
}
user@host# show protocols ospf
graceful-restart {
 no-strict-lsa-checking;
}
area 0.0.0.0 {
 interface fe-1/1/1.0;
 interface fe-1/1/2.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 3907](#)
- [Verifying Graceful Restart Status on page 3907](#)

### **Verifying the OSPF Graceful Restart Configuration**

**Purpose** Verify information about your OSPF graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

### **Verifying Graceful Restart Status**

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

**Related Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)
- [Graceful Restart Concepts on page 2269](#) in the *Junos OS High Availability Configuration Guide*

## **OSPF Traffic Engineering Configuration**

- [Examples: Configuring OSPF Traffic Engineering on page 3907](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 3916](#)

### **Examples: Configuring OSPF Traffic Engineering**

- [OSPF Support for Traffic Engineering on page 3908](#)
- [Example: Enabling OSPF Traffic Engineering Support on page 3910](#)
- [Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface on page 3914](#)

## OSPF Support for Traffic Engineering

---

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path.

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of OSPF. When traffic engineering is enabled on the routing device, you can enable OSPF traffic engineering support. When you enable traffic engineering for OSPF, the shortest-path-first (SPF) algorithm takes into account the various label-switched paths (LSPs) configured under MPLS and configures OSPF to generate opaque link-state advertisements (LSAs) that carry traffic engineering parameters. The parameters are used to populate the traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. The Constrained Shortest Path First (CSPF) algorithm uses the traffic engineering database to compute the paths that MPLS LSPs take. RSVP uses this path information to set up LSPs and to reserve bandwidth for them.

By default, traffic engineering support is disabled. To enable traffic engineering, include the **traffic-engineering** statement. You can also configure the following OSPF traffic engineering extensions:

- **advertise-unnumbered-interfaces**—(OSPFv2 only) Advertises the link-local identifier in the link-local traffic engineering LSA packet. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477, *Signalling Unnumbered Links in Resource Reservation Protocol - Traffic Engineering (RSVP-TE)*.
- **credibility-protocol-preference**—(OSPFv2 only) Assigns a credibility value to OSPF routes in the traffic engineering database. By default, Junos OS prefers IS-IS routes in the traffic engineering database over other interior gateway protocol (IGP) routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure OSPF to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration.
- **ignore-lsp-metrics**—Ignores RSVP LSP metrics in OSPF traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs. This option avoids mutual dependency between OSPF and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date. In addition, If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.



- **multicast-rpf-routes**—(OSPFv2 only) Installs unicast IPv4 routes (not LSPs) in the multicast routing table (**inet.2**) for multicast reverse-path forwarding (RPF) checks. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check if the packet is coming in on an interface that is also sending data back to the packet source.
- **no-topology**—(OSPFv2 only) To disable the dissemination of link-state topology information. If disabled, traffic engineering topology information is no longer distributed within the OSPF area.
- **shortcuts**—Configures OSPF to use MPLS LSPs as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the **inet.3** routing table, and shortcut routes calculated through OSPFv3 are installed in the **inet6.3** routing table.



**NOTE:** Whenever possible, use OSPF IGP shortcuts configured at the `[edit protocols mpls traffic-engineering bgp-igp]` hierarchy level instead of traffic engineering shortcuts configured at the `[edit protocols (ospf | ospf3) traffic-engineering shortcuts]` hierarchy level.

If you configure OSPF IGP shortcuts, **inet.3** routes are moved into the **inet.0** routing table. In addition, you can verify the data path using `ping` or `traceroute` commands since the ping and traceroute packets get tunneled into the LSP. In case of a VPN enabled device, we recommend using `[edit protocols mpls traffic-engineering bgp-igp-both-ribs]` because BGP next-hop resolution for VPN prefixes relies on entries in the **inet.3** table.

If you configure traffic engineering shortcuts, OSPF treats the MPLS LSP as a candidate next hop and installs the routes in the **inet.3** (for OSPFv2) and **inet6.3** (for OSPFv3) routing tables. The only use for these tables is to allow BGP to perform next-hop resolution. In addition, you cannot verify the data path of these routes using `ping` or `traceroute` commands because the ping and traceroute packets get tunneled into the LSP.

- **lsp-metric-info-summary**—Advertises the LSP metric in summary LSAs to treat the LSP as a link. This configuration allows other routing devices in the network to use this LSP. To accomplish this, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

When you enable traffic engineering on the routing device, you can also configure an OSPF metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the traffic engineering database. Its value does not affect normal OSPF forwarding.



**CAUTION:** When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

---

### Example: Enabling OSPF Traffic Engineering Support

---

This example shows how to enable OSPF traffic engineering support to advertise the label-switched path (LSP) metric in summary link-state advertisements (LSAs).

- [Requirements on page 3910](#)
- [Overview on page 3910](#)
- [Configuration on page 3911](#)
- [Verification on page 3914](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure BGP per your network requirements. See the *Junos OS Routing Protocols Configuration Guide*
- Configure MPLS per your network requirements. See the *Junos OS MPLS Applications Configuration Guide*.

#### Overview

You can configure OSPF to treat an LSP as a link and have other routing devices in the network use this LSP. To accomplish this, you configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

In this example, there are four routing devices in area 0.0.0.0, and you want OSPF to treat the LSP named R1-to-R4 that goes from the ingress Device R1 to the egress Device R4 as a link.

For OSPF, you enable traffic engineering on all four routing devices in the area by including the **traffic-engineering** statement. This configuration ensures that the shortest-path-first (SPF) algorithm takes into account the LSPs configured under MPLS and configures OSPF to generate LSAs that carry traffic engineering parameters. You further ensure that OSPF uses the MPLS LSP as the next hop and advertises the LSP metric in summary LSAs, by including the optional **shortcuts lsp-metric-into-summary** statement on the ingress Device R1.

For MPLS, you enable traffic engineering so that MPLS performs traffic engineering on both BGP and IGP destinations by including the **traffic-engineering bgp-igp** statement, and you include the LSP named R1-to-R4 by including the **label-switched-path lsp-path-name to address** statement on the ingress Device R1. The address specified in the **to** statement on the ingress Device R1 must match the router ID of the egress Device R4 for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. In this example, the router ID of the egress Device R4 is 10.0.0.4.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

**CLI Quick Configuration** To quickly enable OSPF traffic engineering support to advertise the LSP metric in summary LSAs, copy the following commands and paste them into the CLI.

Configuration on R1:

```
[edit]
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

Configuration on R2:

```
[edit]
set routing-options router-id 10.0.0.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R3:

```
[edit]
set routing-options router-id 10.0.0.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R4:

```
[edit]
set routing-options router-id 10.0.0.4
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

**Step-by-Step Procedure** To enable OSPF traffic engineering support to advertise LSP metrics in summary LSAs:

1. Configure the router ID.

```
[edit]
```

```
user@R1# set routing-options router-id 10.0.0.1
[edit]
user@R2# set routing-options router-id 10.0.0.2
[edit]
user@R3# set routing-options router-id 10.0.0.3
[edit]
user@R4# set routing-options router-id 10.0.0.4
```

2. Configure the OSPF area and add the interfaces.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface all
user@R1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface all
user@R2# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface all
user@R3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface all
user@R4# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

3. Enable OSPF traffic engineering.

```
[edit]
user@R1 set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
[edit]
user@R2 set protocols ospf traffic-engineering
[edit]
user@R3 set protocols ospf traffic-engineering
[edit]
user@R4 set protocols ospf traffic-engineering
```

4. On Device R1, configure MPLS traffic engineering.

```
[edit]
user@R1 set protocol mpls traffic-engineering bgp-igp
user@R1 set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

5. If you are done configuring the devices, commit the configuration.

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

**Results** Confirm your configuration by entering the **show routing-options**, **show protocols ospf**, and **show protocols mpls** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@host# show routing-options
router-id 10.0.0.1;

user@host# show protocols ospf
 traffic-engineering {
 shortcuts lsp-metric-into-summary;
 }
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
}

user@host# show protocols mpls
 traffic-engineering bgp-igp;
 label-switched-path R1-to-R4 {
 to 10.0.0.4;
 }
}
```

Output for R2:

```
user@host# show routing-options
router-id 10.0.0.2;

user@host# show protocols ospf
 traffic-engineering;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
}
```

Output for R3:

```
user@host# show routing-options
router-id 10.0.0.3;

user@host# show protocols ospf
 traffic-engineering;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
}
```

Output for R4:

```
user@host# show routing-options
router-id 10.0.0.4;
```

```
user@host# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options**, **show protocols ospf3**, and **show protocols mpls** commands.

### ***Verification***

Confirm that the configuration is working properly.

- [Verifying the Traffic Engineering Capability for OSPF on page 3914](#)
- [Verifying OSPF Entries in the Traffic Engineering Database on page 3914](#)
- [Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF on page 3914](#)

### ***Verifying the Traffic Engineering Capability for OSPF***

**Purpose** Verify that traffic engineering has been enabled for OSPF. By default, traffic engineering is disabled.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** for OSPFv3.

### ***Verifying OSPF Entries in the Traffic Engineering Database***

**Purpose** Verify the OSPF information in the traffic engineering database. The Protocol field displays OSPF and the area from which the information was learned.

**Action** From operational mode, enter the **show ted database** command.

### ***Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF***

**Purpose** Verify that OSPF is reporting node information. The Protocol name field displays OSPF and the area from which the information was learned.

**Action** From operational mode, enter the **show ted protocol** command.

### **Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface**

This example shows how to configure the OSPF metric value used for traffic engineering.

- [Requirements on page 3915](#)
- [Overview on page 3915](#)
- [Configuration on page 3915](#)
- [Verification on page 3916](#)

### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure OSPF for traffic engineering. See “[Example: Enabling OSPF Traffic Engineering Support](#)” on page 3910

### Overview

You can configure an OSPF metric that is used exclusively for traffic engineering. To modify the default value of the traffic engineering metric, include the **te-metric** statement. The OSPF traffic engineering metric does not affect normal OSPF forwarding. By default, the traffic engineering metric is the same value as the OSPF metric. The range is 1 through 65,535.

In this example, you configure the OSPF traffic engineering metric on OSPF interface **fe-0/1/1** in area 0.0.0.0.

### Configuration

**CLI Quick Configuration** To quickly configure the OSPF traffic engineering metric for a specific interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/1/1 te-metric 10
```

**Step-by-Step Procedure** To configure an OSPF traffic engineering metric for a specific interface used only for traffic engineering:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the traffic engineering metric of the OSPF network segments.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/1/1 te-metric 10
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
 interface fe-0/1/1.0 {
 te-metric 10;
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the Configured Traffic Engineering Metric**

**Purpose** Verify the traffic engineering metric value. Confirm that Metric field displays the configured traffic engineering metric.

**Action** From operational mode, enter the **show ted database extensive** command.

**Related Documentation**

- [OSPF Configuration Overview](#)
- [Junos OS MPLS Applications Configuration Guide](#)

## **Example: Configuring OSPF Passive Traffic Engineering Mode**

- [OSPF Passive Traffic Engineering Mode on page 3916](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 3917](#)

### **OSPF Passive Traffic Engineering Mode**

---

Ordinarily, interior routing protocols such as OSPF are not run on links between autonomous systems. However, for inter-AS traffic engineering to function properly, information about the inter-AS link—in particular, the address on the remote interface—must be made available inside the autonomous system (AS). This information is not normally included either in the external BGP (EBGP) reachability messages or in the OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include it in the traffic engineering database. OSPF traffic engineering mode allows MPLS label-switched paths (LSPs) to dynamically discover OSPF AS boundary routers and to allow routers to establish a traffic engineering LSP across multiple autonomous systems.



### Example: Configuring OSPF Passive Traffic Engineering Mode

This example shows how to configure OSPF passive mode for traffic engineering on an inter-AS interface. The AS boundary router link between the EBGP peers must be a directly connected link and must be configured as a passive traffic engineering link.

- [Requirements on page 3917](#)
- [Overview on page 3917](#)
- [Configuration on page 3918](#)
- [Verification on page 3918](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure BGP per your network requirements. See the *Junos OS Routing Protocols Configuration Guide*.
- Configure the LSP per your network requirements. See the *Junos OS MPLS Applications Configuration Guide*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

#### Overview

You can configure OSPF passive mode for traffic engineering on an inter-AS interface. The address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGP link. In this example, you configure interface **so-1/1/0** in area 0.0.0.1 as the inter-AS link to distribute traffic engineering information with OSPF within the AS and include the following settings:

- **passive**—Advertises the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.
- **traffic-engineering**—Configures an interface in OSPF passive traffic-engineering mode to enable dynamic discovery of OSPF AS boundary routers. By default, OSPF passive traffic-engineering mode is disabled.
- **remote-node-id**—Specifies the IP address at the far end of the inter-AS link. In this example, the remote IP address is 192.168.207.2.

### Configuration

To quickly configure OSPF passive mode for traffic engineering, copy the following command, remove any line breaks, and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

#### Step-by-Step Procedure

To configure OSPF passive traffic engineering mode:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1
```

2. Configure interface **so-1/1/0** as a passive interface configured for traffic engineering, and specify the IP address at the far end of the inter-AS link.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

#### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
 interface so-1/1/0.0 {
 passive {
 traffic-engineering {
 remote-node-id 192.168.207.2;
 }
 }
 }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

### Verifying the Status of OSPF Interfaces

|                              |                                                                                                                                                                                                                       |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>               | Verify the status of OSPF interfaces. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.                         |
| <b>Action</b>                | From operational mode, enter the <b>show ospf interface detail</b> command for OSPFv2, and enter the <b>show ospf3 interface detail</b> command for OSPFv3.                                                           |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>OSPF Configuration Overview</i></li> <li>• <a href="#">About OSPF Interfaces on page 3838</a></li> <li>• <i>Junos OS MPLS Applications Configuration Guide</i></li> </ul> |

## OSPF Database Protection Configuration

- [Example: Configuring OSPF Database Protection on page 3919](#)

### Example: Configuring OSPF Database Protection

- [OSPF Database Protection Overview on page 3919](#)
- [Configuring OSPF Database Protection on page 3920](#)

### OSPF Database Protection Overview

OSPF database protection allows you to limit the number of link-state advertisements (LSAs) not generated by the local router in a given OSPF routing instance, helping to protect the link-state database from being flooded with excessive LSAs. This feature is particularly useful if VPN routing and forwarding is configured on your provider edge and customer edge routers using OSPF as the routing protocol. An overrun link-state database on the customer edge router can exhaust resources on the provider edge router and impact the rest of the service provider network.

When you enable OSPF database protection, the maximum number of LSAs you specify includes all LSAs whose advertising router ID is not equal to the local router ID (nonself-generated LSAs). These might include external LSAs as well as LSAs with any scope such as the link, area, and autonomous system (AS).

Once the specified maximum LSA count is exceeded, the database typically enters into the ignore state. In this state, all neighbors are brought down, and nonself-generated LSAs are destroyed. In addition, the database sends out hellos but ignores all received packets. As a result, the database does not form any full neighbors, and therefore does not learn about new LSAs. However, if you have configured the **warning-only** option, only a warning is issued and the database does not enter the ignore state but continues to operate as before.

You can also configure one or more of the following options:

- A warning threshold for issuing a warning message before the LSA limit is reached.
- An ignore state time during which the database must remain in the ignore state and after which normal operations can be resumed.
- An ignore state count that limits the number of times the database can enter the ignore state, after which it must enter the isolate state. The isolate state is very similar to the ignore state, but has one important difference: once the database enters the isolate state, it must remain there until you issue a command to clear database protection before it can return to normal operations.
- A reset time during which the database must stay out of the ignore or isolate state before it is returned to a normal operating state.

### Configuring OSPF Database Protection

---

By configuring OSPF database protection, you can help prevent your OSPF link-state database from being overrun with excessive LSAs that are not generated by the local router. You specify the maximum number of LSAs whose advertising router ID is not the same as the local router ID in an OSPF instance. This feature is particularly useful if your provider edge and customer edge routers are configured with VPN routing and forwarding using OSPF.

OSPF database protection is supported on:

- Logical systems
- All routing instances supported by OSPFv2 and OSPFv3
- OSPFv2 and OSPFv3 topologies
- OSPFv3 realms

To configure OSPF database protection:

1. Include the **database-protection** statement at one of the following hierarchy levels:
  - [edit protocols ospf | ospf3]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf |ospf3)]
  - [edit routing-instances *routing-instance-name* protocols (ospf |ospf3)]
  - [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast)]
2. Include the **maximum-lsa *number*** statement.



**NOTE:** The **maximum-lsa** statement is mandatory, and there is no default value for it. If you omit this statement, you cannot configure OSPF database protection.

---

3. (Optional) Include the following statements:

- **ignore-count *number***—Specify the number of times the database can enter the ignore state before it goes into the isolate state.
  - **ignore-time *seconds***—Specify the time limit the database must remain in the ignore state before it resumes regular operations.
  - **reset-time *seconds***—Specify the time during which the database must operate without being in either the ignore or isolate state before it is reset to a normal operating state.
  - **warning-threshold *percent***—Specify the percent of the maximum LSA number that must be exceeded before a warning message is issued.
4. (Optional) Include the **warning-only** statement to prevent the database from entering the ignore state or isolate state when the maximum LSA count is exceeded.



**NOTE:** If you include the **warning-only** statement, values for the other optional statements at the same hierarchy level are not used when the maximum LSA number is exceeded.

5. Verify your configuration by checking the database protection fields in the output of the **show ospf overview** command.

**Related  
Documentation**

- [OSPF Overview on page 3798](#)
- [OSPF Configuration Overview](#)

## OSPF Policy Configuration

- [Examples: Configuring OSPF Routing Policy on page 3921](#)
- [Examples: Configuring Routing Policy for Network Summaries on page 3937](#)

### Examples: Configuring OSPF Routing Policy

- [Understanding OSPF Routing Policy on page 3921](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 3923](#)
- [Example: Redistributing Static Routes into OSPF on page 3927](#)
- [Example: Configuring an OSPF Import Policy on page 3929](#)
- [Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF on page 3933](#)

#### Understanding OSPF Routing Policy

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration. Once a policy is created and named, it must be applied before it is active.

In the **import** statement, you list the name of the routing policy used to filter OSPF external routes from being installed into the routing tables of OSPF neighbors. You can filter the routes, but not link-state address (LSA) flooding. An external route is a route that is outside the OSPF Autonomous System (AS). The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF.

By default, if a routing device has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

This topic describes the following information:

- [Routing Policy Terms on page 3922](#)
- [Routing Policy Match Conditions on page 3922](#)
- [Routing Policy Actions on page 3923](#)

### ***Routing Policy Terms***

Routing policies are made up of one or more terms. A term is a named structure in which match conditions and actions are defined. You can define one or more terms. The name can contain letters, numbers, and hyphens ( - ) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

### ***Routing Policy Match Conditions***

A match condition defines the criteria that a route must match for an action to take place. You can define one or more match conditions for each term. If a route matches all of the match conditions for a particular term, the actions defined for that term are processed.

Each term can include two statements, **from** and **to**, that define the match conditions:

- In the **from** statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The **from** statement is optional. If you omit the **from** and the **to** statements, all routes are considered to match.



**NOTE:** In export policies, omitting the **from** statement from a routing policy term might lead to unexpected results. For more information, see *Applying Routing Policies and Policy Chains to Routing Protocols* in the *Routing Policy Configuration Guide*.

- In the **to** statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The order of the match conditions in a term is not important because a route must match all match conditions in a term for an action to be taken.

For a complete list of match conditions, see *Configuring Match Conditions in Routing Policy Terms* in the *Routing Policy Configuration Guide*.

### **Routing Policy Actions**

An action defines what the routing device does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.

The **then** statement is optional. If you omit it, one of the following occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the **accept** or **reject** action specified by the default policy is executed.

For a complete list of routing policy actions, see *Configuring Actions in Routing Policy Terms* in the *Routing Policy Configuration Guide*.

### **Example: Injecting OSPF Routes into the BGP Routing Table**

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 3924](#)
- [Overview on page 3924](#)
- [Configuration on page 3924](#)

- [Verification on page 3926](#)
- [Troubleshooting on page 3926](#)

### **Requirements**

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See “[Example: Configuring External BGP Point-to-Point Peer Sessions](#)” on page 3150.
- Configure interior gateway protocol (IGP) sessions between peers.

### **Overview**

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

### **Configuration**

- [Configuring the Routing Policy on page 3924](#)
- [Configuring Tracing for the Routing Policy on page 3925](#)

#### **Configuring the Routing Policy**

#### **CLI Quick Configuration**

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

#### **Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```

2. Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

3. Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```



- Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

- Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

**Results** Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
 term injectterm1 {
 from {
 protocol ospf;
 area 0.0.0.1;
 }
 then accept;
 }
}

user@host# show protocols bgp
export injectpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Configuring Tracing for the Routing Policy*

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

- Include a trace action in the policy.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# then trace
```

- Configure the tracing file for the output.

```
[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
```

```
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy
```

**Results** Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
 term injectterm1 {
 then {
 trace;
 }
 }
}

user@host# show routing-options
traceoptions {
 file ospf-bgp-policy-log size 5m files 5;
 flag policy;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying That the Expected BGP Routes Are Present**

**Purpose** Verify the effect of the export policy.

**Action** From operational mode, enter the **show route** command.

### **Troubleshooting**

- [Using the show log Command to Examine the Actions of the Routing Policy on page 3926](#)

### **Using the show log Command to Examine the Actions of the Routing Policy**

**Problem** The routing table contains unexpected routes, or routes are missing from the routing table.

**Solution** If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

### Example: Redistributing Static Routes into OSPF

This example shows how to create a policy that redistributes static routes into OSPF.

- [Requirements on page 3927](#)
- [Overview on page 3927](#)
- [Configuration on page 3927](#)
- [Verification on page 3929](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure static routes. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Configuration Guide*.

#### Overview

In this example, you create a routing policy called `exportstatic1` and a routing term called `exportstatic1`. The policy injects static routes into OSPF. This example includes the following settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens ( - ) and be up to 255 characters long.
- **term**—Defines the match condition and applicable actions for the routing policy. The term name can contain letters, numbers, and hyphens ( - ) and be up to 255 characters long. You specify the name of the term and define the criteria that an incoming route must match by including the **from** statement and the action to take if the route matches the conditions by including the **then** statement. In this example you specify the static protocol match condition and the accept action.
- **export**—Applies the export policy you created to be evaluated when routes are being exported from the routing table into OSPF.

#### Configuration

##### CLI Quick Configuration

To quickly create a policy that injects static routes into OSPF, copy the following commands and paste them into the CLI.

[edit]

```
set policy-options policy-statement exportstatic1 term exportstatic1 from protocol static
set policy-options policy-statement exportstatic1 term exportstatic1 then accept
set protocols ospf export exportstatic1
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To inject static routes into OSPF:

1. Create the routing policy.  

```
[edit]
user@host# edit policy-options policy-statement exportstatic1
```
2. Create the policy term.  

```
[edit policy-options policy-statement exportstatic1]
user@host# set term exportstatic1
```
3. Specify static as a match condition.  

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set from protocol static
```
4. Specify that the route is to be accepted if the previous condition is matched.  

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set then accept
```
5. Apply the routing policy to OSPF.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

---

- ```
[edit]  
user@host# set protocols ospf export exportstatic1
```
6. If you are done configuring the device, commit the configuration.

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

Results Confirm your configuration by entering the `show policy-options` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options  
policy-statement exportstatic1 {  
  term exportstatic1 {  
    from protocol static;  
    then accept;  
  }  
}  
  
user@host# show protocols ospf  
export exportstatic1;
```

To confirm your OSPFv3 configuration, enter the `show policy-options` and the `show protocols ospf3` commands.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Static Routes Are Present on page 3929](#)
- [Verifying That AS External LSAs Are Added to the Routing Table on page 3929](#)

Verifying That the Expected Static Routes Are Present

Purpose Verify the effect of the export policy.

Action From operational mode, enter the **show route** command.

Verifying That AS External LSAs Are Added to the Routing Table

Purpose On the routing device where you configured the export policy, verify that the routing device originates an AS external LSA for the static routes that are added to the routing table.

Action From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

Example: Configuring an OSPF Import Policy

This example shows how to create an OSPF import policy. OSPF import policies apply to external routes only. An external route is a route that is outside the OSPF autonomous system (AS).

- [Requirements on page 3929](#)
- [Overview on page 3929](#)
- [Configuration on page 3930](#)
- [Verification on page 3933](#)

Requirements

Before you begin:

- Configure static routes. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Configuration Guide*.
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 3815](#).

Overview

External routes are learned by AS boundary routers. External routes can be advertised throughout the OSPF domain if you configure the AS boundary router to redistribute the

route into OSPF. An external route might be learned by the AS boundary router from a routing protocol other than OSPF, or the external route might be a static route that you configure on the AS boundary router.

For OSPFv3, the link-state advertisement (LSA) is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An area border router (ABR) originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements. The filtering is done only on external routes in OSPF. The intra-area and interarea routes are not considered for filtering. The default action is to accept the route when the route does not match the policy.

This example includes the following OSPF policy settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens (-) and be up to 255 characters long.
- **export**—Applies the export policy you created to be evaluated when network summary LSAs are flooded into an area. In this example, the export policy is named `export_static`.
- **import**—Applies the import policy you created to prevent external routes from being added to the routing table. In this example, the import policy is named `filter_routes`.

The devices you configure in this example represent the following functions:

- **R1**—Device R1 is in area 0.0.0.0 and has a direct connection to device R2. R1 has an OSPF export policy configured. The export policy redistributes static routes from R1's routing table into R1's OSPF database. Because the static route is in R1's OSPF database, the route is advertised in an LSA to R1's OSPF neighbor. R1's OSPF neighbor is device R2.
- **R2**—Device R2 is in area 0.0.0.0 and has a direct connection to device R1. R2 has an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network and prevents the static route from being installed in R2's routing table. R2's OSPF neighbor is device R1.

Configuration

CLI Quick Configuration

To quickly configure an OSPF import policy, copy the following commands, removing any line breaks, and then paste the commands into the CLI.

Configuration on Device R1:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
set protocols ospf export export_static
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement export_static from protocol static
set policy-options policy-statement export_static then accept
```

Configuration on Device R2:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
set protocols ospf import filter_routes
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement filter_routes from route-filter 10.0.16.0/30 exact
set policy-options policy-statement filter_routes then reject
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure an OSPF import policy:

1. Configure the interfaces.

```
[edit]
user@R1# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R2# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

3. On R1, redistribute the static route into OSPF.

```
[edit]
user@R1# set protocols ospf export export_static
user@R1# set policy-options policy-statement export_static from protocol static
user@R1# set policy-options policy-statement export_static then accept
```

4. On R2, configure the OSPF import policy.

```
[edit]
user@R2# set protocols ospf import filter_routes
user@R2# set policy-options policy-statement filter_routes from route-filter
10.0.16.0/30 exact
user@R2# set policy-options policy-statement filter_routes then reject
```

5. If you are done configuring the devices, commit the configuration.

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

Results Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display

the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}

user@R1# show policy-options
policy-statement export_static {
  from protocol static;
  then accept;
}

user@R1# show protocols ospf
export export_static;
area 0.0.0.0 {
  interface so-0/2/0.0;
}
```

Output for R2:

```
user@R2# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.2.2/30;
    }
  }
}

user@R2# show policy-options
policy-statement filter_routes {
  from {
    route-filter 10.0.16.0/30 exact;
  }
  then reject;
}

user@R2# show protocols ospf
import filter_routes;
area 0.0.0.0 {
  interface so-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, **show routing-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 3933](#)
- [Verifying the Routing Table on page 3933](#)

Verifying the OSPF Database

Purpose Verify that OSPF is advertising the static route in the OSPF database.

Action From operational mode, enter the **show ospf database** for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

Verifying the Routing Table

Purpose Verify the entries in the routing table.

Action From operational mode, enter the **show route** command.

Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF

This example shows how to create an OSPF import policy that prioritizes specific prefixes learned through OSPF.

- [Requirements on page 3933](#)
- [Overview on page 3933](#)
- [Configuration on page 3934](#)
- [Verification on page 3937](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Junos® OS Network Interfaces*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 3810.
- Control OSPF designated router election See “[Example: Controlling OSPF Designated Router Election](#)” on page 3812
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 3815 .
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 3817.

Overview

In a network with a large number of OSPF routes, it can be useful to control the order in which routes are updated in response to a network topology change. In Junos OS

Release 9.3 and later, you can specify a priority of high, medium, or low for prefixes included in an OSPF import policy. In the event of an OSPF topology change, high priority prefixes are updated in the routing table first, followed by medium and then low priority prefixes.

OSPF import policy can only be used to set priority or to filter OSPF external routes. If an OSPF import policy is applied that results in a **reject** terminating action for a nonexternal route, then the **reject** action is ignored and the route is accepted anyway. By default, such a route is now installed in the routing table with a priority of low. This behavior prevents traffic black holes, that is, silently discarded traffic, by ensuring consistent routing within the OSPF domain.

In general, OSPF routes that are not explicitly assigned a priority are treated as priority medium, except for the following:

- Summary discard routes have a default priority of low.
- Local routes that are not added to the routing table are assigned a priority of low.
- External routes that are rejected by import policy and thus not added to the routing table are assigned a priority of low.

Any available match criteria applicable to OSPF routes can be used to determine the priority. Two of the most commonly used match criteria for OSPF are the **route-filter** and **tag** statements.

In this example, the routing device is in area 0.0.0.0, with interfaces **fe-0/1/0** and **fe-1/1/0** connecting to neighboring devices. You configure an import routing policy named **ospf-import** to specify a priority for prefixes learned through OSPF. Routes associated with these prefixes are installed in the routing table in the order of the prefixes' specified priority. Routes matching **200.3.0.0/16 orlonger** are installed first because they have a priority of **high**. Routes matching **200.2.0.0/16 orlonger** are installed next because they have a priority of **medium**. Routes matching **200.1.0.0/16 orlonger** are installed last because they have a priority of **low**. You then apply the import policy to OSPF.



NOTE: The priority value takes effect when a new route is installed, or when there is a change to an existing route.

Configuration

CLI Quick Configuration

To quickly configure an OSPF import policy that prioritizes specific prefixes learned through OSPF, copy the following commands, removing any line breaks, and then paste the commands into the CLI.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.5/30
set policy-options policy-statement ospf-import term t1 from route-filter 200.1.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t1 then priority low
set policy-options policy-statement ospf-import term t1 then accept
```

```

set policy-options policy-statement ospf-import term t2 from route-filter 200.2.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t2 then priority medium
set policy-options policy-statement ospf-import term t2 then accept
set policy-options policy-statement ospf-import term t3 from route-filter 200.3.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t3 then priority high
set policy-options policy-statement ospf-import term t3 then accept
set protocols ospf import ospf-import
set protocols ospf area 0.0.0.0 interface fe-0/1/0
set protocols ospf area 0.0.0.0 interface fe-1/1/0

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure an OSPF import policy that prioritizes specific prefixes:

1. Configure the interfaces.

[edit]

```

user@host# set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
user@host# set interfaces fe-0/2/0 unit 0 family inet address 192.168.8.5/30

```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

[edit]

```

user@host# set protocols ospf area 0.0.0.0 interface fe-0/1/0
user@host# set protocols ospf area 0.0.0.0 interface fe-0/2/0

```

3. Configure the policy to specify the priority for prefixes learned through OSPF.

[edit]

```

user@host# set policy-options policy-statement ospf-import term t1 from route-filter
  200.1.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t1 then priority
  low
user@host# set policy-options policy-statement ospf-import term t1 then accept
user@host# set policy-options policy-statement ospf-import term t2 from route-filter
  200.2.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t2 then priority
  medium
user@host# set policy-options policy-statement ospf-import term t2 then accept
user@host# set policy-options policy-statement ospf-import term t3 from route-filter
  200.3.0.0/16 orlonger
user@host# set policy-options policy-statement ospf-import term t3 then priority
  high
user@host# set policy-options policy-statement ospf-import term t3 then accept

```

4. Apply the policy to OSPF.

[edit]

```
user@host# set protocols ospf import ospf-import
```

5. If you are done configuring the device, commit the configuration.

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

Results Confirm your configuration by entering the **show interfaces**, **show policy-options**, and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces  
fe-0/1/0 {  
  unit 0 {  
    family inet {  
      address 192.168.8.4/30;  
    }  
  }  
}  
fe-0/2/0 {  
  unit 0 {  
    family inet {  
      address 192.168.8.5/30;  
    }  
  }  
}  
  
user@host# show policy-options  
policy-statement ospf-import {  
  term t1 {  
    from {  
      route-filter 200.1.0.0/16 orlonger;  
    }  
    then {  
      priority low;  
      accept;  
    }  
  }  
  term t2 {  
    from {  
      route-filter 200.2.0.0/16 orlonger;  
    }  
    then {  
      priority medium;  
      accept;  
    }  
  }  
  term t3 {  
    from {  
      route-filter 200.3.0.0/16 orlonger;  
    }  
    then {  
      priority high;  
      accept;  
    }  
  }  
}
```

```

}

user@host# show protocols ospf
import ospf-import;
area 0.0.0.0 {
    interface fe-0/1/0.0;
    interface fe-0/2/0.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands.

Verification

Confirm that the configuration is working properly.

Verifying the Prefix Priority in the OSPF Routing Table

Purpose	Verify the priority assigned to the prefix in the OSPF routing table.
Action	From operational mode, enter the show ospf route detail for OSPFv2, and enter the show ospf3 route detail command for OSPFv3.
Related Documentation	<ul style="list-style-type: none"> • OSPF Overview on page 3798 • OSPF Configuration Overview • Configuring Match Conditions in Routing Policy Terms in the Routing Policy Configuration Guide • Configuring Actions in Routing Policy Terms in the Routing Policy Configuration Guide

Examples: Configuring Routing Policy for Network Summaries

- [Import and Export Policies for Network Summaries Overview on page 3937](#)
- [Example: Configuring an OSPF Export Policy for Network Summaries on page 3938](#)
- [Example: Configuring an OSPF Import Policy for Network Summaries on page 3946](#)

Import and Export Policies for Network Summaries Overview

By default, OSPF uses network-summary link-state advertisements (LSAs) to transmit route information across area boundaries. Each area border router (ABR) floods network-summary LSAs to other routing devices in the same area. The ABR also controls which routes from the area are used to generate network-summary LSAs into other areas. Each ABR maintains a separate topological database for each area to which they are connected. In Junos OS Release 9.1 and later, you can configure export and import policies for OSPFv2 and OSPFv3 that enable you to control how network-summary LSAs, which contain information about interarea OSPF prefixes, are distributed and generated. For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

The export policy enables you to specify which summary LSAs are flooded into an area. The import policy enables you to control which routes learned from an area are used to generate summary LSAs into other areas. You define a routing policy at the **[edit policy-options policy-statement *policy-name*]** hierarchy level. As with all OSPF export policies, the default for network-summary LSA export policies is to reject everything. Similarly, as with all OSPF import policies, the default for network-summary LSA import policies is to accept all OSPF routes.

Example: Configuring an OSPF Export Policy for Network Summaries

This example shows how to create an OSPF export policy to control the network-summary (Type 3) LSAs that the ABR floods into an OSPF area.

- [Requirements on page 3938](#)
- [Overview on page 3938](#)
- [Configuration on page 3940](#)
- [Verification on page 3945](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#)

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.



NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

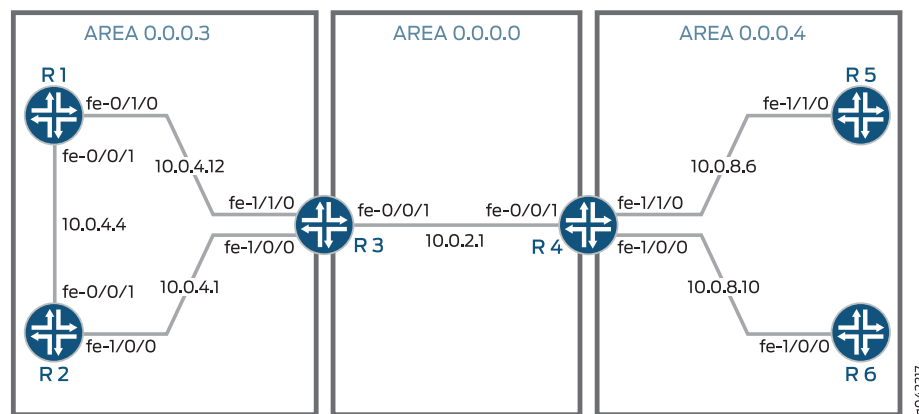
The following guidelines apply to export network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.

- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

Figure 128 on page 3939 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 128: Sample Topology Used for an OSPF Export Network Summary Policy



In this example, you configure R4 with an export network summary policy named `export-policy` that only allows routes that match the `10.0.4.4` prefix from area 3 into area 4. The export policy controls the network-summary LSAs that R4 floods into area 4. This results in only the allowed interarea route to enter area 4, and all other interarea routes to be purged from the OSPF database and the routing table of the devices in area 4. You first define the policy and then apply it to the ABR by including the `network-summary-export` statement for OSPFv2 or the `inter-area-prefix-export` statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface `fe-0/1/0` has an IP address of `10.0.4.13/30` and connects to R3. Interface `fe-0/0/1` has an IP address of `10.0.4.5/30` and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface `fe-0/0/1` has an IP address of `10.0.4.6/30` and connects to R1. Interface `fe-1/0/0` has an IP address of `10.0.4.1` and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface `fe-1/0/0` has an IP address of `10.0.4.2/30` and connects to R2. Interface `fe-1/1/0` has an IP address of `10.0.4.14/30` and connects to R1. Interface `fe-0/0/1` has an IP address of `10.0.2.1/30` and connects to R4.

- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface **fe-0/0/1** has an IP address of 10.0.2.4/30 and connects to R3. Interface **fe-1/1/0** has an IP address of 10.0.8.6/30 and connects to R5. Interface **fe-1/0/0** has an IP address of 10.0.8.9/30 and connects to R6.
- R5—Device R5 is an internal router in area 4. Interface **fe-1/1/0** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **fe-1/0/0** has an IP address of 10.0.8.10/30 and connects to R4.

Configuration

CLI Quick Configuration To quickly configure an OSPF export policy for network summaries, copy the following commands, removing any line breaks, and then paste the commands into the CLI.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set policy-options policy-statement export-policy term term1 from route-filter 10.0.4.4/30
  prefix-length-range /30-/30
set policy-options policy-statement export-policy term term1 then accept
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0
set protocols ospf area 0.0.0.4 network-summary-export export-policy
```


Configuration on Device R5:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface fe-0/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration in CLI User Guide*.

To configure an OSPF export policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
```

```
[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
```

```
[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
```

```
[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
```

```
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

```
[edit]
```

```
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
```

```
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
[edit]
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
```

```
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
[edit]
```

```
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

```
[edit]
```

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
[edit]
```

```
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

3. On R4, configure the export network summary policy.

```
[edit]
```

```
user@R4# set policy-options policy-statement export-policy term term1 from  
route-filter 10.0.4.4/30 prefix-length-range /30-/30
```

```
user@R4# set policy-options policy-statement export-policy term term1 then accept
```

4. On R4, apply the export network summary policy to OSPF.



NOTE: For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```
[edit]
```

```
user@R4# set protocols ospf area 0.0.0.4 network-summary-export export-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
```

```
user@host# commit
```

Results Confirm your configuration by entering the **show interfaces**, **show policy-options**, and **show protocols ospf** commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.4.5/30;  
    }  
  }  
}
```

```

}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}

user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}

```

Output for R2:

```

user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.1/30;
    }
  }
}

user@R2# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}

```

Output for R3:

```

user@R3# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.2/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {

```

```
        family inet {
            address 10.0.4.14/30;
        }
    }
}
```

```
user@R3# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.3 {
    interface fe-1/0/0.0;
    interface fe-1/1/0.0;
}
```

Output for R4:

```
user@R4# show interfaces
fe-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.2.1/30;
        }
    }
}
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.8.9/30;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.8.6/30;
        }
    }
}
```

```
user@R4# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.4 {
    network-summary-export export-policy;
    interface fe-1/0/0.0;
    interface fe-1/1/0.0;
}
```

```
user@R4# show policy-options
policy-statement export-policy {
    term term1 {
        from {
            route-filter 10.0.4.4/30 prefix-length-range /30-/30;
        }
        then accept;
    }
}
```

```

    }
  }

```

Output for R5:

```

user@R5# show interfaces
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.4 {
  interface fe-1/1/0.0;
}

```

Output for R6:

```

user@R6# show interfaces
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.10/30;
    }
  }
}

user@R6# show protocols ospf
area 0.0.0.4 {
  interface fe-1/0/0.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 3945](#)
- [Verifying the Routing Table on page 3946](#)

Verifying the OSPF Database

Purpose Verify that the OSPF database for the devices in area 4 includes the interarea route that we permitted on the ABR R4. The other interarea routes that are not specified should age out or no longer be present in the OSPF database.

Action From operational mode, enter the **show ospf database netsummary area 0.0.0.4** command for OSPFv2, and enter the **show ospf3 database inter-area-prefix area 0.0.0.4** command for OSPFv3.

Verifying the Routing Table

- Purpose** Verify that the routes corresponding to the rejected network summaries are no longer present in R4's, R5's, or R6's routing table.
- Action** From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

Example: Configuring an OSPF Import Policy for Network Summaries

This example shows how to create an OSPF import policy to control the network-summary (Type 3) LSAs that the ABR advertises out of an OSPF area.

- [Requirements on page 3946](#)
- [Overview on page 3946](#)
- [Configuration on page 3948](#)
- [Verification on page 3953](#)

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 3810](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 3812](#).

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.




NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to import network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

Figure 129 on page 3947 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 129: Sample Topology Used for an OSPF Import Network Summary Policy



In this example, you configure R3 with an import network summary policy named `import-policy` so R3 only generates network summaries for the route 10.0.4.12/30. The import policy controls the routes and therefore the network summaries that R3 advertises out of area 3, so applying this policy means that R3 only advertises route 10.0.4.12/30 out of area 3. This results in existing network summaries from other interarea routes getting purged from the OSPF database in area 0 and area 4, as well as the routing tables of the devices in areas 0 and area 4. You first define the policy and then apply it to the ABR by including the **network-summary-import** statement for OSPFv2 or the **inter-area-prefix-import** statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface **fe-0/1/0** has an IP address of 10.0.4.13/30 and connects to R3. Interface **fe-0/0/1** has an IP address of 10.0.4.5/30 and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface **fe-0/0/1** has an IP address of 10.0.4.6/30 and connects to R1. Interface **fe-1/0/0** has an IP address of 10.0.4.1/30 and connects to R3.

- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface **fe-1/0/0** has an IP address of 10.0.4.2/30 and connects to R2. Interface **fe-1/1/0** has an IP address of 10.0.4.14/30 and connects to R1. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R3. Interface **fe-1/1/0** has an IP address of 10.0.8.6/30 and connects to R5. Interface **fe-1/0/0** has an IP address of 10.0.8.9/30 and connects to R6.
- R5—Device R5 is an internal router in area 4. Interface **fe-1/1/0** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **fe-1/0/0** has an IP address of 10.0.8.10/30 and connects to R4.

Configuration

CLI Quick Configuration To quickly configure an OSPF import policy for network summaries, copy the following commands, removing any line breaks, and then paste the commands into CLI.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set policy-options policy-statement import-policy term term1 from route-filter 10.0.4.12/30
  prefix-length-range /30-/30
set policy-options policy-statement import-policy term term1 then accept
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```



```

set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-1/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0

```

Configuration on Device R5:

```

[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface fe-1/1/0

```

Configuration on Device R6:

```

[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration in CLI User Guide*.

To configure an OSPF import policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```

[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30

[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30

[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30

[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30

[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30

[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30

```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
[edit]
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

3. On R3, configure the import network summary policy.

```
[edit ]
user@R3# set policy-options policy-statement import-policy term term1 from
route-filter 10.0.4.12/30 prefix-length-range /30-/30
user@R3# set policy-options policy-statement import-policy term term1 then accept
```

4. On R3, apply the import network summary policy to OSPF.



NOTE: For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

5. If you are done configuring the devices, commit the configuration.

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

Results Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```

user@R1# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}

user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}

```

Output for R2:

```

user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.1/30;
    }
  }
}

user@R2# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}

```

Output for R3:

```

user@R3# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {

```

```
        family inet {
            address 10.0.4.2/30;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.4.14/30;
        }
    }
}

user@R3# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.3 {
    network-summary-import import-policy;
    interface fe-1/0/0.0;
    interface fe-1/1/0.0;
}

user@R3# show policy-options
policy-statement import-policy {
    term term1 {
        from {
            route-filter 10.0.4.12/30 prefix-length-range /30-/30;
        }
        then accept;
    }
}
```

Output for R4:

```
user@R4# show interfaces
fe-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.2.1/30;
        }
    }
}
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.8.9/30;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.8.6/30;
        }
    }
}
```

```

user@R4# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/1.0;
}
area 0.0.0.4 {
    interface fe-0/1/0.0;
    interface fe-1/0/0.0;
}

```

Output for R5:

```

user@R5# show interfaces
fe-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.8.5/30;
        }
    }
}

user@R5# show protocols ospf
area 0.0.0.4 {
    interface fe-1/1/0.0;
}

```

Output for R6:

```

user@R6# show interfaces
fe-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.8.10/30;
        }
    }
}

user@R6# show protocols ospf
area 0.0.0.4 {
    interface fe-1/0/0.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 3953](#)
- [Verifying the Routing Table on page 3954](#)

Verifying the OSPF Database

Purpose Verify that the OSPF database for the devices in area 4 includes the interarea route that we are advertising from R3. Any other routes from area 3 should not be advertised into area 4, so those entries should age out or no longer be present in the OSPF database.

Action From operational mode, enter the **show ospf database netsummary area 0.0.0.4** command for OSPFv2, and enter the **show ospf3 database inter-area-prefix area 0.0.0.4** command for OSPFv3.

Verifying the Routing Table

Purpose Verify that the specified route is included in R4's, R5's, or R6's routing table. Any other routes from area 3 should not be advertised into area 4.

Action From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

- Related Documentation**
- [OSPF Overview on page 3798](#)
 - [OSPF Configuration Overview](#)
 - [Configuring Match Conditions in Routing Policy Terms in the Routing Policy Configuration Guide](#)
 - [Configuring Actions in Routing Policy Terms in the Routing Policy Configuration Guide](#)

OSPF Monitoring Configuration

- [Example: Configuring OSPF Trace Options on page 3954](#)

Example: Configuring OSPF Trace Options

- [Tracing OSPF Protocol Traffic on page 3954](#)
- [Example: Tracing OSPF Protocol Traffic on page 3956](#)

Tracing OSPF Protocol Traffic

Tracing operations record detailed messages about the operation of OSPF. You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace.

You can specify the following OSPF protocol-specific trace options:

- **database-description**—All database description packets, which are used in synchronizing the OSPF topological database
- **error**—OSPF error packets
- **event**—OSPF state transitions
- **flooding**—Link-state flooding packets
- **graceful-restart**—Graceful-restart events
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable
- **ldp-synchronization**—Synchronization events between OSPF and LDP

- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database
- **nsr-synchronization**—Nonstop routing synchronization events
- **on-demand**—Trace demand circuit extensions
- **packet-dump**—Dump the contents of selected packet types
- **packets**—All OSPF packets
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events
- **spf**—Shortest path first (SPF) calculations

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



NOTE: Use the **detail** flag modifier with caution as it might cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the OSPF protocol using the **traceoptions flag** statement included at the **[edit protocols ospf]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



NOTE: Use the trace flag all with caution as it might cause the CPU to become very busy.

Example: Tracing OSPF Protocol Traffic

This example shows how to trace OSPF protocol traffic.

- [Requirements on page 3956](#)
- [Overview on page 3956](#)
- [Configuration on page 3957](#)
- [Verification on page 3960](#)

Requirements

This example assumes that OSPF is properly configured and running in your network, and you want to trace OSPF protocol traffic for debugging purposes.

Overview

You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace. All files are placed in a directory on the routing device's hard disk. On M Series and T Series routers, trace files are stored in the /var/log directory.

This example shows a few configurations that might be useful when debugging OSPF protocol issues. The verification output displayed is specific to each configuration.



TIP: To keep track of your log files, create a meaningful and descriptive name so it is easy to remember the content of the trace file. We recommend that you place global routing protocol tracing output in the file `routing-log`, and OSPF tracing output in the file `ospf-log`.

In the first example, you globally enable tracing operations for all routing protocols that are actively running on your routing device to the file `routing-log`. With this configuration, you keep the default settings for the trace file size and the number of trace files. After enabling global tracing operations, you enable tracing operations to provide detailed information about OSPF packets, including link-state advertisements, requests, and updates, database description packets, and hello packets to the file `ospf-log`, and you configure the following options:

- **size**—Specifies the maximum size of each trace file, in KB, MB, or GB. In this example, you configure 10 KB as the maximum size. When the file reaches its maximum size, it is renamed with a .0 extension. When the file again reaches its maximum size, it is renamed with a .1 extension, and the newly created file is renamed with a .0 extension. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option. You specify **k** for

KB, **m** for MB, and **g** for GB. By default, the trace file size is 128 KB. The file size range is 10 KB through the maximum file size supported on your system.

- **files**—Specifies the maximum number of trace files. In this example, you configure a maximum of 5 trace files. When a trace file reaches its maximum size, it is renamed with a .0 extension, then a .1 extension, and so on until the maximum number of trace files is reached. When the maximum number of files is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option. By default, there are 10 files. The range is 2 through 1000 files.

In the second example, you trace all SPF calculations to the file `ospf-log` by including the **spf** flag. You keep the default settings for the trace file size and the number of trace files.

In the third example, you trace the creation, receipt, and retransmission of all LSAs to the file `ospf-log` by including the **lsa-request**, **lsa-update**, and **lsa-ack** flags. You keep the default settings for the trace file size and the number of trace files.

Configuration

- [Configuring Global Tracing Operations and Tracing OSPF Packet Information on page 3957](#)
- [Tracing SPF Calculations on page 3959](#)
- [Tracing Link-State Advertisements on page 3959](#)

Configuring Global Tracing Operations and Tracing OSPF Packet Information

CLI Quick Configuration

To quickly enable global tracing operations for all routing protocols actively running on your routing device and to trace detailed information about OSPF packets, copy the following commands and paste them into the CLI.

```
[edit]
set routing-options traceoptions file routing-log
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions file files 5 size 10k
set protocols ospf traceoptions flag lsa-ack
set protocols ospf traceoptions flag database-description
set protocols ospf traceoptions flag hello
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-request
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure global routing tracing operations and tracing operations for OSPF packets:

1. Configure tracing at the routing options level to collect information about the active routing protocols on your routing device.

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the filename for the global trace file.

```
[edit routing-options traceoptions]
user@host# set file routing-log
```

3. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

4. Configure the maximum number of trace files.

```
[edit protocols ospf traceoptions]
user@host# set file files 5
```

5. Configure the maximum size of each trace file.

```
[edit protocols ospf traceoptions]
user@host# set file size 10k
```

6. Configure tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-ack
user@host# set flag database-description
user@host# set flag hello
user@host# set flag lsa-update
user@host# set flag lsa-request
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results Confirm your configuration by entering the `show routing-options` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
traceoptions {
  file routing-log;
}

user@host# show protocols ospf
traceoptions {
  file ospf-log size 10k files 5;
  flag lsa-ack;
  flag database-description;
  flag hello;
  flag lsa-update;
  flag lsa-request;
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options** and the **show protocols ospf3** commands.

Tracing SPF Calculations

CLI Quick Configuration To quickly trace SPF calculations, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag spf
```

Step-by-Step Procedure To configure SPF tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the SPF tracing flag.

```
[edit protocols ospf traceoptions]
user@host# set flag spf
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log ;
  flag spf;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Tracing Link-State Advertisements

CLI Quick Configuration To quickly trace the creation, receipt, and retransmission of all LSAs, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag lsa-request
```

```
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-ack
```

Step-by-Step Procedure To configure link-state advertisement tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the link-state advertisement tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-request
user@host# set flag lsa-update
user@host# set flag lsa-ack
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log;
  flag lsa-request;
  flag lsa-update;
  flag lsa-ack;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

Confirm that the configuration is working properly.

Verifying Trace Operations

Purpose Verify that the Trace options field displays the configured trace operations, and verify that the Trace file field displays the location on the routing device where the file is saved, the name of the file to receive the output of the tracing operation, and the size of the file.

Action From operational mode, enter the `show ospf overview extensive` command for OSPFv2, and enter the `show ospf3 overview extensive` command for OSPFv3.

- Related Documentation**
- [OSPF Overview on page 3798](#)
 - [OSPF Configuration Overview](#)
 - [Junos OS Tracing and Logging Operations in the Junos OS System Basics Configuration Guide](#)
 - [Example: Tracing Global Routing Protocol Operations in the Junos OS Routing Protocols Configuration Guide](#)

Configuration Statements

- [area on page 3963](#)
- [area-range on page 3965](#)
- [authentication \(Protocols OSPF\) on page 3967](#)
- [backup-spf-options on page 3968](#)
- [bandwidth-based-metrics on page 3969](#)
- [bfd-liveness-detection \(Protocols OSPF\) on page 3971](#)
- [context-identifier \(Protocols OSPF\) on page 3974](#)
- [database-protection on page 3975](#)
- [dead-interval on page 3977](#)
- [default-lsa on page 3978](#)
- [disable \(OSPF\) on page 3979](#)
- [export \(Protocols OSPF\) on page 3981](#)
- [external-preference \(Protocols OSPF\) on page 3982](#)
- [flood-reduction on page 3983](#)
- [graceful-restart \(Protocols OSPF\) on page 3984](#)
- [hello-interval \(Protocols OSPF\) on page 3986](#)
- [helper-disable \(Multiple Protocols\) on page 3987](#)
- [ignore-lsp-metrics on page 3988](#)
- [import \(Protocols OSPF\) on page 3989](#)
- [inter-area-prefix-export on page 3990](#)
- [inter-area-prefix-import on page 3991](#)
- [interface \(Protocols OSPF\) on page 3992](#)
- [interface-type \(Protocols OSPF\) on page 3995](#)
- [lsa-refresh-interval on page 3996](#)
- [metric \(Protocols OSPF Interface\) on page 3997](#)
- [no-eligible-backup \(Protocols OSPF\) on page 3999](#)
- [no-nssa-abr on page 4000](#)
- [no-rfc-1583 on page 4001](#)

- [no-strict-lsa-checking](#) on page 4002
- [node-link-protection](#) (Protocols OSPF) on page 4003
- [notify-duration](#) on page 4004
- [nssa](#) on page 4005
- [ospf](#) on page 4006
- [ospf3](#) on page 4006
- [overload](#) (Protocols OSPF) on page 4007
- [passive](#) (Protocols OSPF) on page 4009
- [preference](#) (Protocols OSPF) on page 4010
- [prefix-export-limit](#) (Protocols OSPF) on page 4011
- [priority](#) (Protocols OSPF) on page 4012
- [realm](#) on page 4013
- [reference-bandwidth](#) (Protocols OSPF) on page 4014
- [retransmit-interval](#) (OSPF) on page 4015
- [rib-group](#) (Protocols OSPF) on page 4016
- [shortcuts](#) (Protocols OSPF) on page 4017
- [spf-options](#) (Protocols OSPF) on page 4018
- [stub](#) on page 4020
- [summaries](#) on page 4021
- [topology](#) (OSPF) on page 4022
- [traceoptions](#) (Protocols OSPF) on page 4023
- [traffic-engineering](#) (OSPF) on page 4026
- [transmit-interval](#) (Protocols OSPF) on page 4028
- [transit-delay](#) (OSPF) on page 4029

area

Syntax	<pre> area <i>area-id</i> { interface <i>interface-name</i> { passive; topology (ipv4-multicast <i>name</i>) { disable; } } virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> { topology (ipv4-multicast <i>name</i>) { disable; } } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.</p> <p>Specify multiple area statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the area-range statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the virtual-link statement.</p> <p>To specify that the routing device is directly connected to the OSPF backbone, include the area 0.0.0.0 statement.</p> <p>All routing devices on the backbone must be contiguous. If they are not, use the virtual-link statement to create the appearance of connectivity to the backbone.</p>

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



NOTE: If you configure an interface with the **passive** statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

Options	area-id —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number 0.0.0.0 is reserved for the OSPF backbone area.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview on page 3803• <i>Understanding Multiple Address Families for OSPFv3</i>• <i>virtual-link</i>

area-range

Syntax	area-range <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple area-range statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple area-range statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>
Default	By default, area border routers do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.
Options	<p>exact—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p>mask-length—Number of significant bits in the network mask.</p> <p>network—IP address. You can specify one or more IP addresses.</p>

override-metric *metric*—(Optional) Override the metric for the IP address range and configure a specific metric value.

restrict—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

Range: 1 through 16,777,215

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements on page 3852
------------------------------	--

authentication (Protocols OSPF)

Syntax	<pre> authentication { md5 key-identifier { key key-value; start-time YYYY-MM-DD.hh:mm; } simple-password key; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> interface interface-name],</p> <p>[edit protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface interface-name],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure an authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding OSPFv2 Authentication</i> • <i>Example: Configuring MD5 Authentication for OSPFv2 Exchanges</i> • <i>Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface</i> • <i>Example: Configuring Simple Authentication for OSPFv2 Exchanges</i>

backup-spf-options

Syntax	<pre> backup-spf options { disable; downstream-paths-only; no-install; } </pre>
Hierarchy Level	<pre> [edit protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit protocols ospf topology (default <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf topology (default <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default <i>name</i>)]; [edit protocols ospf3 realm ipv4-unicast], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm ipv4-unicast], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast] </pre>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm for backup next hops for protected OSPF interfaces. Use these options to override the default behavior of having Junos OS calculate backup paths for all the topologies in an OSPF instance when at least one OSPF interface is configured with link protection or node-link protection. These options also enable you to change the default behavior for a specific topology in an OSPF instance.</p>
Options	<p>disable—Do not calculate backup next hops for the specified OSPF instance or topology.</p> <p>downstream-paths-only—Calculate and install only downstream paths as defined in RFC 5286, <i>Basic Specification for IP Fast Reroute: Loop-Free Alternates</i> for the specified OSPF instance or topology.</p> <p>no-install—Do not install the backup next hops for the specified OSPF instance or topology.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Backup SPF Options for Protected OSPF Interfaces</i> • <i>link-protection</i> • node-link-protection on page 4003

bandwidth-based-metrics

Syntax	<pre>bandwidth-based-metrics { bandwidth <i>value</i>; metric <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.</p>
Options	<p>bandwidth <i>value</i>—Specify the bandwidth threshold in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000,000</p> <p>metric <i>number</i>—Specify a metric value to associate with a specific bandwidth value.</p> <p>Range: 1 through 65,535</p>



NOTE: You must also configure a static metric value for the OSPF interface or topology with the `metric` statement. Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for

the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

.....

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 3867• metric on page 3997• Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 3867 |
|------------------------------|---|

bfd-liveness-detection (Protocols OSPF)

Syntax `bfd-liveness-detection {`
 `authentication {`
 `algorithm algorithm-name;`
 `key-chain key-chain-name;`
 `loose-check;`
 `}`
 `detection-time {`
 `threshold milliseconds;`
 `}`
 `full-neighbors-only`
 `minimum-interval milliseconds;`
 `minimum-receive-interval milliseconds;`
 `multiplier number;`
 `no-adaptation;`
 `transmit-interval {`
 `minimum-interval milliseconds;`
 `threshold milliseconds;`
 `}`
 `version (1 | automatic);`
 `}`

Hierarchy Level `[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast |`
 `ipv4-multicast | ipv6-multicast) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `(ospf | ospf3) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface`
 `interface-name],`
 `[edit protocols (ospf | ospf3) area area-id interface interface-name],`
 `[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id`
 `interface interface-name],`
 `[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface`
 `interface-name],`
 `[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |`
 `ipv4-multicast | ipv6-multicast) area area-id interface interface-name]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 detection-time threshold and **transmit-interval threshold** options added in Junos OS Release 8.2.
 Support for logical systems introduced in Junos OS Release 8.3.
 no-adaptation option introduced in Junos OS Release 9.0.
 no-adaptation option introduced in Junos OS Release 9.0 for EX Series switches.
 Support for OSPFv3 introduced in Junos OS Release 9.3.
 Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.
 full-neighbors-only option introduced in Junos OS Release 9.5.
 full-neighbors-only option introduced in Junos OS Release 9.5 for EX Series switches.

authentication algorithm, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately.

Options **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

full-neighbors-only—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval** *minimum-interval* and **minimum-receive-interval** statements.

Range: 1 through 255,000 milliseconds

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000 milliseconds

multiplier *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: **automatic**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring BFD for OSPF on page 3883](#)
- [Example: Configuring BFD Authentication for OSPF on page 3887](#)

context-identifier (Protocols OSPF)

Syntax context-identifier *identifier*

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],
[edit protocols (ospf | ospf3) *area area-id*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure OSPF context-identifier information.

Options *identifer*—IPv4 address that defines a protection pair. The context identifier is manually configured on both the primary and protector provider edge (PE) devices.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [show ospf context-identifier on page 4048](#)

database-protection

Syntax	<pre>database-protection { ignore-count <i>number</i>; ignore-time <i>seconds</i>; maximum-lsa <i>number</i>; reset-time <i>seconds</i>; warning-only; warning-threshold <i>percent</i>; }</pre>
Hierarchy Level	<pre>[edit protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-unicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the maximum number of link-state advertisements (LSAs) that are not generated by the router or switch in a given OSPF instance.
Default	By default, OSPF database protection is not enabled.
Options	<p>ignore-count <i>number</i>—Configure the number of times the database can enter the ignore state. When the ignore count is exceeded, the database enters the isolate state.</p> <p>Range: 1 through 32</p> <p>Default: 5</p> <p>ignore-time <i>seconds</i>—Configure the time the database must remain in the ignore state before it resumes regular operations (enters retry state).</p> <p>Range: 30 through 3,600 seconds</p> <p>Default: 300 seconds</p> <p>maximum-lsa <i>number</i>—Configure the maximum number of LSAs whose advertising router ID is different from the local router ID in a given OSPF instance. This includes external LSAs as well as LSAs with any scope, such as the link, area, and autonomous system (AS). This value is mandatory.</p> <p>Range: 1 through 1,000,000</p> <p>Default: None</p> <p>reset-time <i>seconds</i>—Configure the time period during which the database must operate without being in the ignore or isolate state before it is reset to a normal operating state.</p> <p>Range: 60 through 86,400 seconds</p> <p>Default: 600 seconds</p>

warning-only—Specify that only a warning should be issued when the maximum LSA number is exceeded. If configured, no other action is taken against the database.

warning-threshold *percent*—Configure the percentage of the maximum number of LSAs to be exceeded before a warning message is logged.

Range: 30 through 100 percent

Default: 75 percent

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• OSPF Database Protection Overview on page 3919• Configuring OSPF Database Protection on page 3920
------------------------------	--

dead-interval

Syntax	<code>dead-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF Timers on page 3876 • Configuring RSVP and OSPF for LMP Peer Interfaces

- [hello-interval on page 3986](#)

default-lsa

Syntax	<pre>default-lsa { default-metric <i>metric</i>; metric-type <i>type</i>; type-7; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>On area border routers only, for a not-so-stubby area (NSSA), inject a default link-state advertisement (LSA) with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview on page 3803 • Example: Configuring OSPF Not-So-Stubby Areas on page 3826 • nssa on page 4005 • stub on page 4020

disable (OSPF)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols ospf <i>area</i> <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Disable OSPF, an OSPF interface, or an OSPF virtual link.</p> <p>By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that</p>

are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id *router-id* transit-area *area-id* topology *name*]** hierarchy level.



NOTE: If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id *router-id* transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>OSPF Configuration Overview</i>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i>

export (Protocols OSPF)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply one or more policies to routes being exported from the routing table into OSPF.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding OSPF Routing Policy on page 3921 • Import and Export Policies for Network Summaries Overview on page 3937 • import on page 3989 • <i>Routing Policy Configuration Guide</i>

external-preference (Protocols OSPF)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Set the route preference for OSPF external routes.
Options	<i>preference</i> —Preference value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 150
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Controlling OSPF Route Preferences on page 3869• preference on page 4010

flood-reduction

Syntax	flood-reduction;
Hierarchy Level	<p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interfaces <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interfaces <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interfaces <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interfaces <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interfaces <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>transit-area</i> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote <i>address</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Specify to send self-generated link-state advertisements (LSAs) with the DoNotAge bit set. As a result, self-originated LSAs are not reflooded every 30 minutes, as required by OSPF by default. An LSA is refreshed only when the content of the LSA changes, which reduces OSPF traffic overhead in stable topologies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 3860

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Graceful Restart for OSPF on page 3894 • Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 3898 • Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 3901 • Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 3904 • Configuring Graceful Restart for QFabric Systems on page 1375 • <i>Junos OS High Availability Configuration Guide</i>
------------------------------	--

hello-interval (Protocols OSPF)

Syntax	<code>hello-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.
Options	<p>seconds—Time between hello packets, in seconds.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF Timers on page 3876 • Configuring RSVP and OSPF for LMP Peer Interfaces

- [dead-interval on page 3977](#)

helper-disable (Multiple Protocols)

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart], [edit protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart.
Default	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Protocols Graceful Restart on page 2285 • Configuring Graceful Restart for MPLS-Related Protocols

ignore-lsp-metrics

Syntax	ignore-lsp-metrics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for (OSPFv3) introduced in Junos OS Release 9.4. Support for (OSPFv3) introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling OSPF Traffic Engineering Support on page 3910

import (Protocols OSPF)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Filter OSPF routes from being added to the routing table.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding OSPF Routing Policy on page 3921 • Import and Export Policies for Network Summaries Overview on page 3937 • export on page 3981 • <i>Routing Policy Configuration Guide</i>

inter-area-prefix-export

Syntax	<code>inter-area-prefix-export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>area area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area area-id</i>],</code> <code>[edit protocols ospf3 <i>area area-id</i>],</code> <code>[edit protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area area-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.
Options	<i>policy-name</i> —Name of a policy configured at the <code>[edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>]</code> hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Import and Export Policies for Network Summaries Overview on page 3937• inter-area-prefix-import on page 3991• <i>Routing Policy Configuration Guide</i>

inter-area-prefix-import

Syntax	<code>inter-area-prefix-import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit protocols ospf3 area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ip4-unicast ipv4-multicast ipv6-multicast)], area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.
Options	<i>policy-name</i> —Name of a policy configured at the [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>] hierarchy level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Import and Export Policies for Network Summaries Overview on page 3937 • inter-area-prefix-export on page 3990 • Routing Policy Configuration Guide

interface (Protocols OSPF)

Syntax interface *interface-name* {
 disable;
 authentication key <key-id identifier>;
 bfd-liveness-detection {
 authentication {
 algorithm *algorithm-name*;
 key-chain *key-chain-name*;
 loose-check;
 }
 detection-time {
 threshold *milliseconds*;
 }
 minimum-interval *milliseconds*;
 minimum-receive-interval *milliseconds*;
 transmit-interval {
 threshold *milliseconds*;
 minimum-interval *milliseconds*;
 }
 multiplier *number*;
 }
 dead-interval *seconds*;
 demand-circuit;
 hello-interval *seconds*;
 ipsec-sa *name*;
 interface-type *type*;
 ldp-synchronization {
 disable;
 hold-time *seconds*;
 }
 metric *metric*;
 neighbor *address* <eligible>;
 no-interface-state-traps;
 passive;
 poll-interval *seconds*;
 priority *number*;
 retransmit-interval *seconds*;
 te-metric *metric*;
 topology (ipv4-multicast | *name*) {
 metric *metric*;
 }
 transit-delay *seconds*;
 transmit-interval *seconds*;
}

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* protocols ospf3 *realm* (ipv4-unicast |
 ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 ospf3 *realm* (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit protocols (ospf | ospf3) *area area-id*],

```
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |
  ipv4-multicast | ipv6-multicast) area area-id]
```

Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the topology statement introduced in Junos OS Release 9.0.</p> <p>Support for the topology statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the no-interface-state-traps statement introduced in Junos OS Release 10.3.</p> <p>This statement is supported only for OSPFv2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Enable OSPF routing on a routing device interface.</p> <p>You must include at least one interface statement in the configuration to enable OSPF on the routing device.</p>
Options	<p>interface-name—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify all. Specifying a particular interface and all produces an invalid configuration.</p>
<p>Required Privilege Level</p> <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>	



NOTE: For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

The remaining statements are explained separately.



NOTE: You cannot run both OSPF and ethernet-tcc encapsulation between two Juniper Networks routing devices.

- Related Documentation**
- *OSPF Configuration Overview*
 - *Example: Configuring Multitopology Routing Based on Applications*
 - *Example: Configuring Multitopology Routing Based on a Multicast Source*
 - *Example: Configuring Multiple Address Families for OSPFv3*
 - *neighbor*

interface-type (Protocols OSPF)

Syntax	<code>interface-type (nbma p2mp p2p);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4. You cannot configure other interface types for OSPFv3.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Specify the type of interface.</p> <p>By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in Nonbroadcast multiaccess (NBMA) mode, configure the nbma interface type, using the IP address of the local ATM interface.</p> <p>In Junos OS Release 9.3 and later, a point-to-point interface can be an Ethernet interface without a subnet.</p>
Default	The software chooses the correct interface type based on the type of physical interface.
Options	<p>nbma (OSPFv2 only)—Nonbroadcast multiaccess (NBMA) interface.</p> <p>p2mp (OSPFv2 only)—Point-to-multipoint interface.</p> <p>p2p—Point-to-point interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [About OSPF Interfaces on page 3838](#)
 - [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network on page 3841](#)

lsa-refresh-interval

Syntax	lsa-refresh-interval <i>minutes</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.2.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Configure the refresh interval for all self-generated link-state advertisement (LSAs). The OSPF standard requires that every LSA be refreshed every 30 minutes. The Juniper Networks implementation refreshes LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. By using this configuration, you can specify when self-originated LSAs are refreshed.</p> <p>You can override the default behavior by globally configuring the OSPF LSA refresh interval at the [edit protocols ospf ospf3] hierarchy level. However, if you also have OSPF flood reduction configured for a specific interface in an OSPF area at the [edit protocols ospf ospf3 area <i>area-id</i> interface <i>interface-name</i>] hierarchy level, the flood reduction configuration takes precedence for that specific interface.</p>
Options	<p>minutes—Time between an LSA refresh, in minutes.</p> <p>Range: 25 through 50 minutes (1,500 through 3,000 seconds)</p> <p>Default: 50 minutes</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 3860

metric (Protocols OSPF Interface)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation.</p> <p>To set the cost of routes exported into OSPF, configure the appropriate routing policy.</p>
Options	<p><i>metric</i>—Cost of the route.</p> <p>Range: 1 through 65,535</p> <p>Default: By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific</p>

value you configure for the **metric** overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Example: Controlling the Cost of Individual OSPF Network Segments on page 3863• <i>Example: Configuring OSPFv2 Sham Links</i>• <i>Example: Configuring Multitopology Routing Based on Applications</i>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>• bandwidth-based-metrics on page 3969• reference-bandwidth on page 4014
------------------------------	--

no-eligible-backup (Protocols OSPF)

Syntax	no-eligible-backup;
Hierarchy Level	<p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (default <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (default <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (default <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (default <i>name</i>)],</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Exclude the specified interface as a backup interface for OSPF interfaces on which link protection or node-link protection is enabled.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Excluding an OSPF Interface as a Backup for a Protected Interface</i> • <i>link-protection</i> • node-link-protection on page 4003

no-nssa-abr

Syntax	no-nssa-abr;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring OSPF Not-So-Stubby Areas on page 3826


no-rfc-1583

Syntax	no-rfc-1583;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.
Default	Compatibility with RFC 1583 is enabled by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 3836

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router.
Default	By default, LSA checking is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• <i>maximum-neighbor-recovery-time</i>• <i>recovery-time</i>

node-link-protection (Protocols OSPF)

Syntax	node-link-protection;
Hierarchy Level	<p>[edit protocols (ospf ospf3) protocols area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm ipv4-unicast area <i>area-id</i> interface <i>interface-name</i>],</p>
Release Information	<p>Statement introduced in Junos OS Release 10.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Enable node-link protection on the specified OSPF interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop router altogether and establishes a path through a different router.</p>
<div>  <p>NOTE: This feature is not supported for the OSPF IPv4 multicast topology or for the OSPFv3 IPv4 multicast or IPv6 multicast topologies because node-link protection creates alternate next-hop paths only for unicast routes.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Node-Link Protection for OSPF link-protection

notify-duration

Syntax	notify-duration <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart], [edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the length of time the router notifies helper OSPF routers that it has completed graceful restart.
Options	<i>seconds</i> —Length of time in the router notifies helper OSPF routers that it has completed graceful restart. Range: 1 through 3600 Default: 30
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2285• Configuring Graceful Restart for QFabric Systems on page 1375• restart-duration on page 2316

nssa

Syntax	<pre> nssa { area-range <i>network/mask-length</i> <restrict> <exact> <override-metric <i>metric</i>>; default-lsa { default-metric <i>metric</i>; metric-type <i>type</i>; type-7; } (no-summaries summaries); } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3) area <i>area-id</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.</p> <p>You cannot configure an area as being both a stub area and an NSSA.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview on page 3803 • Example: Configuring OSPF Not-So-Stubby Areas on page 3826 • stub on page 4020

ospf

Syntax	<code>ospf { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPF routing on the routing device. You must include the ospf statement to enable OSPF on the routing device.
Default	OSPF is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>OSPF Configuration Overview</i>• <i>[edit protocols ospf] Hierarchy Level</i>

ospf3

Syntax	<code>ospf3 { ... }</code>
Hierarchy Level	[edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPFv3 routing on the routing device. You must include the ospf3 statement to enable OSPFv3.
Default	OSPFv3 is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>OSPF Configuration Overview</i>• <i>[edit protocols ospf3] Hierarchy Level</i>

overload (Protocols OSPF)

Syntax	<pre>overload { timeout <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf <i>topology</i> (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <i>topology</i> (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (<i>ospf</i> <i>ospf3</i>)], [edit protocols ospf <i>topology</i> (default ipv4-multicast <i>name</i>)], [edit protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf <i>topology</i> (default ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)],</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>



NOTE: Traffic destined to directly attached interfaces continues to reach the routing device.

Options	<p>timeout <i>seconds</i>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the overload statement is deleted or a timeout is set.</p> <p>Range: 60 through 1800 seconds</p> <p>Default: 0 seconds</p>
----------------	--



NOTE: Multitopology Routing does not support the timeout option.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 3872• <i>Example: Configuring Multitopology Routing Based on Applications</i>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i> |
|------------------------------|--|

passive (Protocols OSPF)

Syntax	<pre> passive { traffic-engineering { remote-node-id address; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.</p> <p>To configure an interface in OSPF passive traffic engineering mode, include the traffic-engineering statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.</p> <p>Enable OSPF on an interface by including the interface statement at the [edit protocols (ospf ospf3) area <i>area-id</i>] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>] hierarchy levels. Disable it by including the disable statement. To prevent OSPF from running on an interface, include the passive statement. These three states are mutually exclusive.</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring a Passive OSPF Interface on page 3848• Example: Configuring OSPF Passive Traffic Engineering Mode on page 3917• disable on page 3979

preference (Protocols OSPF)

Syntax	preference <i>preference</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Set the route preference for OSPF internal routes.
Options	preference —Preference value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 10
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Controlling OSPF Route Preferences on page 3869• external-preference on page 3982

prefix-export-limit (Protocols OSPF)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a limit to the number of prefixes exported into OSPF.
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Limiting the Number of Prefixes Exported to OSPF on page 3858 • Example: Configuring Multitopology Routing Based on Applications • Example: Configuring Multitopology Routing Based on a Multicast Source


priority (Protocols OSPF)

Syntax	<code>priority</code> <i>number</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Specify the routing device's priority for becoming the designated routing device. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.
Options	<p>number—Routing device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router.</p> <p>Range: 0 through 255</p> <p>Default: 128</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Designated Router Overview on page 3809 • Example: Controlling OSPF Designated Router Election on page 3812

realm

Syntax	<pre> realm (ipv4-unicast ipv4-multicast ipv6-unicast) { area area-id { interface interface-name; } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols ospf3], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3], [edit protocols ospf3], [edit routing-instances <i>routing-instance-name</i> protocols ospf3] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Configure OSPFv3 to advertise address families other than unicast IPv6. Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.</p>
Options	<p>ipv4-unicast—Configure a realm for IPv4 unicast routes.</p> <p>ipv4-multicast—Configure a realm for IPv4 multicast routes.</p> <p>ipv6-multicast—Configure a realm for IPv6 multicast routes.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring Multiple Address Families for OSPFv3</i>

reference-bandwidth (Protocols OSPF)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$
Options	<p><i>reference-bandwidth</i>—Reference bandwidth, in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bits</p> <p>Default: 100 Mbps (100,000,000 bits)</p>
<div>  <p>NOTE: The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Controlling the Cost of Individual OSPF Network Segments on page 3863 • metric on page 3997

retransmit-interval (OSPF)

Syntax	<code>retransmit-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 5 seconds</p>



NOTE: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring OSPF Timers on page 3876](#)
- [Configuring RSVP and OSPF for LMP Peer Interfaces](#)

rib-group (Protocols OSPF)

Syntax `rib-group group-name;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols ([ospf](#) | ospf3)],
[edit logical-systems *logical-system-name* protocols ospf3 [realm](#) (ipv4-unicast |
ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
([ospf](#) | ospf3)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
ospf3 [realm](#) (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols ([ospf](#) | ospf3)],
[edit protocols ospf3 [realm](#) (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols ([ospf](#) | ospf3)],
[edit routing-instances *routing-instance-name* protocols ospf3 [realm](#) (ipv4-unicast |
ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for the **realm** statement introduced in Junos OS Release 9.2.
Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.

Options *group-name*—Name of the routing table group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Exporting Specific Routes from One Routing Table Into Another Routing Table](#)
- [Example: Importing Direct and Static Routes Into a Routing Instance](#)
- [Understanding Multiprotocol BGP](#)
- [interface-routes on page 2898](#)
- [rib-group on page 2935](#)

shortcuts (Protocols OSPF)

Syntax	shortcuts { lsp-metric-into-summary; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering], [edit protocols (ospf ospf3) traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the inet.3 routing table, and shortcut routes calculated through OSPFv3 are installed in the inet6.3 routing table.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Enabling OSPF Traffic Engineering Support on page 3910

spf-options (Protocols OSPF)

Syntax	<pre> spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm. You can configure the following:</p> <ul style="list-style-type: none"> • A delay for when to run the SPF algorithm after a network topology change is detected. • The maximum number of times the SPF algorithm can run in succession. • A hold-down interval after the SPF algorithm runs the maximum number of times. <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p>

Options	delay <i>milliseconds</i> —Time interval between the detection of a topology change and when the SPF algorithm runs. Range: 50 through 8000 milliseconds Default: 200 milliseconds
	holddown <i>milliseconds</i> —Time interval to hold down, or to wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession. Range: 2000 through 20,000 milliseconds Default: 5000 milliseconds
	rapid-runs <i>number</i> —Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold down interval begins. Range: 1 through 10 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring SPF Algorithm Options for OSPF</i>• <i>Example: Configuring Multitopology Routing Based on Applications</i>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>

stub

Syntax	<code>stub <default-metric <i>metric</i>> <(no-summaries summaries)>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (ospf ospf3) area <i>area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3) area <i>area-id</i>],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Specify that this area not be flooded with AS external link-state advertisements (LSAs). You must include the stub statement when configuring all routing devices that are in the stub area. The backbone cannot be configured as a stub area. You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).
Options	no-summaries —(Optional) Do not advertise routes into the stub area. If you include the default-metric option, only the default route is advertised. summaries —(Optional) Flood summary LSAs into the stub area. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview on page 3803• Example: Configuring OSPF Stub and Totally Stubby Areas on page 3822• nssa on page 4005

summaries

Syntax	(summaries no-summaries);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for QFX Series.</p>
Description	<p>Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):</p> <ul style="list-style-type: none"> • summaries—Flood summary link-state advertisements (LSAs) into the NSSA. • no-summaries—Prevent area border routers from advertising summaries into an NSSA. If default-metric is configured for an NSSA, a Type 3 LSA is injected into the area by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview on page 3803 • Example: Configuring OSPF Not-So-Stubby Areas on page 3826 • nssa on page 4005 • stub on page 4020

topology (OSPF)

Syntax	<pre>topology (default ipv4-multicast <i>name</i>) { <i>spf-options</i> { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; } topology-id <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf], [edit protocols ospf], [edit routing-instances <i>routing-instance-name</i> protocols ospf]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable a topology for OSPF multitenancy routing. You must first configure one or more topologies under the [edit routing-options] hierarchy level.
Options	<p>default—Name of the default topology. This topology is automatically created, and all routes that correspond to it are automatically added to the inet.0 routing table. You can modify certain default parameters, such as for the SPF algorithm.</p> <p>ipv4-multicast—Name of the topology for IPv4 multicast traffic.</p> <p>name—Name of a topology you configured at the [edit routing-options] hierarchy level to create a topology for a specific type of traffic, such as voice or video.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multitenancy Routing Based on Applications</i>• <i>Example: Configuring Multitenancy Routing Based on a Multicast Source</i>

traceoptions (Protocols OSPF)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default OSPF protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place OSPF tracing output in the file ospf-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p>

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Tracing OSPF Protocol Traffic on page 3956

traffic-engineering (OSPF)

Syntax	<pre>traffic-engineering { <advertise-unnumbered-interfaces>; <credibility-protocol-preference>; ignore-lsp-metrics; multicast-rpf-routes; no-topology; shortcuts { lsp-metric-into-summary; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit protocols (ospf ospf3)]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>multicast-rpf-routes option introduced in Junos OS Release 7.5.</p> <p>advertise-unnumbered-interfaces option introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable the OSPF traffic engineering features.
Default	Traffic engineering support is disabled.
Options	<p>advertise-unnumbered-interfaces—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p>credibility-protocol-preference—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.</p>

multicast-rpf-routes—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



NOTE: You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

no-topology—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.



CAUTION: When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation [• Example: Enabling OSPF Traffic Engineering Support on page 3910](#)

transmit-interval (Protocols OSPF)

Syntax	<code>transmit-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</code> <code>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Set the interval at which OSPF packets are transmitted on an interface.
Options	<i>milliseconds</i> —Transmission interval, in milliseconds. Range: 1 through 4,294,967 milliseconds Default: 30 milliseconds
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring OSPF Timers on page 3876

transit-delay (OSPF)

Syntax	<code>transit-delay seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<p>Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.</p> <p>You should never have to modify the transit delay time.</p>
Options	<p>seconds—Estimated time, in seconds.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 1 second</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF Timers on page 3876 • Configuring RSVP and OSPF for LMP Peer Interfaces

CHAPTER 44

Administration

- [Routine Monitoring on page 4031](#)
- [Operational Commands on page 4031](#)

Routine Monitoring

- [Monitoring OSPF Routing Information on page 4031](#)

Monitoring OSPF Routing Information

Purpose Use the monitoring functionality to monitor OSPF routing information on routing devices.

Action To view OSPF routing information in the CLI, enter the following CLI commands:

- `show ospf neighbor`
- `show ospf interface`
- `show ospf statistics`

Related Documentation

- [show \(ospf | ospf3\) interface on page 4069](#)
- [clear \(ospf | ospf3\) neighbor on page 4038](#)
- [show \(ospf | ospf3\) statistics on page 4097](#)

Operational Commands

- `clear (ospf | ospf3) database`
- `clear (ospf | ospf3) database-protection`
- `clear (ospf | ospf3) io-statistics`
- `clear (ospf | ospf3) neighbor`
- `clear (ospf | ospf3) overload`
- `clear (ospf | ospf3) statistics`
- `show (ospf | ospf3) backup coverage`
- `show (ospf | ospf3) backup neighbor`

- `show ospf context-identifier`
- `show ospf database`
- `show ospf3 database`
- `show (ospf | ospf3) interface`
- `show (ospf | ospf3) io-statistics`
- `show (ospf | ospf3) log`
- `show (ospf | ospf3) neighbor`
- `show (ospf | ospf3) overview`
- `show (ospf | ospf3) route`
- `show (ospf | ospf3) statistics`

clear (ospf | ospf3) database

List of Syntax [Syntax on page 4033](#)
 [Syntax \(EX Series Switch and QFX Series\) on page 4033](#)

Syntax clear (ospf | ospf3) database
 <advertising-router (*router-id* | self)>
 <area *area-id*>
 <asbrsummary>
 <external>
 <instance *instance-name*>
 <inter-area-prefix>
 <inter-area-router>
 <intra-area-prefix>
 <link-local>
 <logical-system (all | *logical-system-name*)>
 <lsa-id *lsa-id*>
 <netsummary>
 <network>
 <nssa>
 <opaque-area>
 <purge>
 <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
 <router>

Syntax (EX Series Switch and QFX Series) clear (ospf | ospf3) database
 <advertising-router (*router-id* | self)>
 <area *area-id*>
 <asbrsummary>
 <external>
 <instance *instance-name*>
 <inter-area-prefix>
 <inter-area-router>
 <intra-area-prefix>
 <link-local>
 <lsa-id *lsa-id*>
 <netsummary>
 <network>
 <nssa>
 <opaque-area>
 <purge>
 <router>

Release Information Command introduced before Junos OS Release 7.4.
 advertising-router *router-id*, **area** *area-id*, **asbrsummary**, **external**, **inter-area-prefix**, **inter-area-router**, **intra-area-prefix**, **link-local**, **lsa-id** *lsa-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 realm option added in Junos OS Release 9.2.
 advertising-router (*router-id* | **self**) option added in Junos OS Release 9.5.
 advertising-router (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.



CAUTION: This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

advertising-router (*router-id* | **self**)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

area *area-id*—(Optional) Discard entries for the LSAs in the specified area.

asbrsummary—(Optional) Discard summary AS boundary router LSA entries.

external—(Optional) Discard external LSAs.

instance *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

inter-area-prefix—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPFv3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

link-local—(Optional) Delete link-local LSAs.

lsa-id *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

netsummary—(Optional) Discard summary network LSAs.

network—(Optional) Discard network LSAs.

nssa—(Optional) Discard not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Discard opaque area-scope LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level

clear

Related Documentation

- [show ospf database on page 4050](#)
- [show ospf3 database on page 4058](#)

List of Sample Output [clear ospf database on page 4035](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear ospf database](#)

```
user@host> clear ospf database
```

clear (ospf | ospf3) database-protection

Syntax	clear (ospf ospf3) database-protection <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Open Shortest Path First (OSPF) link-state database from its isolated state. Reset the ignore count, ignore timer, and reset timer, and resume normal operations.
Options	instance <i>instance-name</i> —(Optional) Clear the OSPF link-state database for the specified routing instance only.
Required Privilege Level	clear
Output Fields	This command produces no output.

Sample Output

clear ospf database-protection

```
user@host> clear ospf database-protection
```


clear (ospf | ospf3) io-statistics

List of Syntax	Syntax on page 4037 Syntax (EX Series Switch and QFX Series) on page 4037
Syntax	clear (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) input and output statistics.
Options	none —Clear OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf io-statistics on page 4037
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

List of Syntax	Syntax on page 4038 Syntax (EX Series Switch and QFX Series) on page 4038
Syntax	<pre>clear (ospf ospf3) neighbor <area <i>area-id</i>> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear (ospf ospf3) neighbor <area <i>area-id</i>> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Tear down Open Shortest Path First (OSPF) neighbor connections.
Options	<p>none—Tear down OSPF connections with all neighbors for all routing instances.</p> <p>area <i>area-id</i>—(Optional) Tear down neighbor connections for the specified area only.</p> <p>instance <i>instance-name</i>—(Optional) Tear down neighbor connections for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Tear down neighbor connections for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) Clear the state of the specified neighbor only.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show (ospf ospf3) neighbor on page 4080
List of Sample Output	clear ospf neighbor on page 4039

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor

```
user@host> clear ospf neighbor
```

clear (ospf | ospf3) overload

List of Syntax	Syntax on page 4040 Syntax (EX Series Switches) on page 4040
Syntax	<code>clear (ospf ospf3) overload</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>clear (ospf ospf3) overload</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).
Options	none —Clear the overload bit and rebuild LSAs for all routing instances. instance <i>instance-name</i> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf overload on page 4040
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```

clear (ospf | ospf3) statistics

List of Syntax	Syntax on page 4041 Syntax (EX Series Switch and QFX Series) on page 4041
Syntax	clear (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) statistics.
Options	none —Clear OSPF statistics. instance <i>instance-name</i> —(Optional) Clear statistics for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show (ospf ospf3) statistics on page 4097
List of Sample Output	clear ospf statistics on page 4041
Output Fields	See show (ospf ospf3) statistics for an explanation of output fields.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

Packet type	Total	Last 5 seconds
-------------	-------	----------------

	Sent	Received	Sent	Received
Hello	3254	2268	3	1
DbD	41	46	0	0
LSReq	8	7	0	0
LSUpdate	212	154	0	0
LSAck	65	98	0	0

DBDs retransmitted	:	3, last 5 seconds	:	0
LSAs flooded	:	12, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	3, last 5 seconds	:	0
LSAs requested	:	5, last 5 seconds	:	0
LSAs acknowledged	:	19, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:

626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3	1	3	1
DbD	0	0	0	0
LSReq	0	0	0	0
LSUpdate	0	0	0	0
LSAck	0	0	0	0

DBDs retransmitted	:	0, last 5 seconds	:	0
LSAs flooded	:	0, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	0, last 5 seconds	:	0
LSAs requested	:	0, last 5 seconds	:	0
LSAs acknowledged	:	0, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:

None

show (ospf | ospf3) backup coverage

Syntax	<pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> < logical-system (all <i>logical-system-name</i>)> <realm (ipv4-unicast ipv6-unicast)> <topology <i>topology-name</i>></pre>
Syntax (QFX Series)	<pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about the level of backup coverage available for all the nodes and prefixes in the network.
Options	<p>none—Display information about the level backup coverage for all OSPF routing instances in all logical systems.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.</p> <p>realm (ipv4-unicast ipv6-unicast)—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.</p> <p>topology (default <i>topology-name</i>)—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show (ospf ospf3) backup lsp
List of Sample Output	show ospf backup coverage on page 4044 show ospf3 backup coverage on page 4044
Output Fields	<p>Table 304 on page 4043 lists the output fields for the show (ospf ospf3) backup coverage command. Output fields are listed in the approximate order in which they appear.</p>

Table 304: show (ospf | ospf3) backup coverage Output Fields

Field Name	Field Description
Node Coverage	Information about backup coverage for each OSPF node.
Area	Area number. Area 0.0.0.0 is the backbone.

Table 304: show (ospf | ospf3) backup coverage Output Fields (*continued*)

Field Name	Field Description
Covered Nodes	Number of nodes for which backup coverage is available.
Total Nodes	Total number of OSPF nodes.
Route Coverage	Information about backup coverage for each type of OSPF route.
Path Type	Type of OSPF path: Intra , Inter , Ext1 , Ext2 , and All .
Covered Routes	For each path type, the number of routes for which backup coverage is available.
Total Routes	For each path type, the total number of configured routes.
Percent Covered	For all nodes and for each path type, the percentage for which backup coverage is available.

Sample Output

show ospf backup coverage

```
user@host> show ospf backup coverage
Topology default coverage:
```

Node Coverage:

Area	Covered Nodes	Total Nodes	Percent Covered
0.0.0.0	4	5	80.00%

Route Coverage:

Path Type	Covered Routes	Total Routes	Percent Covered
Intra	8	14	57.14%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	1	1	100.00%
All	9	15	60.00%

show ospf3 backup coverage

```
user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:
```

Area	Covered Nodes	Total Nodes	Percent Covered
0.0.0.0	4	5	80.00%

Route Coverage:

Path Type	Covered Routes	Total Routes	Percent Covered
-----------	-------------------	-----------------	--------------------

Intra	4	6	66.67%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	1	1	100.00%
All	5	7	71.43%

show (ospf | ospf3) backup neighbor

Syntax	<pre>show (ospf ospf3) backup neighbor <area <i>area-id</i>> <instance (default <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <topology (default ipv4-multicast <i>topology-name</i>)></pre>
Syntax (QFX Series)	<pre>show (ospf ospf3) backup neighbor <area <i>area-id</i>> <instance <i>instance-name</i>> <topology (default ipv4-multicast <i>topology-name</i>)></pre>
Release Information	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the neighbors through which direct next hops for the backup paths are available.
Options	<p>none—Display all neighbors that have direct next hops for backup paths.</p> <p>area <i>area-id</i>—(Optional) Display the area information.</p> <p>instance (default <i>instance-name</i>)—(Optional) Display information about the default routing instance or a particular routing instance.</p> <p>logical-system (default ipv4-multicast <i>logical-system-name</i>)—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.</p> <p>topology (default ipv4-multicast <i>topology-name</i>)—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show (ospf ospf3) backup spf
List of Sample Output	show ospf backup neighbor on page 4047
Output Fields	Table 305 on page 4046 lists the output fields for the show (ospf ospf3) backup neighbor command. Output fields are listed in the approximate order in which they appear.

Table 305: show (ospf | ospf3) backup neighbor Output Fields

Field Name	Field Description	Level of Output
Neighbor to Self Metric	Metric from the backup neighbor to the OSPF node.	All levels
Self to Neighbor Metric	Metric from the OSPF node to the backup neighbor.	All levels

Table 305: show (ospf |ospf3) backup neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Direct next-hop	Interface and address of the direct next hop.	All levels

Sample Output

show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

show ospf context-identifier

List of Syntax	Syntax on page 4048 Syntax (EX Series Switches and QFX Series) on page 4048
Syntax	<pre>show ospf context-identifier <brief detail> <area <i>area-id</i>> <context-id> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf context-identifier <brief detail> <area <i>area-id</i>> <context-id> <instance <i>instance-name</i>></pre>
Release Information	Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.
Options	<p>none—Display information about all context identifiers.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the context identifier for the specified area.</p> <p>context-id—(Optional) Display information about the specified context identifier.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the context identifier for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><i>egress-protection (Layer 2 circuit)</i> in the <i>Junos OS VPNs Configuration Guide</i><i>egress-protection (MPLS)</i> in the <i>Junos OS VPNs Configuration Guide</i>
List of Sample Output	show ospf context-identifier on page 4049 show ospf context-identifier detail on page 4049
Output Fields	Table 306 on page 4049 lists the output fields for the show ospf context-identifier command. Output fields are listed in the approximate order in which they appear.

Table 306: show ospf context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.	All levels
Status	State of the path: active or inactive .	All levels
Metric	Advertised OSPF metric.	All levels
Area	OSPF area number.	All levels
Other Advertisements	Other advertisements received by the OSPF node: <ul style="list-style-type: none"> • Advertising router—Address of the device that sent the advertisement. • Type—Type of OSPF path: inter-area and stub. • Metric—Advertised OSPF metric. • None—No additional advertisements were received by the OSPF node. 	detail

Sample Output

show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
Other Advertisements:
Advertising router: 8.8.8.103
Type: stub link
Metric: 65534
```

```
show ospf database
```

List of Syntax Syntax on page 4050 Syntax (EX Series Switches and QFX Series) on page 4050	
Syntax	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p>

area *area-id*—(Optional) Display the LSAs in a particular area.

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- [clear \(ospf | ospf3\) database on page 4033](#)

List of Sample Output

[show ospf database on page 4053](#)
[show ospf database brief on page 4053](#)
[show ospf database detail on page 4053](#)
[show ospf database extensive on page 4055](#)
[show ospf database summary on page 4057](#)

Output Fields

[Table 307 on page 4051](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 307: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: ASBRSum , Extern , Network , NSSA , OpaqArea , Router , or Summary .	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels

Table 307: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Seq	Link sequence number of the advertisement.	All levels
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: <ul style="list-style-type: none"> bits—Flags describing the routing device that generated the LSP. link count—Number of links in the advertisement. id—ID of a routing device or subnet on the link. data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. TOS count—Number of type-of-service (ToS) entries in the advertisement. TOS 0 metric—Metric for ToS 0. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. 	detail extensive
Network	Network link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. attached router—ID of the attached neighbor. 	detail extensive
Summary	Summary link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. 	detail extensive
Gen timer	How long until the LSA is regenerated.	extensive
Aging timer	How long until the LSA expires.	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed.	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires.	extensive
sent <i>hh:mm:ss</i> ago	How long ago the LSA was sent.	extensive
Last changed <i>hh:mm:ss</i> ago	How long ago the route was changed.	extensive

Table 307: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Change count	Number of times the route has changed.	extensive
Ours	Indicates that this is a local advertisement.	extensive
Router LSAs	Number of router link-state advertisements in the link-state database.	summary
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

Sample Output

show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   215    0x20 0x4112  48
Router     *10.255.71.242  10.255.71.242 0x80000002   214    0x20 0x11b1  48
Summary    *23.1.1.0       10.255.71.242 0x80000002   172    0x20 0x6d72  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000002   217    0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.52    10.255.71.52  0x80000004   174    0x20 0xd021  36
Router     *10.255.71.242  10.255.71.242 0x80000003   173    0x20 0xe191  36
Network    *23.1.1.1       10.255.71.242 0x80000002   173    0x20 0x9c76  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222    0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179    0x20 0x3942  36
Router     *10.255.71.242  10.255.71.242 0x80000003   177    0x20 0xf37d  36
Network    *24.1.1.1       10.255.71.242 0x80000002   177    0x20 0xc591  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *23.1.1.0       10.255.71.242 0x80000002   172    0x20 0x6d72  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222    0x28 0xeb3b  36

```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 4053](#).

show ospf database detail

```

user@host> show ospf database detail

```

```

    OSPF link state database, Area 0.0.0.1
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.70.103  10.255.70.103  0x80000002  261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000002  260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0      10.255.71.242  0x80000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000002  263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.2
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.52   10.255.71.52   0x80000004  220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1      10.255.71.242  0x80000002  219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0      10.255.71.242  0x80000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000001  268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.3
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.238  10.255.71.238  0x80000003  225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1      10.255.71.242  0x80000002  223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242

```

```

    attached router 10.255.71.238
Summary *12.1.1.0      10.255.71.242    0x80000001    263    0x20 0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
Summary *23.1.1.0      10.255.71.242    0x80000002    218    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
NSSA   *33.1.1.1      10.255.71.242    0x80000001    268    0x28 0xeb3b    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

show ospf database extensive

```

user@host> show ospf database extensive
    OSPF link state database, Area 0.0.0.1
Type      ID          Adv Rtr      Seq          Age    Opt  Cksum  Len
Router    10.255.70.103    10.255.70.103  0x80000002    286    0x20 0x4112  48
    bits 0x0, link count 2
    id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:14
    Installed 00:04:43 ago, expires in 00:55:14
    Last changed 00:04:43 ago, Change count: 2
Router    *10.255.71.242  10.255.71.242  0x80000002    285    0x20 0x11b1  48
    bits 0x3, link count 2
    id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:15
    Aging timer 00:55:15
    Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
    Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:57
    Aging timer 00:55:57
    Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248    0x20 0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA     *33.1.1.1      10.255.71.242    0x80000002    288    0x28 0x73bd    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
    Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
Type      ID          Adv Rtr      Seq          Age    Opt  Cksum  Len
Router    10.255.71.52     10.255.71.52    0x80000004    245    0x20 0xd021  36
    bits 0x0, link count 1

```

```

id 23.1.1.1, data 23.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:55
Installed 00:04:02 ago, expires in 00:55:55
Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 244 0x20 0xe191 36
bits 0x3, link count 1
id 23.1.1.1, data 23.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1 10.255.71.242 0x80000002 244 0x20 0x9c76 32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.52
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0 10.255.71.242 0x80000001 288 0x20 0xfeec 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0 10.255.71.242 0x80000002 248 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000001 293 0x28 0xe047 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
Last changed 00:04:53 ago, Change count: 1, Ours

OSPF link state database, Area 0.0.0.3
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 250 0x20 0x3942 36
bits 0x0, link count 1
id 24.1.1.1, data 24.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:50
Installed 00:04:07 ago, expires in 00:55:50
Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 248 0x20 0xf37d 36
bits 0x3, link count 1
id 24.1.1.1, data 24.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1 10.255.71.242 0x80000002 248 0x20 0xc591 32

```

```

mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.238
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288    0x20 0xfeec    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1      10.255.71.242    0x80000001    293    0x28 0xeb3b    36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours

```

show ospf database summary

```

user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:

```


advertising-router (*address* | *self*)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (*all* | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (*ipv4-multicast* | *ipv4-unicast* | *ipv6-multicast*)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- [clear \(ospf | ospf3\) database on page 4033](#)

List of Sample Output

[show ospf3 database brief on page 4064](#)
[show ospf3 database extensive on page 4064](#)
[show ospf3 database summary on page 4067](#)

Output Fields

[Table 308 on page 4059](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 308: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area <i>area-number</i>	Entries in the link-state database for this area.	brief detail extensive

Table 308: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	brief detail extensive
OSPF Link-Local link state database, interface interface-name	Entries in the link-local link-state database for this interface.	brief detail extensive
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: Extern , InterArPfx , InterArRtr , IntraArPrx , Link , Network , NSSA , or Router .	brief detail extensive
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	brief detail extensive
Adv Rtr	Address of the routing device that sent the advertisement.	brief detail extensive
Seq	Link sequence number of the advertisement.	brief detail extensive
Age	Time elapsed since the LSA was originated, in seconds.	brief detail extensive
Cksum	Checksum value of the LSA.	brief detail extensive
Len	Length of the advertisement, in bytes.	brief detail extensive
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	detail extensive
Options	Option bits carried in the router LSA.	detail extensive
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link. 	detail extensive
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	detail extensive
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	detail extensive
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	detail extensive
Metric	Cost of the router link.	detail extensive

Table 308: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	detail extensive
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	detail extensive
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	detail extensive

Table 308: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
options	Optional capabilities supported by the routing device.	detail extensive
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Prefix	IPv6 address prefix.	extensive
Prefix-options	Option bit associated with the prefix.	extensive
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of the route, which depends on the value of Type .	detail extensive
Type <i>n</i>	Type of external metric: Type 1 or Type 2 .	detail extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	detail extensive
Options	Option bits carried in the link LSA.	detail extensive
priority	Router priority of the interface attaching the originating routing device to the link.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive

Table 308: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. <ul style="list-style-type: none"> Router—Address prefixes are associated with a router LSA. Network—Address prefixes are associated with a network LSA. 	detail extensive
Ref-lsa-id	Link-state ID of the referenced LSA.	detail extensive
Ref-router-id	Advertising router ID of the referenced LSA.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>hh:mm:ss</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
<i>n</i> Router LSAs	Number of router LSAs in the link-state database.	summary
<i>n</i> Network LSAs	Number of network LSAs in the link-state database.	summary

Table 308: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>n</i> InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	summary
<i>n</i> InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
<i>n</i> IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
<i>n</i> Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface <i>interface-name</i>	Name of the interface for which link-local LSA information is displayed.	summary
<i>n</i> Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```

user@host> show ospf3 database brief
    OSPF3 link state database, area 0.0.0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  Router    0.0.0.1        10.255.4.85  0x80000003   885  0xa697  40
  Router    *0.0.0.1        10.255.4.93  0x80000002   953  0xc677  40
  InterArPfx *0.0.0.2        10.255.4.93  0x80000001   910  0xb96f  44
  InterArRtr *0.0.0.1        10.255.4.93  0x80000001   910  0xe159  32
  IntraArPfx *0.0.0.1        10.255.4.93  0x80000002   432  0x788f  72

    OSPF3 link state database, area 0.0.0.1
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  Router    *0.0.0.1        10.255.4.93  0x80000003   916  0xea40  40
  Router    0.0.0.1        10.255.4.97  0x80000006   851  0xc95b  40
  Network    0.0.0.2        10.255.4.97  0x80000002   916  0x4598  32
  InterArPfx *0.0.0.1        10.255.4.93  0x80000002   117  0xa980  44
  InterArPfx *0.0.0.2        10.255.4.93  0x80000002    62  0xd47e  44
  NSSA      0.0.0.1        10.255.4.97  0x80000002   362  0x45ee  44
  IntraArPfx 0.0.0.1        10.255.4.97  0x80000006   851  0x2f77  52

    OSPF3 AS SCOPE link state database
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  Extern    0.0.0.1        10.255.4.85  0x80000002    63  0x9b86  44
  Extern    *0.0.0.1        10.255.4.93  0x80000001   910  0x59c9  44

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  Link      *0.0.0.2        10.255.4.93  0x80000003   916  0x4dab  64

```

show ospf3 database extensive

```

user@host> show ospf3 database extensive
    OSPF3 link state database, area 0.0.0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  Router    0.0.0.1        10.255.4.85  0x80000003  1028  0xa697  40

```

```

bits 0x2, Options 0x13
Type PointToPoint (1), Metric 10
  Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
Aging timer 00:42:51
Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router *0.0.0.1 10.255.4.93 0x80000002 1096 0xc677 40
bits 0x3, Options 0x13
Type PointToPoint (1), Metric 10
  Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
Gen timer 00:00:40
Aging timer 00:41:44
Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
Ours
InterArPfx *0.0.0.2 10.255.4.93 0x80000001 1053 0xb96f 44
Prefix feee::10:10:2:0/126
Prefix-options 0x0, Metric 10
Gen timer 00:17:02
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
InterArPfx *0.0.0.3 10.255.4.93 0x80000001 1053 0x71d3 44
Prefix feee::10:255:4:97/128
Prefix-options 0x0, Metric 10
Gen timer 00:21:07
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
InterArRtr *0.0.0.1 10.255.4.93 0x80000001 1053 0xe159 32
Dest-router-id 10.255.4.97, Options 0x19, Metric 10
Gen timer 00:29:18
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
IntraArPfx 0.0.0.1 10.255.4.85 0x80000002 1028 0x2403 72
Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
Prefix-count 2
Prefix feee::10:255:4:85/128
  Prefix-options 0x2, Metric 0
Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
Aging timer 00:42:51
Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1 10.255.4.93 0x80000002 575 0x788f 72
Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
Prefix-count 2
Prefix feee::10:255:4:93/128
  Prefix-options 0x2, Metric 0
Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
Gen timer 00:33:23
Aging timer 00:50:24
Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
OSPF3 link state database, area 0.0.0.1
Type ID Adv Rtr Seq Age Cksum Len
Router *0.0.0.1 10.255.4.93 0x80000003 1059 0xea40 40
bits 0x3, Options 0x19
Type Transit (2), Metric 10
  Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
Gen timer 00:08:51
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago

```

```

Router      0.0.0.1          10.255.4.97      0x80000006   994  0xc95b  40
  bits 0x2, Options 0x19
  Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network     0.0.0.2          10.255.4.97      0x80000002   1059 0x4598  32
  Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1          10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001   1089 0x9bbb  44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA        0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx  0.0.0.1          10.255.4.97      0x80000006   994  0x2f77  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
    Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx  0.0.0.3          10.255.4.97      0x80000002   1059 0x4446  52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  OSPF3 AS SCOPE link state database
  Type      ID      Adv Rtr      Seq      Age  Cksum  Len
Extern     0.0.0.1      10.255.4.85  0x80000002  206  0x9b86  44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern     *0.0.0.1      10.255.4.93  0x80000001  1053 0x59c9  44
  Prefix feee::200:200:1:0/124

```

```

Prefix-options 0x0, Metric 10, Type 2,
Gen timer 00:25:12
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

```

```

OSPF3 Link-Local link state database, interface ge-1/3/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      *0.0.0.2      10.255.4.93  0x80000003  1059 0x4dab  64
fe80::290:69ff:fe39:1cdb
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Gen timer 00:12:56
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link      0.0.0.2      10.255.4.97  0x80000003  205  0xa87d  64
fe80::290:69ff:fe38:883e
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Aging timer 00:56:35
Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

```

```

OSPF3 Link-Local link state database, interface so-2/2/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.2      10.255.4.85  0x80000002  506  0x42bb  64
fe80::280:42ff:fe10:f169
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Aging timer 00:51:34
Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link      *0.0.0.3      10.255.4.93  0x80000002  505  0x6b7a  64
fe80::280:42ff:fe10:f177
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Gen timer 00:37:28
Aging timer 00:51:35
Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
Ours

```

show ospf3 database summary

```

user@host> show ospf3 database summary
Area 0.0.0.0:
  2 Router LSAs
  1 InterArPfx LSAs
  1 InterArRtr LSAs
  1 IntraArPfx LSAs
Area 0.0.0.1:
  2 Router LSAs
  1 Network LSAs
  2 InterArPfx LSAs
  1 NSSA LSAs
  1 IntraArPfx LSAs
Externals:
  2 Extern LSAs
Interface ge-1/3/0.0:
  1 Link LSAs
Interface lo0.0:

```

Interface so-2/2/0.0:
1 Link LSAs

show (ospf | ospf3) interface

List of Syntax	Syntax on page 4069 Syntax (EX Series Switches and QFX Series) on page 4069
Syntax	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>area option introduced in Junos OS Release 9.2.</p> <p>area option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the status of OSPF interfaces.
Options	<p>none—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><i>interface-name</i>—(Optional) Display information for the specified interface.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view

List of Sample Output [show ospf interface brief on page 4072](#)
[show ospf interface detail on page 4072](#)
[show ospf3 interface detail on page 4072](#)
[show ospf interface detail\(When Multiarea Adjacency Is Configured\) on page 4072](#)
[show ospf interface area area-id on page 4074](#)
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 4074](#)
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 4074](#)

Output Fields [Table 309 on page 4070](#) lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 309: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: LAN , NBMA , P2MP , P2P , or Virtual .	detail extensive
Address	IP address of the neighbor.	detail extensive
Mask	Netmask of the neighbor.	detail extensive
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	detail extensive
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	detail extensive
MTU	Interface maximum transmission unit (MTU).	detail extensive
Cost	Interface cost (metric).	detail extensive
DR addr	Address of the designated router.	detail extensive
BDR addr	Address of the backup designated router.	detail extensive
Adj count	Number of adjacent neighbors.	detail extensive
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	detail extensive

Table 309: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the DoNotAge bit set. As a result, LSAs are refreshed only when a change occurs.	extensive
Priority	Router priority used in designated router (DR) election on this interface.	detail extensive
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	extensive
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	extensive
Descriptor list	List of packet descriptors.	extensive
Hello	Configured value for the hello timer.	detail extensive
Dead	Configured value for the dead timer.	detail extensive
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured. 	detail extensive
Topology	(Multiarea adjacency) Name of topology: default or name .	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync , in holddown , and not supported .	extensive
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	extensive
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the remaining field displays the number of seconds that remain until the configured hold timer expires.	extensive
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	detail extensive
Active key ID	(OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key.	detail extensive
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST .	detail extensive

Table 309: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief
Intf           State   Area      DR ID      BDR ID      Nbrs
at-5/1/0.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
ge-2/3/0.0     DR      0.0.0.0    192.168.4.16 192.168.4.15 1
lo0.0          DR      0.0.0.0    192.168.4.16 0.0.0.0     0
so-0/0/0.0     Down    0.0.0.0    0.0.0.0    0.0.0.0     0
so-6/0/1.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
so-6/0/2.0     Down    0.0.0.0    0.0.0.0    0.0.0.0     0
so-6/0/3.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
```

show ospf interface detail

```
user@host> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/1.0     BDR    0.0.0.0    192.168.37.12 10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa
```

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
Interface      State   Area      DR-ID      BDR-ID      Nbrs
so-0/0/3.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub
```

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
regress@router> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0    10.255.245.2 0.0.0.0     0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 127.0.0.1, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
```

```

Topology default (ID 0) -> Cost: 0
100.0          DR          0.0.0.0          10.255.245.2    0.0.0.0          0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 10.255.245.2, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 0
so-0/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary

```

```
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
```

show ospf interface extensive (When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

show ospf interface extensive (When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID
Nbrs
so-1/0/3.0     Down    0.0.0.0   0.0.0.0    0.0.0.0
0
Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddtime: 10 seconds, remaining: 1
```

show (ospf | ospf3) io-statistics

List of Syntax	Syntax on page 4075 Syntax (EX Series Switch and QFX Series) on page 4075
Syntax	show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) input and output statistics.
Options	none —Display OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear (ospf ospf3) statistics on page 4041
List of Sample Output	show ospf io-statistics on page 4076
Output Fields	Table 310 on page 4075 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear.

Table 310: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Sample Output

show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1  
Receive errors:  
None
```


show (ospf | ospf3) log

List of Syntax	Syntax on page 4077 Syntax (EX Series Switch and QFX Series) on page 4077
Syntax	<pre>show (ospf ospf3) log <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <topology <i>topology-name</i>></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) log <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>topology option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.
Options	<p>none—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology <i>topology-name</i>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf log on page 4078 show ospf log topology voice on page 4078
Output Fields	<p>Table 311 on page 4077 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear.</p>

Table 311: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (w) and days (d), since the SPF calculation was made.

Table 311: show (ospf | ospf3) log Output Fields (*continued*)

Field Name	Field Description
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field.

Sample Output

show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF            0.000070
1w4d 17:25:58 Stub            0.000019
1w4d 17:25:58 Interarea       0.000054
1w4d 17:25:58 External        0.000005
1w4d 17:25:58 Cleanup         0.000203
1w4d 17:25:58 Total          0.000537
1w4d 17:24:48 SPF            0.000125
1w4d 17:24:48 Stub            0.000017
1w4d 17:24:48 SPF            0.000100
1w4d 17:24:48 Stub            0.000016
1w4d 17:24:48 Interarea       0.000056
1w4d 17:24:48 External        0.000005
1w4d 17:24:48 Cleanup         0.000238
1w4d 17:24:48 Total          0.000600
...

```

show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF            0.000116
00:06:11      Stub            0.000114
00:06:11      Interarea       0.000126
00:06:11      External        0.000067
00:06:11      NSSA            0.000037
00:06:11      Cleanup         0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF            0.000140
00:13:33      Stub            0.000116
00:13:43      Interarea       0.000128
00:13:33      External        0.000075
00:13:38      NSSA            0.000039
00:13:53      Cleanup         0.000657

    Last 100 events

```

When	Type	Elapsed
00:13:53	SPF	0.000090
00:13:53	Stub	0.000041
00:13:53	Interarea	0.000123
00:13:53	External	0.000040
00:13:53	NSSA	0.000038
00:13:53	Cleanup	0.000657
00:13:53	Total	0.001252
.		
.		
00:06:11	SPF	0.000116
00:06:11	Stub	0.000114
00:06:11	Interarea	0.000126
00:06:11	External	0.000067
00:06:11	NSSA	0.000037
00:06:11	Cleanup	0.000186
00:06:11	Total	0.000818

show (ospf | ospf3) neighbor

List of Syntax [Syntax on page 4080](#)
 [Syntax \(EX Series Switches and QFX Series\) on page 4080](#)

Syntax `show (ospf | ospf3) neighbor`
 `<brief | detail | extensive>`
 `<area area-id>`
 `<instance (all | instance-name)>`
 `<interface interface-name>`
 `<logical-system (all | logical-system-name)>`
 `<neighbor>`
 `<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>`

Syntax (EX Series Switches and QFX Series) `show (ospf | ospf3) neighbor`
 `<brief | detail | extensive>`
 `<area area-id>`
 `<instance (all | instance-name)>`
 `<interface interface-name>`
 `<neighbor>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 instance all option introduced in Junos OS Release 9.1.
 instance all option introduced in Junos OS Release 9.1 for EX Series switches.
 area, **interface**, and **realm** options introduced in Junos OS Release 9.2.
 area and **interface** options introduced in Junos OS Release 9.2 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display information about OSPF neighbors.

CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the **show (ospf | ospf3) neighbor** command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the **show (ospf | ospf3) neighbor** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.

Options **none**—Display standard information about all OSPF neighbors for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the OSPF neighbors for the specified area.

instance (all | *instance-name*)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

interface *interface-name*—(Optional) Display information about OSPF neighbors for the specified logical interface.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) neighbor on page 4038](#)

List of Sample Output [show ospf neighbor brief on page 4083](#)
[show ospf neighbor detail on page 4083](#)
[show ospf neighbor extensive on page 4084](#)
[show ospf3 neighbor detail on page 4085](#)
[show ospf neighbor area area-id on page 4085](#)
[show ospf neighbor interface interface-name on page 4085](#)
[show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 4085](#)

Output Fields [Table 312 on page 4081](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 312: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

Table 312: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the neighbor:</p> <ul style="list-style-type: none"> • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater. 	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	extensive
Link state retransmission list	<p>Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed:</p> <ul style="list-style-type: none"> • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpaqueArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement. 	detail extensive

Table 312: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address.	detail extensive
area	Area that the neighbor is in.	detail extensive
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	detail extensive
opt	Option bits received in the hello packets from the neighbor.	detail extensive
DR or DR-ID	Address of the designated router.	detail extensive
BDR or BDR-ID	Address of the backup designated router.	detail extensive
Up	Length of time since the neighbor came up.	detail extensive
adjacent	Length of time since the adjacency with the neighbor was established.	detail extensive

Sample Output

show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2      ge-1/2/0.10     ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11     Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128   33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2     ge-1/2/0.1    Full      10.5.1.2 128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.56.0    172.25.27.82 0x8000004d
Router  10.5.1.94    10.5.1.94    0x8000005c
Network 10.5.24.2    10.5.1.94    0x80000036
Summary 10.8.57.0    172.25.27.82 0x80000024
Extern  1.10.90.0   10.8.1.2     0x80000041
Extern  1.4.109.0    10.6.1.2     0x80000041
Router  10.5.1.190    10.5.1.190   0x8000005f
Network 10.5.48.2    10.5.1.190   0x8000003d
Summary 10.8.58.0    172.25.27.82 0x8000004d
Extern  1.10.91.0    10.8.1.2     0x80000041
Extern  1.4.110.0    10.6.1.2     0x80000041
Router  10.5.1.18     10.5.1.18    0x8000005f
Network 10.5.5.2     10.5.1.18    0x80000033
Summary 10.8.59.0    172.25.27.82 0x8000003a
Summary 10.8.62.0    172.25.27.82 0x80000025

10.5.10.2    ge-1/2/0.10    ExStart  10.5.1.38      128   38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2    ge-1/2/0.11    Full      10.5.1.42      128   33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.58.0    172.25.27.82 0x8000004d

```


Extern	1.10.91.0	10.8.1.2	0x80000041
Extern	1.1.247.0	10.5.1.2	0x8000003f
Extern	1.4.110.0	10.6.1.2	0x80000041
Router	10.5.1.18	10.5.1.18	0x8000005f
Network	10.5.5.2	10.5.1.18	0x80000033
Summary	10.8.59.0	172.25.27.82	0x8000003a

show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
Area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43
```

show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0          Full     10.255.245.5 128    37
Area 1.1.1.1
```

show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    32
Area 2.2.2.2
```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

List of Syntax	Syntax on page 4086 Syntax (EX Series Switch and QFX Series) on page 4086
Syntax	<code>show (ospf ospf3) overview</code> <code><brief extensive></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><realm (ipv4-multicast ipv4-unicast ipv6-multicast)></code>
Syntax (EX Series Switch and QFX Series)	<code>show (ospf ospf3) overview</code> <code><brief extensive></code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) overview information.
Options	none —Display standard information about all OSPF neighbors for all routing instances. brief extensive —(Optional) Display the specified level of output. instance <i>instance-name</i> —(Optional) Display all OSPF interfaces under the named routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	view
List of Sample Output	show ospf overview on page 4088 show ospf overview (With Database Protection) on page 4089 show ospf3 overview (With Database Protection) on page 4089 show ospf overview extensive on page 4089
Output Fields	Table 225 on page 2364 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.

Table 313: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 313: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the [edit protocols ospf area area-id] level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

show ospf overview (With Database Protection)

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed

```

show ospf3 overview (With Database Protection)

```

user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed

```

show ospf overview extensive

```

user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes

```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

show (ospf | ospf3) route

List of Syntax [Syntax on page 4091](#)
 [Syntax \(EX Series Switch and QFX Series\) on page 4091](#)

Syntax show (ospf | ospf3) route
 <brief | detail | extensive>
 <abr | asbr | extern | inter | intra>
 <destination>
 <instance (default | ipv4-multicast | *instance-name*)>
 <logical-system (default | ipv4-multicast | *logical-system-name*)>
 <network>
 <no-backup-coverage>
 <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
 <router>
 <topology (default | ipv4-multicast | *topology-name*)>
 <transit>

Syntax (EX Series Switch and QFX Series) show (ospf | ospf3) route
 <brief | detail | extensive>
 <abr | asbr | extern | inter | intra>
 <destination>
 <instance *instance-name*>
 <network>
 <no-backup-coverage>
 <router>
 <topology (default | ipv4-multicast | *topology-name*)>
 <transit>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 topology option introduced in Junos OS Release 9.0.
 realm option introduced in Junos OS Release 9.2.
 Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display the entries in the Open Shortest Path First (OSPF) routing table.

Options **none**—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.

destination—Display routes to the specified IP address (with optional destination prefix length).

brief | detail | extensive—(Optional) Display the specified level of output.

abr—(Optional) Display routes to area border routers.

asbr—(Optional) Display routes to autonomous system border routers.

extern—(Optional) Display external routes.

inter—(Optional) Display interarea routes.

intra—(Optional) Display intra-area routes.

instance (**default** | **ipv4-multicast** | **instance-name**)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (**default** | **ipv4-multicast** | **logical-system-name**)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (**default** | **ipv4-multicast** | **topology-name**)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level

view

List of Sample Output

[show ospf route on page 4094](#)
[show ospf route detail on page 4094](#)
[show ospf3 route on page 4094](#)
[show ospf3 route detail on page 4095](#)
[show ospf route topology voice on page 4095](#)

Output Fields

[Table 314 on page 4092](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 314: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels

Table 314: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
Path type	How the route was learned: <ul style="list-style-type: none"> • Inter—Interarea route • Ext1—External type 1 route • Ext2—External type 2 route • Intra—Intra-area route 	All levels
Route type	The type of routing device from which the route was learned: <ul style="list-style-type: none"> • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard. 	All levels
NH Type	Next-hop type: LSP or IP .	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the NH Type is IP , then it is the address of the next hop. If the NH Type is LSP , then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	detail
Origin	Router from which the route was learned.	detail
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	detail
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	detail
Fwd NZ	Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA.	detail

Table 314: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR , AS BR , Area/AS BR , or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router. 	detail
priority	The priority assigned to the prefix: <ul style="list-style-type: none"> • high • medium • low <p>NOTE: The priority field applies only to routes of type Network.</p>	detail

Sample Output

show ospf route

```

user@host> show ospf route
Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type                Interface  addr/label
10.255.71.12          Intra  Router   IP     1        fe-0/0/2.0  192.16.22.86
10.255.71.13/32       Intra  Network  IP     0         lo0.0
192.168.222.84/30     Intra  Network  LSP    1        fe-0/0/2.0   1sp-ab

```

show ospf route detail

```

user@host> show ospf route detail
Topology default Route Table:

Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type                Interface  addr/label
10.255.14.174          Inter  AS BR     IP     210      t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185
10.255.14.178          Intra  Router   IP     200      t3-3/1/3.0
  area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30          Intra  Network  IP     10       t3-3/1/2.0
  area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32           Inter  Network  IP     210      t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24           Ext2   Network  IP     0        t1-3/0/1.0
  area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30           Inter  Network  IP     220      t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority high

```

show ospf3 route

```

user@host> show ospf3 route
Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type                Interface  addr/label

```

```

10.255.71.13      Intra Router      IP      1
NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1      Intra Router      IP      40 fxp1.1      192.168.36.17

    area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3      Intra AS BR      IP      1 fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32   Intra Network    IP      40 fxp1.1      192.168.36.17

    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32   Intra Network    IP      0 lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32   Intra Network    IP      1 fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3, priority low
    Intra Transit      IP      1
NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network    IP      1
NH-interface fe-0/0/2.0
abcd::71:12/128   Intra Network    IP      0
NH-interface lo0.0
abcd::71:13/128   Intra Network    LSP     1
NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix
Path      Route      NH      Metric
type      type
10.255.14.174 Intra Area/AS BR IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178 Intra Router IP 200
NH-interface t3-3/1/3.0
Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2 Intra Transit IP 200
NH-interface t1-3/0/1.0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128 Inter Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48 Ext1 Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128 Intra Network IP 0
NH-interface lo0.0
Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

show ospf route topology voice

```

user@host show ospf route topology voice

```

Topology voice Route Table:

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	Nexthop addr/label
10.255.8.2	Intra	Router	IP	1	so-0/2/0.0	
10.255.8.3	Intra	Router	IP	2	so-0/2/0.0	
10.255.8.1/32	Intra	Network	IP	0	lo0.0	
10.255.8.2/32	Intra	Network	IP	1	so-0/2/0.0	
10.255.8.3/32	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.0/29	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.44/30	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.46/32	Intra	Network	IP	1	so-0/2/0.0	
192.168.8.48/30	Intra	Network	IP	1	so-0/2/1.0	
192.168.8.52/30	Intra	Network	IP	2	so-0/2/0.0	
192.168.9.44/30	Intra	Network	IP	1	so-0/2/0.0	
192.168.9.45/32	Intra	Network	IP	2	so-0/2/0.0	

show (ospf | ospf3) statistics

List of Syntax	Syntax on page 4097 Syntax (EX Series Switch and QFX Series) on page 4097
Syntax	show (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display OSPF statistics.
Options	none —Display OSPF statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display all statistics for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear (ospf ospf3) statistics on page 4041
List of Sample Output	show ospf statistics on page 4099 show ospf statistics logical-system all on page 4099 show ospf3 statistics on page 4100
Output Fields	Table 315 on page 4097 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear.

Table 315: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.

Table 315: show (ospf | ospf3) statistics Output Fields (*continued*)

Field Name	Field Description
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	<p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.
Receive errors	<p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches <p>If there are no receive errors, the output displays none.</p>

Sample Output

show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                31        14
  DbD                 9         10
  LSReq               2          2
LSUpdate             8         16
  LSAck              9          9
                   Last 5 seconds
                   Sent      Received
Hello                2          2
  DbD                 0          0
  LSReq               0          0
LSUpdate             0          0
  LSAck              0          0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
  115923 no virtual link found

```

show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
  DbD                3          2
  LSReq              1          1
LSUpdate           2752       1825
  LSAck             1821       2747
                   Last 5 seconds
                   Sent      Received
Hello                1          0
  DbD                 0          0
  LSReq               0          0
LSUpdate            0          0
  LSAck              0          0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```

```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                313698      313695         0         0
  DbD                  2         3         0         0
  LSReq                1         1         0         0
LSUpdate             1825      2752         0         0
LSAck                2747      1821         0         0

DBDs retransmitted   :                0, last 5 seconds :      0
LSAs flooded         :                1825, last 5 seconds :      0
LSAs flooded high-prio :                10, last 5 seconds :      0
LSAs retransmitted   :                0, last 5 seconds :      0
LSAs transmitted to nbr:                1, last 5 seconds :      0
LSAs requested       :                2, last 5 seconds :      0
LSAs acknowledged   :                2748, last 5 seconds :      0

Flood queue depth    :                0
Total rexmit entries :                0
db summaries         :                0
lsreq entries        :                0

Receive errors:
None
-----

```

show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                0         0         0         0
  DbD                  0         0         0         0
  LSReq                0         0         0         0
LSUpdate             0         0         0         0
LSAck                0         0         0         0

DBDs retransmitted   :                0, last 5 seconds :      0
LSAs flooded         :                0, last 5 seconds :      0
LSAs flooded high-prio :                0, last 5 seconds :      0
LSAs retransmitted   :                0, last 5 seconds :      0
LSAs transmitted to nbr:                0, last 5 seconds :      0
LSAs requested       :                0, last 5 seconds :      0
LSAs acknowledged   :                0, last 5 seconds :      0

Flood queue depth    :                0
Total rexmit entries :                0
db summaries         :                0
lsreq entries        :                0

Receive errors:
None

```


PART 15

Routing Information Protocol

- [Overview on page 4103](#)
- [Configuration on page 4109](#)
- [Administration on page 4197](#)

Overview

- [RIP Overview on page 4103](#)

RIP Overview

- [RIP Overview on page 4103](#)

RIP Overview

RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using the hop count as the metric.

In a RIP network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.



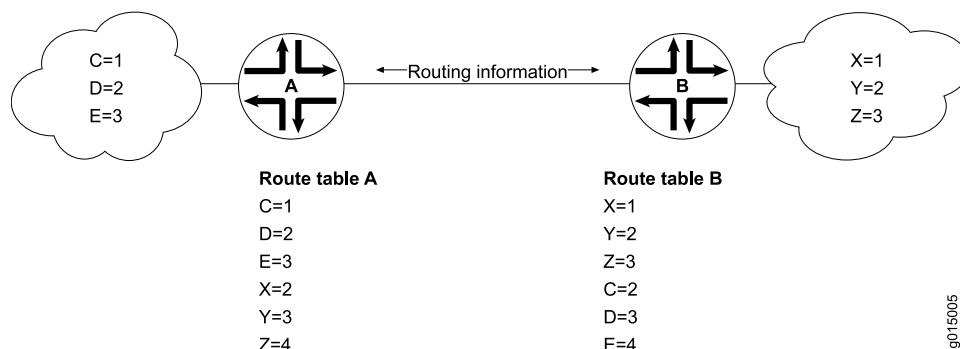
NOTE: In general, the term *RIP* refers to RIP version 1 and RIP version 2.

This topic contains the following sections:

- [Distance-Vector Routing Protocols on page 4103](#)
- [RIP Protocol Overview on page 4104](#)
- [RIP Packets on page 4105](#)
- [Maximizing Hop Count on page 4106](#)
- [Split Horizon and Poison Reverse Efficiency Techniques on page 4106](#)
- [Limitations of Unidirectional Connectivity on page 4107](#)

Distance-Vector Routing Protocols

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database. [Figure 130 on page 4104](#) shows how distance-vector routing works.

Figure 130: Distance-Vector Protocol

In [Figure 130 on page 4104](#), Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When Router A receives routing information from Router B, it adds 1 to the hop count to determine the new hop count. For example, Router X has a hop count of 1, but when Router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to Router X through Router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

RIP Protocol Overview

The RIP IGP uses the Bellman-Ford, or *distance-vector*, algorithm to determine the best route to a destination. RIP uses the hop count as the metric. RIP enables hosts and routers to exchange information for computing routes through an IP-based network. RIP is intended to be used as an IGP in reasonably homogeneous networks of moderate size.

The Junos® operating system (Junos OS) supports RIP versions 1 and 2.



NOTE: RIP is not supported for multipoint interfaces.

RIP version 1 packets contain the minimal information necessary to route packets through a network. However, this version of RIP does not support authentication or subnetting.

RIP uses User Datagram Protocol (UDP) port 520.

RIP has the following architectural limitations:

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIP depends on counting to infinity to resolve certain unusual situations—When the network consists of several hundred routers, and when a routing loop has formed, the amount of time and network bandwidth required to resolve a next hop might be great.
- RIP uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

RIP Packets

RIP packets contain the following fields:

- Command—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically and also when a request message is received. Periodic response messages are called *update messages*. Update messages contain the command and version fields and 25 destinations (by default), each of which includes the destination IP address and the metric to reach that destination.



NOTE: Beginning with Junos OS Release 11.1, three additional command field types are available to support RIP demand circuits. When you configure an interface for RIP demand circuits, the command field indicates whether the packet is an update request, update response, or update acknowledge message. Neighbor interfaces send updates on demand, not periodically. These command field types are only valid on interfaces configured for RIP demand circuits. For more detailed information, see *RIP Demand Circuits Overview*.

- Version number—Version of RIP that the originating router is running.
- Address family identifier—Address family used by the originating router. The family is always IP.
- Address—IP address included in the packet.
- Metric—Value of the metric advertised for the address.
- Mask—Mask associated with the IP address (RIP version 2 only).
- Next hop—IP address of the next-hop router (RIP version 2 only).

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If Router A is many hops away from a new host, Router B, the route to B might take significant time to propagate through the network and be imported into Router A's routing table. If the two routers are 5 hops away from each other, Router A cannot import the route to Router B until 2.5 minutes after Router B is online (30 seconds per hop). For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the *network diameter*.


Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.

If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as *split horizon*, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned.

[Figure 131 on page 4106](#) shows an example of the split horizon technique.


Figure 131: Split Horizon Example



In [Figure 131 on page 4106](#), Router A advertises routes to Routers C, D, and E to Router B. In this example, Router A can reach Router C in 2 hops. When Router A advertises the route to Router B, Router B imports it as a route to Router C through Router A in 3 hops. If Router B then readvertised this route to Router A, Router A would import it as a route to Router C through Router B in 4 hops. However, the advertisement from Router B to Router A is unnecessary, because Router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.

Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If Router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. [Figure 132 on page 4107](#) shows an example of the poison reverse technique.

Figure 132: Poison Reverse Example




In [Figure 132 on page 4107](#), Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Routers C, D, and E are definitely not reachable through Router A.

Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As [Figure 133 on page 4107](#) shows, RIP networks are limited by their unidirectional connectivity.

Figure 133: Limitations of Unidirectional Connectivity



In [Figure 133 on page 4107](#), Routers A and D flood their routing table information to Router B. Because the path to Router E has the fewest hops when routed through Router A, that route is imported into Router B's forwarding table. However, suppose that Router A can transmit traffic but is not receiving traffic from Router B because of an unavailable link or invalid routing policy. If the only route to Router E is through Router A, any traffic destined for Router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake.

**Related
Documentation**

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- *RIP Configuration Overview*
- [Example: Configuring RIP on page 4109](#)

CHAPTER 46

Configuration

- [RIP Configuration Tasks on page 4109](#)
- [RIP Configuration Statements on page 4171](#)

RIP Configuration Tasks

- [Example: Configuring RIP on page 4109](#)
- [Example: Configuring Authentication for RIP Routes on page 4116](#)
- [Example: Configuring BFD for RIP on page 4122](#)
- [Example: Configuring BFD Authentication for RIP on page 4128](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4136](#)
- [Examples: Controlling Traffic with Metrics in a RIP Network on page 4142](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4150](#)
- [Example: Redistributing Routes Among RIP Instances on page 4154](#)
- [Example: Configuring RIP Timers on page 4159](#)
- [Example: Tracing RIP Protocol Traffic on page 4166](#)

Example: Configuring RIP

- [Understanding Basic RIP Routing on page 4109](#)
- [Example: Configuring a Basic RIP Network on page 4110](#)

Understanding Basic RIP Routing

RIP is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). By default, RIP does not advertise the subnets that are directly connected through the device's interfaces. For traffic to pass through a RIP network, you must create a routing policy to export these routes. Advertising only the direct routes propagates the routes to the immediately adjacent RIP-enabled router only. To propagate all routes through the entire RIP network, you must configure the routing policy to export the routes learned through RIP.

Example: Configuring a Basic RIP Network

This example shows how to configure a basic RIP network.

- [Requirements on page 4110](#)
- [Overview on page 4110](#)
- [Configuration on page 4110](#)
- [Verification on page 4113](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


In this example, you configure a basic RIP network, create a RIP group called **rip-group**, and add the directly connected interfaces to the RIP group. Then you configure a routing policy to advertise direct routes using policy statement **advertise-routes-through-rip**.

By default, Junos OS does not advertise RIP routes, not even routes that are learned through RIP. To advertise RIP routes, you must configure and apply an export routing policy that advertises RIP-learned and direct routes.

In Junos OS, you do not need to configure the RIP version. RIP version 2 is used by default.

To use RIP on the device, you must configure RIP on all of the RIP interfaces within the network. [Figure 134 on page 4110](#) shows the topology used in this example.

Figure 134: Sample RIP Network Topology



"CLI Quick Configuration" on page 4110 shows the configuration for all of the devices in [Figure 134 on page 4110](#). The section "[Step-by-Step Procedure](#)" on page 4111 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
```

```

set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a basic RIP network:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

```

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Routing Table on page 4113](#)
- [Looking at the Routes That Device R1 Is Advertising to Device R2 on page 4113](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4114](#)
- [Verifying the RIP-Enabled Interfaces on page 4114](#)
- [Verifying the Exchange of RIP Messages on page 4114](#)
- [Verifying Reachability of All Hosts in the RIP Network on page 4115](#)

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes..

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 00:59:15, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32    *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32    *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32   *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32   *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32     *[RIP/100] 00:45:09, metric 1
                 MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol rip** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

Looking at the Routes That Device R1 Is Advertising to Device R2

Purpose Verify that Device R1 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R1> show route advertising-protocol rip 10.0.0.1
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32    *[Direct/0] 05:18:26
                 > via lo0.1
192.168.1.1/32   *[Direct/0] 05:18:25
                 > via lo0.1
```

Meaning Device R1 is sending routes to its directly connected networks.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30          *[RIP/100] 02:31:22, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32       *[RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32       *[RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32      *[RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32      *[RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Meaning Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIP.

Verifying the RIP-Enabled Interfaces

Purpose Verify that all RIP-enabled Interfaces are available and active.

Action From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-1/2/0.1	Up	10.0.0.1	224.0.0.9	mcast	both	1

Meaning The output shows that the RIP-enabled interface on Device R1 is operational.

In general for this command, the output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Local State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

Verifying the Exchange of RIP Messages

Purpose Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

Action From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter              Total    Last 5 min    Last minute
-----
Updates Sent          2669         10           2
Triggered Updates Sent    2           0           0
Responses Sent          0           0           0
Bad Messages           0           0           0
RIPv1 Updates Received    0           0           0
RIPv1 Bad Route Entries    0           0           0
RIPv1 Updates Ignored     0           0           0
RIPv2 Updates Received   2675        11           2
RIPv2 Bad Route Entries    0           0           0
RIPv2 Updates Ignored     0           0           0
Authentication Failures    0           0           0
RIP Requests Received     0           0           0
RIP Requests Ignored      0           0           0
none                     0           0           0

```

Meaning The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

Verifying Reachability of All Hosts in the RIP Network

Purpose Use the **traceroute** command on each loopback address in the network to verify that all hosts in the RIP network are reachable from each Juniper Networks device.

Action From operational mode, enter the **traceroute** command.

```

user@R1> traceroute 192.168.3.3
traceroute to 192.168.3.3 (192.168.3.3), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.094 ms  1.028 ms  0.957 ms
 2  192.168.3.3 (192.168.3.3)  1.344 ms  2.245 ms  2.125 ms

```

Meaning Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

**Related
Documentation**

- *Example: Configuring Point-to-Multipoint RIP Networks*

Example: Configuring Authentication for RIP Routes

- [Understanding RIP Authentication on page 4116](#)
- [Example: Configuring Route Authentication for RIP on page 4116](#)
- [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 4121](#)
- [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 4121](#)

Understanding RIP Authentication

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. Authentication provides an additional layer of security on the network beyond the other security features. By default, this authentication is disabled.

Authentication keys can be specified in either plain-text or MD5 form. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

This type of authentication is not supported on RIPv1 networks.

Example: Configuring Route Authentication for RIP

This example shows how to configure authentication for a RIP network.

- [Requirements on page 4116](#)
- [Overview on page 4116](#)
- [Configuration on page 4117](#)
- [Verification on page 4120](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

This example shows MD5 authentication.

Figure 135 on page 4117 shows the topology used in this example.

Figure 135: RIP Authentication Network Topology



"CLI Quick Configuration" on page 4117 shows the configuration for all of the devices in Figure 135 on page 4117. The section "Step-by-Step Procedure" on page 4118 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

Device R2
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2

```

```
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$Lf1Xds2gJDHmoJCu1hKvoJGUjq"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$G.UkP5T39tOz3K87V4oz36/Cu"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RIP authentication:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
```

```
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Require MD5 authentication for RIP route queries received on an interface.

The passwords must match on neighboring RIP routers. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Do not enter the password as shown here. The password shown here is the encrypted password that is displayed in the configuration after the actual password is already configured.

```
[edit protocols rip]
user@R1# set authentication-type md5
user@R1# set authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"
```

6. Configure tracing operations to track authentication.

```
[edit protocols rip traceoptions]
user@R1# set file rip-authentication-messages
user@R1# set flag auth
user@R1# set flag packets
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  traceoptions {
    file rip-authentication-messages;
    flag auth;
```

```

        flag packets;
    }
    authentication-type md5;
    authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"; ## SECRET-DATA
    group rip-group {
        export advertise-routes-through-rip;
        neighbor fe-1/2/0.1;
    }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking for Authentication Failures on page 4120](#)
- [Verifying That MD5 Authentication Is Enabled in RIP Update Packets on page 4121](#)

Checking for Authentication Failures

Purpose Verify that there are no authentication failures.

Action From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          2669         10          2
Triggered Updates Sent      2          0          0
Responses Sent          0          0          0
Bad Messages           0          0          0
RIPv1 Updates Received     0          0          0
RIPv1 Bad Route Entries    0          0          0
RIPv1 Updates Ignored      0          0          0
RIPv2 Updates Received    2675         11          2
RIPv2 Bad Route Entries    0          0          0
RIPv2 Updates Ignored      0          0          0
Authentication Failures      0          0          0
RIP Requests Received      0          0          0
RIP Requests Ignored        0          0          0
none                     0          0          0

```

Meaning The output shows that there are no authentication failures.

Verifying That MD5 Authentication Is Enabled in RIP Update Packets

Purpose Use tracing operations to verify that MD5 authentication is enabled in RIP updates.

Action From operational mode, enter the **show log** command.

```
user@R1> show log rip-authentication-messages | match md5
Feb 15 15:45:13.969462      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:45:43.229867      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:13.174410      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:42.716566      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:47:11.425076      sending msg 0xb9a8c04, 3 rtes (needs MD5)
...
```

Meaning The **(needs MD5)** output shows that all route updates require MD5 authentication.

Enabling Authentication with Plain-Text Passwords (CLI Procedure)

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 316 on page 4121](#).
3. If you are finished configuring the router, commit the configuration.

Table 316: Configuring Simple RIP Authentication

Task	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to simple .	Set the authentication type to simple : set authentication-type simple
Set the authentication key to a simple-text password. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the authentication key to a simple-text password: set authentication-key <i>password</i>

Enabling Authentication with MD5 Authentication (CLI Procedure)

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 317 on page 4122](#).
3. If you are finished configuring the router, commit the configuration.

Table 317: Configuring MD5 RIP Authentication

Task	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to MD5 .	Set the authentication type to md5 : set authentication-type md5
Set the MD5 authentication key (password). The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the MD5 authentication key: set authentication-key <i>password</i>

Related Documentation

- [Example: Configuring RIP on page 4109](#)

Example: Configuring BFD for RIP

- [Understanding BFD for RIP on page 4122](#)
- [Example: Configuring BFD for RIP on page 4123](#)

Understanding BFD for RIP

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 4123](#)
- [Overview on page 4123](#)
- [Configuration on page 4125](#)
- [Verification on page 4127](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  version (1 | automatic);
}
```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.


To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 136 on page 4125 shows the topology used in this example.

Figure 136: RIP BFD Network Topology



"CLI Quick Configuration" on page 4125 shows the configuration for all of the devices in Figure 136 on page 4125. The section "Step-by-Step Procedure" on page 4126 describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip

```

```
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Configure tracing operations to track BFD messages.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
```

```

unit 1 {
    family inet {
        address 10.0.0.1/30;
    }
}

user@R1# show protocols
bfd {
    traceoptions {
        file bfd-trace;
        flag all;
    }
}
rip {
    group rip-group {
        export advertise-routes-through-rip;
        bfd-liveness-detection {
            minimum-interval 600;
        }
        neighbor fe-1/2/0.1;
    }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 4127](#)
- [Checking the BFD Trace File on page 4128](#)

Verifying That the BFD Sessions Are Up

Purpose Make sure that the BFD sessions are operating.

Action From operational mode, enter the **show bfd session** command.

```

user@R1> show bfd session

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

```

1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

```

Meaning The output shows that there are no authentication failures.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

Meaning The output shows the normal functioning of BFD.

Related Documentation

- [Example: Configuring RIP on page 4109](#)
- [Example: Configuring Authentication for RIP Routes on page 4116](#)
- [Example: Configuring Point-to-Multipoint RIP Networks](#)

Example: Configuring BFD Authentication for RIP

- [Understanding BFD Authentication for RIP on page 4128](#)
- [Example: Configuring BFD Authentication for RIP on page 4130](#)

Understanding BFD Authentication for RIP

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over RIP. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and the level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4129](#)
- [Security Authentication Keychains on page 4129](#)
- [Strict Versus Loose Authentication on page 4130](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and

associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for a RIP network.

- [Requirements on page 4130](#)
- [Overview on page 4130](#)
- [Configuration on page 4131](#)
- [Verification on page 4135](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

The devices must be running Junos OS Release 9.6 or later.


Overview

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the RIP protocol.
2. Associate the authentication keychain with the RIP protocol.
3. Configure the related security authentication keychain.

[Figure 137 on page 4131](#) shows the topology used in this example.

Figure 137: RIP BFD Authentication Network Topology



"CLI Quick Configuration" on page 4131 shows the configuration for all of the devices in Figure 137 on page 4131. The section "Step-by-Step Procedure" on page 4132 describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
    "$9$dlV2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
    "2012-2-16.12:00:00 -0800"

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct

```

```
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
  keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD authentication:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```


5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.



NOTE: Nonstop active routing is not supported with **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication algorithm keyed-md5
```

7. Specify the keychain to be used to associate BFD sessions on RIP with the unique security authentication keychain attributes.

The keychain you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication key-chain bfd-rip
```

8. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication loose-check
```

9. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 7.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-rip]
user@R1# set key 53 secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"
user@R1# set key 53 start-time "2012-2-16.12:00:00 -0800"
```

10. Configure tracing operations to track BFD authentication.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}

user@R1# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
rip {
  group rip-group {
    export advertise-routes-through-rip;
    bfd-liveness-detection {
      minimum-interval 600;
    }
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

user@R1# show security
authentication-key-chains {
  key-chain bfd-rip {
    key 53 {
      secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"; ## SECRET-DATA
      start-time "2012-2-16.12:00:00 -0800";
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Authenticated on page 4135](#)
- [Viewing Extensive Information About the BFD Authentication on page 4135](#)
- [Checking the BFD Trace File on page 4136](#)

Verifying That the BFD Sessions Are Authenticated

Purpose Make sure that the BFD sessions are authenticated.

Action From operational mode, enter the **show bfd session detail** command.

```
user@R1> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**
 Session up time 01:39:34
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53

1 sessions, 1 clients
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

Meaning **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive Information About the BFD Authentication

Purpose View the keychain name, the authentication algorithm and mode for each client in the session, and the BFD authentication configuration status.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**
keychain bfd-rip, algo keyed-md5, mode loose
 Session up time 01:46:29
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53
 Min async interval 0.600, min slow interval 1.000
 Adaptive async TX interval 0.600, RX interval 0.600
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 3
 Remote min TX interval 0.600, min RX interval 0.600, multiplier 3
 Local discriminator 225, remote discriminator 226
 Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-rip, algo keyed-md5, mode loose
 Session ID: 0x300501

1 sessions, 1 clients
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

Meaning The output shows the keychain name, the authentication algorithm and mode for the client in the session, and the BFD authentication configuration status.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

Meaning The output shows the normal functioning of BFD.

- Related Documentation**
- [Example: Configuring BFD for RIP on page 4122](#)
 - [Example: Configuring Authentication for RIP Routes on page 4116](#)
 - [Example: Configuring RIP on page 4109](#)

Example: Applying Policies to RIP Routes Imported from Neighbors

- [Understanding RIP Import Policy on page 4136](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4136](#)

Understanding RIP Import Policy

The default RIP import policy is to accept all received RIP routes that pass a sanity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement, and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes.

Example: Applying Policies to RIP Routes Imported from Neighbors

This example shows how to configure an import policy in a RIP network.

- [Requirements on page 4137](#)
- [Overview on page 4137](#)

- [Configuration on page 4137](#)
- [Verification on page 4140](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, Device R1 has an import policy that accepts the 10/8 and 192.168/16 RIP routes and rejects all other RIP routes. This means that the 172.16/16 RIP routes are excluded from Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 138 on page 4137](#) shows the topology used in this example.

Figure 138: RIP Import Policy Network Topology



"[CLI Quick Configuration](#)" on page 4137 shows the configuration for all of the devices in [Figure 138 on page 4137](#). The section "[Step-by-Step Procedure](#)" on page 4138 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip import rip-import
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set policy-options policy-statement rip-import term 1 from protocol rip
set policy-options policy-statement rip-import term 1 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement rip-import term 1 from route-filter 192.168.0.0/16
  orlonger
  
```

```
set policy-options policy-statement rip-import term 1 then accept
set policy-options policy-statement rip-import term 2 then reject
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP import policy:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled.

You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
```

```

user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Configure the import policy.

```

[edit policy-options policy-statement rip-import]
user@R1# set term 1 from protocol rip
user@R1# set term 1 from route-filter 10.0.0.0/8 orlonger
user@R1# set term 1 from route-filter 192.168.0.0/16 orlonger
user@R1# set term 1 then accept
user@R1# set term 2 then reject

```

6. Apply the import policy.

```

[edit protocols rip]
user@R1# set import rip-import

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
}

user@R1# show protocols
rip {
  import rip-import;
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {

```

```
term 1 {
    from protocol [ direct rip ];
    then accept;
}
}
policy-statement rip-import {
    term 1 {
        from {
            protocol rip;
            route-filter 10.0.0.0/8 orlonger;
            route-filter 192.168.0.0/16 orlonger;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Looking at the Routes That Device R2 Is Advertising to Device R1 on page 4140](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4141](#)
- [Checking the Routing Table on page 4141](#)
- [Testing the Import Policy on page 4141](#)

Looking at the Routes That Device R2 Is Advertising to Device R1

Purpose Verify that Device R2 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R2> show route advertising-protocol rip 10.0.0.2
```

```
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.4/30      *[Direct/0] 2d 01:17:44
                  >   via fe-1/2/0.5
172.16.2.2/32    *[Direct/0] 2d 04:09:52
                  >   via lo0.2
172.16.3.3/32    *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
192.168.2.2/32   *[Direct/0] 2d 04:09:52
                  >   via lo0.2
192.168.3.3/32   *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
```

Meaning Device R2 is sending 172.16/16 routes to Device R1.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30          *[RIP/100] 01:06:03, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32      *[RIP/100] 01:06:03, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32      *[RIP/100] 01:06:03, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Meaning The output shows that the 172.16/16 routes are excluded.

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes.

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip

inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30          *[RIP/100] 00:54:34, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32      *[RIP/100] 00:54:34, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32      *[RIP/100] 00:54:34, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32        *[RIP/100] 00:49:00, metric 1
                    MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you delete or deactivate the import policy, the routing table contains the 172.16/16 routes.

Testing the Import Policy

Purpose By using the **test policy** command, monitor the number of rejected prefixes.

Action From operational mode, enter the **test policy rip-import 172.16/16** command.

```
user@R1> test policy rip-import 172.16/16
Policy rip-import: 0 prefix accepted, 1 prefix rejected
```

Meaning The output shows that the policy rejected one prefix.

Related Documentation

- [Example: Configuring RIP on page 4109](#)

Examples: Controlling Traffic with Metrics in a RIP Network

- [Understanding Traffic Control with Metrics in a RIP Network on page 4142](#)
- [Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 4143](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 4144](#)
- [Example: Configuring the Metric Value Added to Imported RIP Routes on page 4146](#)

Understanding Traffic Control with Metrics in a RIP Network

To tune a RIP network and to control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIP provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. In other words, by default, the metric of routes that RIP imports from a neighbor or exports to a neighbor is incremented by 1. These routes include those learned from RIP as well as those learned from other protocols. The metrics are attributes that specify the cost of any route advertised through a host. By increasing or decreasing the metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

The incoming metric modifies the cost of an individual segment when a route across the segment is imported into the routing table. For example, if you set the incoming metric on the segment to **3**, the individual segment cost along the link is changed from 1 to **3**. The increased cost affects all route calculations through that link. Other routes that were previously excluded because of a high hop count might now be included in the router's forwarding table.

The outgoing metric modifies the path cost for all the routes advertised out of a particular interface. Unlike the incoming metric, the outgoing metric modifies the routes that other routers are learning and thereby controls the way they send traffic.

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIP neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIP neighbor configuration.

Example: Controlling Traffic in a RIP Network with an Incoming Metric

This example shows how to control traffic with an incoming metric.

- [Requirements on page 4143](#)
- [Overview on page 4143](#)
- [Configuration on page 4144](#)
- [Verification on page 4144](#)


Requirements

Before you begin, define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through the RIP routing exchanges. See “[Example: Configuring a Basic RIP Network](#)” on page 4110.

Overview

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces as shown in [Figure 139 on page 4143](#). Because the route through Router B and the route through Router C have the same number of hops, both routes are imported into the forwarding table. However, because the T3 link from Router B to Router D has a higher bandwidth than the T1 link from Router C to Router D, you want traffic to flow from Router A through Router B to Router D.

Figure 139: Controlling Traffic in a RIP Network with the Incoming Metric



To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes only the routes in Router A's routing table, and affects only how Router A sends traffic to Router D. Router D's route selection is based on its own routing table, which, by default, includes no adjusted metric values.

In the example, Router C receives a route advertisement from Router D and readvertises the route to Router A. When Router A receives the route, it applies the incoming metric on the interface. Instead of incrementing the metric by 1 (the default), Router A increments it by 3 (the configured incoming metric), giving the route from Router A to Router D

through Router C a total path metric of 4. Because the route through Router B has a metric of 2, it becomes the preferred route for all traffic from Router A to Router D.

This example uses a RIP group called **alpha 1** on interface **ge3-0/0/0**.

Configuration

Step-by-Step Procedure

To control traffic with an incoming metric:

1. Enable RIP on the interface.

```
[edit protocols rip]  
user@host# set group alpha1 neighbor ge-0/0/0
```
2. Set the incoming metric.

```
[edit protocols rip]  
user@host# set metric-in 3
```
3. If you are done configuring the device, commit the configuration.

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

Verification

To verify that the configuration is working properly, enter the **show route protocols rip** command.

Example: Controlling Traffic in a RIP Network with an Outgoing Metric

This example shows how to control traffic with an outgoing metric.

- [Requirements on page 4144](#)
- [Overview on page 4144](#)
- [Configuration on page 4145](#)
- [Verification on page 4145](#)

Requirements

Before you begin:


- Define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 4110](#).
- Control traffic with an incoming metric. See [“Example: Controlling Traffic in a RIP Network with an Incoming Metric” on page 4143](#).

Overview

In this example, each route from Router A to Router D has two hops as shown in [Figure 140 on page 4145](#). However, because the link from Router A to Router B in the RIP group has a higher bandwidth than the link from Router A to Router C in RIP group Alpha 1, you want traffic from Router D to Router A to flow through Router B. To control the

way Router D sends traffic to Router A, you can alter the routes that Router D receives by configuring the outgoing metric on Router A's interfaces in the Alpha 1 RIP group.

Figure 140: Controlling Traffic in a RIP Network with the Outgoing Metric



If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from Router A through Router C as 4. In contrast, the unchanged default total path metric to Router A through Router B in the RIP group is 2. The fact that Router A's interfaces belong to two different RIP groups allows you to configure two different outgoing metrics on its interfaces, because you configure path metrics at the group level.

By configuring the outgoing metric, you control the way Router A sends traffic to Router D. By configuring the outgoing metric on the same router, you control the way Router D sends traffic to Router A.

This example uses an outgoing metric of 3.

Configuration

Step-by-Step Procedure

To control traffic with an outgoing metric:

1. Set the outgoing metric.


```
[edit protocols rip group alpha1]
user@host# set metric-out 3
```
2. If you are done configuring the device, commit the configuration.


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

Verification

To verify that the configuration is working properly, enter the **show protocols rip** command.

Example: Configuring the Metric Value Added to Imported RIP Routes

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- [Requirements on page 4146](#)
- [Overview on page 4146](#)
- [Configuration on page 4146](#)
- [Verification on page 4149](#)

Requirements


No special configuration beyond device initialization is required before configuring this example.

Overview

Normally, when multiple routes are available, RIP selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

[Figure 141 on page 4146](#) shows the topology used in this example.

Figure 141: RIP Incoming Metrics Network Topology



Device R1 has two potential paths to reach 172.16.2.2/32. The default behavior is to send traffic out the 0.1/30 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. “[CLI Quick Configuration](#)” on [page 4146](#) shows the configuration for all of the devices in [Figure 141 on page 4146](#). The section “[Step-by-Step Procedure](#)” on [page 4147](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces ge-1/2/1 unit 10 description to-R3
```

```

set interfaces ge-1/2/1 unit 10 family inet address 10.0.0.10/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group primary export advertise-routes-through-rip
set protocols rip group primary neighbor ge-1/2/1.10
set protocols rip group secondary export advertise-routes-through-rip
set protocols rip group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor ge-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP metrics:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet address 10.0.0.10/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32

```

```
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

```
[edit protocols rip]
user@R1# set group primary neighbor ge-1/2/1.10
user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip]
user@R1# set group primary export advertise-routes-through-rip
user@R1# set group secondary export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
ge-1/2/1 {
  unit 10 {
    description to-R3;
    family inet {
      address 10.0.0.10/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
```



```

    }
  }
user@R1# show protocols
rip {
  group primary {
    export advertise-routes-through-rip;
    neighbor ge-1/2/1.10;
  }
  group secondary {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      metric-in 4;
    }
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Route Is Active on page 4149](#)
- [Removing the metric-in Statement on page 4149](#)

Verifying That the Expected Route Is Active

Purpose Make sure that to reach 172.16.2.2/32, Device R1 uses the path through Device R3.

Action From operational mode, enter the **show route 172.16.2.2** command.

```

user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      *[RIP/100] 00:15:46, metric 3, tag 0
                  > to 10.0.0.9 via ge-1/2/1.10

```

Meaning The **to 10.0.0.9 via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. The metric for this route is 3.

Removing the metric-in Statement

Purpose Delete or deactivate the **metric-in** statement to see what happens to the 172.16.2.2/32 route.

Action 1. From configuration mode, deactivate the **metric-in** statement.

```
[edit protocols rip group secondary neighbor fe-1/2/0.1]
user@R1# deactivate metric-in
user@R1# commit
```

2. From operational mode, enter the **show route 172.16.2.2** command.

```
user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      * [RIP/100] 00:00:06, metric 2, tag 0
> to 10.0.0.2 via fe-1/2/0.1
```

Meaning The **to 10.0.0.2 via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 172.16.2.2/32. The metric for this route is 2.

Related Documentation

- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4136](#)

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

- [Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4150](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4150](#)

Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets

RIP version 1 (RIPv1) and RIP version 2 (RIPv2) can run simultaneously. This might make sense when you are migrating a RIPv1 network to a RIPv2 network. This also allows interoperation with a device that supports RIPv1 but not RIPv2.

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets. You can configure this behavior by including the [send](#) and [receive](#) statements in the RIP configuration.

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

This example shows how to configure whether the RIP update messages conform to RIP version 1 (RIPv1) only, to RIP version 2 (RIPv2) only, or to both versions. You can also disable the sending or receiving of update messages.

- [Requirements on page 4150](#)
- [Overview on page 4150](#)
- [Configuration on page 4151](#)
- [Verification on page 4153](#)

Requirements


No special configuration beyond device initialization is required before configuring this example.

Overview

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets.

Figure 142 on page 4151 shows the topology used in this example.

Figure 142: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology



In this example, Device R1 is configured to receive only RIPv2 packets.

“CLI Quick Configuration” on page 4151 shows the configuration for all of the devices in Figure 142 on page 4151. The section “Step-by-Step Procedure” on page 4152 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

- Device R1**
- ```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1 receive version-2
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R2**
- ```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R3**
- ```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip

```

```
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP packet versions that can be received:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **receive version-2** setting causes this interface to accept only RIPv2 packets.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1 receive version-2
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
```

```

 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1 {
 receive version-2;
 }
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

#### Verifying That the Receive Mode Is Set to RIPv2 Only

**Purpose** Make sure that the interfacing Device R2 is configured to receive only RIPv2 packets, instead of both RIPv1 and RIPv2 packets.

**Action** From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

| Neighbor   | Local<br>State | Source<br>Address | Destination<br>Address | Send<br>Mode | Receive<br>Mode | In<br>Met |
|------------|----------------|-------------------|------------------------|--------------|-----------------|-----------|
| fe-1/2/0.1 | Up             | 10.0.0.1          | 224.0.0.9              | mcast        | v2 only         | 1         |

**Meaning** In the output, the **Receive Mode** field displays **v2 only**. The default **Receive Mode** is **both**.

**Related Documentation**

- [Example: Configuring RIP on page 4109](#)

## Example: Redistributing Routes Among RIP Instances

- [Understanding Route Redistribution Among RIP instances on page 4154](#)
- [Example: Redistributing Routes Between Two RIP Instances on page 4155](#)

### Understanding Route Redistribution Among RIP instances

---

You can redistribute routes among RIP processes. Another way to say this is to export RIP routes from one RIP instance to other RIP instances.

In Junos OS, route redistribution among routing instances is accomplished by using routing table groups, also called RIB groups. Routing table groups allow you to import and export routes from a protocol within one routing table into another routing table.



**NOTE:** In contrast, the policy-based import and export functions allow you import and export routes between different protocols within the same routing table.

---

Consider the following partial example:

```
protocols {
 rip {
 rib-group inet-to-voice;
 }
}
routing-instances {
 voice {
 protocols {
 rip {
 rib-group voice-to-inet;
 }
 }
 }
}
routing-options {
 rib-groups {
 inet-to-voice {
 import-rib [inet.0 voice.inet.0];
 }
 voice-to-inet {
 import-rib [voice.inet.0 inet.0];
 }
 }
}
```

The way to read the **import-rib** statement is as follows. Take the routes from the protocol (RIP, in this case), and import them into the primary (or local) routing table and also into any other routing tables listed after this. The primary routing table is the routing table where the routing table group is being used. That would be either **inet.0** if used in the main routing instance or **voice.inet.0** if used within the routing instance. In the **inet-to-voice** routing table group, **inet.0** is listed first because this routing table group is used in the

main routing instance. In the **voice-to-inet** routing table group, **voice.inet.0** is listed first because this routing table group is used in the voice routing instance.

### Example: Redistributing Routes Between Two RIP Instances

This example shows how to configure a RIP routing instance and control the redistribution of RIP routes between the routing instance and the master instance.

- [Requirements on page 4155](#)
- [Overview on page 4155](#)
- [Configuration on page 4155](#)
- [Verification on page 4159](#)

#### Requirements


No special configuration beyond device initialization is required before configuring this example.

#### Overview

When you create a routing instance called voice, Junos OS creates a routing table called **voice.inet.0**. The example shows how to install routes learned through the master RIP instance into the **voice.inet.0** routing table. The example also shows how to install routes learned through the voice routing instance into **inet.0**. This is done by configuring routing table groups. RIP routes are installed into each routing table that belongs to a routing table group.

[Figure 143 on page 4155](#) shows the topology used in this example.

**Figure 143: Redistributing Routes Between RIP Instances Network Topology**



[“CLI Quick Configuration” on page 4155](#) shows the configuration for all of the devices in [Figure 143 on page 4155](#). The section [“Step-by-Step Procedure” on page 4156](#) describes the steps on Device R2.

#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
```

```
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip rib-group inet-to-voice
set protocols rip group to-R3 export advertise-routes-through-rip
set protocols rip group to-R3 neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set routing-instances voice protocols rip group to-R1 export advertise-routes-through-rip
set routing-instances voice interface fe-1/2/0.2
set routing-instances voice protocols rip rib-group voice-to-inet
set routing-instances voice protocols rip group to-R1 neighbor fe-1/2/0.2
set routing-options rib-groups inet-to-voice import-rib inet.0
set routing-options rib-groups inet-to-voice import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib inet.0
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To redistribute RIP routes between routing instances:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
```



```
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Create the routing instance, and add one or more interfaces to the routing instance.

```
[edit routing-instances voice]
user@R2# set interface fe-1/2/0.2
```

3. Create the RIP groups and add the interfaces.

```
[edit protocols rip group to-R3]
user@R2# set neighbor fe-1/2/1.5

[edit routing-instances voice protocols rip group to-R1]
user@R2# set neighbor fe-1/2/0.2
```

4. Create the routing table groups.

```
[edit routing-options rib-groups]
user@R2# set inet-to-voice import-rib inet.0
user@R2# set inet-to-voice import-rib voice.inet.0

user@R2# set voice-to-inet import-rib voice.inet.0
user@R2# set voice-to-inet import-rib inet.0
```

5. Apply the routing table groups.

```
[edit protocols rip]
user@R2# set rib-group inet-to-voice

[edit routing-instances voice protocols rip]
user@R2# set rib-group voice-to-inet
```

6. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

7. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group to-R3]
user@R2# set export advertise-routes-through-rip

[edit routing-instances voice protocols rip group to-R1]
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
 unit 2 {
 family inet {
```

```
 address 10.0.0.2/30;
 }
}
fe-1/2/1 {
 unit 5 {
 family inet {
 address 10.0.0.5/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.2/32;
 address 172.16.2.2/32;
 }
 }
}

user@R2# show protocols
rip {
 rib-group inet-to-voice;
 group to-R3 {
 export advertise-routes-through-rip;
 neighbor fe-1/2/1.5;
 }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

user@R2# show routing-instances
voice {
 interface fe-1/2/0.2;
 protocols {
 rip {
 rib-group voice-to-inet;
 group to-R1 {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.2;
 }
 }
 }
}

user@R2# show routing-options
rib-groups {
 inet-to-voice {
 import-rib [inet.0 voice.inet.0];
 }
 voice-to-inet {
```

```

import-rib [voice.inet.0 inet.0];
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

### Checking the Routing Tables

**Purpose** Make sure that the routing tables contain the expected routes.

**Action** From operational mode, enter the **show route protocol rip** command.

```

user@R2> show route protocol rip
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32 * [RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32 * [RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32 * [RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32 * [RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32 * [RIP/100] 01:44:13, metric 1
 MultiRecv

voice.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32 * [RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32 * [RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32 * [RIP/100] 02:06:03, metric 2, tag 0
 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32 * [RIP/100] 01:58:14, metric 2, tag 0
 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32 * [RIP/100] 01:44:13, metric 1
 MultiRecv

```

**Meaning** The output shows that both routing tables contain all of the RIP routes.

**Related Documentation**

- [Example: Configuring RIP on page 4109](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4136](#)

## Example: Configuring RIP Timers

- [Understanding RIP Timers on page 4160](#)
- [Example: Configuring RIP Timers on page 4161](#)

## Understanding RIP Timers

---

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

**update-interval** *seconds*;

*seconds* can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

**route-timeout** *seconds*;

*seconds* can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the **holddown** statement:

**holddown** *seconds*;

*seconds* can be a value from 10 through 180. The default value is 120 seconds.



**NOTE:** In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

---

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. The hold-down timer must be greater than the route timeout. Normally, the default values are best left in effect for standard operations.

### Example: Configuring RIP Timers

This example shows how to configure the RIP update interval and how to monitor the impact of the change.

- [Requirements on page 4161](#)
- [Overview on page 4161](#)
- [Configuration on page 4161](#)
- [Verification on page 4164](#)

#### Requirements

No special configuration beyond device initialization is required before configuring this example.

#### Overview


In this example, Device R2 has an update interval of 60 seconds for its neighbor, Device R1, and an update interval of 10 seconds for its neighbor, Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 144 on page 4161](#) shows the topology used in this example.

**Figure 144: RIP Timers Network Topology**



“CLI Quick Configuration” on [page 4161](#) shows the configuration for all of the devices in [Figure 144 on page 4161](#). The section “Step-by-Step Procedure” on [page 4162](#) describes the steps on Device R2.

#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**      **set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30**  
**set interfaces lo0 unit 1 family inet address 172.16.0.1/32**

```
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2 update-interval 60
set protocols rip group rip-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30
```

```
user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30
```

```
user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Configure different update intervals for the two RIP neighbors.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R2# set neighbor fe-1/2/0.2 update-interval 60
user@R2# set neighbor fe-1/2/1.5 update-interval 10
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
 unit 2 {
 family inet {
 address 10.0.0.2/30;
 }
 }
}
fe-1/2/1 {
 unit 5 {
 family inet {
 address 10.0.0.5/30;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.2.2/32;
 address 172.16.2.2/32;
 }
 }
}

user@R2# show protocols
rip {
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.2 {
 update-interval 60;
 }
 }
}
```

```

 neighbor fe-1/2/1.5 {
 update-interval 10;
 }
 }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the RIP Updates Sent by Device R2 on page 4164](#)
- [Checking the RIP Updates Received by Device R2 on page 4165](#)
- [Checking the RIP Updates Received by Device R3 on page 4165](#)

### Checking the RIP Updates Sent by Device R2

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R2> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```

 rts learned rts held down rqsts dropped resps dropped
 4 2 0 0

```

```
fe-1/2/0.2: 2 routes learned; 5 routes advertised; timeout 180s; update interval 60s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| Updates Sent            | 123   | 5          | 1           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 244   | 10         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

```
fe-1/2/1.5: 2 routes learned; 5 routes advertised; timeout 180s; update interval 10s
```

| Counter | Total | Last 5 min | Last minute |
|---------|-------|------------|-------------|
|---------|-------|------------|-------------|



|                         |     |    |   |
|-------------------------|-----|----|---|
| Updates Sent            | 734 | 32 | 6 |
| Triggered Updates Sent  | 0   | 0  | 0 |
| Responses Sent          | 0   | 0  | 0 |
| Bad Messages            | 0   | 0  | 0 |
| RIPv1 Updates Received  | 0   | 0  | 0 |
| RIPv1 Bad Route Entries | 0   | 0  | 0 |
| RIPv1 Updates Ignored   | 0   | 0  | 0 |
| RIPv2 Updates Received  | 245 | 11 | 2 |
| RIPv2 Bad Route Entries | 0   | 0  | 0 |
| RIPv2 Updates Ignored   | 0   | 0  | 0 |
| Authentication Failures | 0   | 0  | 0 |
| RIP Requests Received   | 0   | 0  | 0 |
| RIP Requests Ignored    | 0   | 0  | 0 |
| none                    | 0   | 0  | 0 |

**Meaning** The **update interval** field shows that the interval is 60 seconds for Neighbor R1 and 10 seconds for Neighbor R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

### *Checking the RIP Updates Received by Device R2*

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0
```

```
fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 312   | 10         | 2           |
| Triggered Updates Sent  | 2     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 181   | 5          | 1           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 1     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

### *Checking the RIP Updates Received by Device R3*

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R3> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped
 5 0 0 0

fe-1/2/0.6: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter Total Last 5 min Last minute

Updates Sent 314 11 2
Triggered Updates Sent 1 0 0
Responses Sent 0 0 0
Bad Messages 0 0 0
RIPv1 Updates Received 0 0 0
RIPv1 Bad Route Entries 0 0 0
RIPv1 Updates Ignored 0 0 0
RIPv2 Updates Received 827 31 6
RIPv2 Bad Route Entries 0 0 0
RIPv2 Updates Ignored 0 0 0
Authentication Failures 0 0 0
RIP Requests Received 0 0 0
RIP Requests Ignored 0 0 0
none 0 0 0

```

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

**Related Documentation**

- [Example: Configuring RIP on page 4109](#)
- [Example: Configuring RIP Demand Circuits](#)

## Example: Tracing RIP Protocol Traffic

- [Understanding RIP Trace Operations on page 4166](#)
- [Example: Tracing RIP Protocol Traffic on page 4167](#)

### Understanding RIP Trace Operations

You can trace various types of RIP protocol traffic to help debug RIP protocol issues.

To trace RIP protocol traffic, include the **traceoptions** statement at the **[edit protocols rip]** hierarchy level:

```

traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}

```

You can specify the following RIP protocol-specific trace options using the **flag** statement:

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing

- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop active routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



**NOTE:** Use the **detail** flag modifier with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIP protocol using the **traceoptions flag** statement included at the **[edit protocols rip]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



**NOTE:** Use the trace flag **all** with caution because this may cause the CPU to become very busy.

### Example: Tracing RIP Protocol Traffic

This example shows how to trace RIP protocol operations.

- [Requirements on page 4168](#)
- [Overview on page 4168](#)

- [Configuration on page 4168](#)
- [Verification on page 4170](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 145 on page 4168](#) shows the topology used in this example.

**Figure 145: RIP Trace Operations Network Topology**



"CLI Quick Configuration" on [page 4168](#) shows the configuration for all of the devices in [Figure 145 on page 4168](#). The section "Step-by-Step Procedure" on [page 4169](#) describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip traceoptions file rip-trace-file
set protocols rip traceoptions flag route
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

#### Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
```

```

set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

```

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Configure the RIP group, and add the interface to the group.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Configure RIP tracing operations.

```

[edit protocols rip traceoptions]
user@R1# set file rip-trace-file
user@R1# set flag route

```

4. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

5. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 traceoptions {
 file rip-trace-file;
 flag route;
 }
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

**Checking the Log File**

**Purpose** Make sure that the RIP route updates are logged in the configured log file.

**Action** 1. Deactivate the extra loopback interface address on Device R3.

```
[edit interfaces lo0 unit 3 family inet]
user@R3# deactivate address 172.16.3.3/32
user@R3# commit
```

2. From operational mode on Device R1, enter the **show log rip-trace-file** command with the **| match 172.16.3.3** option.

```
user@R1> show log rip-trace-file | match 172.16.3.3
Mar 1 11:39:53.975192 Setting RIPv2 rtbit on route 172.16.3.3/32, tsi =
0xbb69228
Mar 1 11:39:59.847118 172.16.3.3/32: metric-in: 16, change: 3 -> 16; # gw:
1, pkt_upd_src 10.0.0.2, inx: 0, rte_upd_src 10.0.0.2
Mar 1 11:39:59.847568 CHANGE 172.16.3.3/32 nhid 591 gw 10.0.0.2
RIP pref 100/0 metric 3/0 fe-1/2/0.1 <Delete Int>
Mar 1 11:39:59.847629 Best route to 172.16.3.3/32 got deleted. Doing route calculation
on the stored rte-info
```

**Meaning** The output shows that the route to 172.16.3.3/32 was deleted.

**Related Documentation**

- [Example: Configuring RIP on page 4109](#)

## RIP Configuration Statements

---


- [any-sender on page 4172](#)
- [authentication-key \(Protocols RIP\) on page 4173](#)
- [authentication-type \(Protocols RIP\) on page 4174](#)
- [bfd-liveness-detection \(Protocols RIP\) on page 4175](#)
- [check-zero on page 4178](#)
- [export \(Protocols RIP\) on page 4179](#)
- [group \(Protocols RIP\) on page 4180](#)
- [holddown \(Protocols RIP\) on page 4182](#)
- [import \(Protocols RIP\) on page 4183](#)
- [message-size on page 4184](#)
- [metric-in \(Protocols RIP\) on page 4185](#)
- [metric-out \(Protocols RIP\) on page 4186](#)
- [neighbor \(Protocols RIP\) on page 4187](#)
- [preference \(Protocols RIP\) on page 4188](#)
- [receive \(Protocols RIP\) on page 4189](#)
- [rib-group \(Protocols RIP\) on page 4190](#)
- [rip on page 4190](#)

- [route-timeout \(Protocols RIP\) on page 4191](#)
- [send \(Protocols RIP\) on page 4192](#)
- [traceoptions \(Protocols RIP\) on page 4193](#)
- [update-interval \(Protocols RIP\) on page 4196](#)

---

## any-sender

---

|                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                    | <code>any-sender;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                           | <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b></code><br><code><i>neighbor-name</i>]</code> |
| <b>Release Information</b>                                                                                                                                                                       | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                                               | <p>Disable strict sender address checks.</p> <p>If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation might cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.</p>                                                                                                                                                                                                 |
| <div> <b>NOTE:</b> The <code>any-sender</code> statement is supported only for peer-to-peer interfaces.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                                                                                                                                                                  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                     | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4109</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## authentication-key (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key password;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Require authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b><i>password</i></b> —Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4116</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## authentication-type (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the type of authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, RIP authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>type</b> —Authentication type: <ul style="list-style-type: none"><li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li><li>• <b>none</b>—Disable authentication. If <b>none</b> is configured, the configured authentication key is ignored.</li><li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.</li></ul>                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4116</a></li><li>• <a href="#">authentication-key on page 4173</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## bfd-liveness-detection (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {     authentication {         algorithm <i>algorithm-name</i>;         key-chain <i>key-chain-name</i>;         loose-check;     }     detection-time {         threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;     }     version (1   automatic); } </pre>                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/> rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/> [edit protocols rip <b>group</b> <i>group-name</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b><br/> <i>neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Options <b>detection-time threshold</b> and <b>transmit-interval threshold</b> introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>Option <b>no-adaptation</b> introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b> —Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, or <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the [edit security] hierarchy level.</p>                                                                                                                                                                                                                           |

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |


- Related Documentation**
- [Example: Configuring BFD for RIP on page 4123](#)
  - [Example: Configuring BFD Authentication for RIP on page 4130](#)

## check-zero

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (check-zero   no-check-zero);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit protocols <i>rip</i>], [edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Some of the reserved fields in RIP version 1 packets must be zero, whereas in RIP version 2 packets, most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</p> <p>If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets even though they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</p> <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"><li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li><li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li></ul> |
| <b>Default</b>                  | check-zero                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4109</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## export (Protocols RIP)

|                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                             | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                    | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</p>                                                                      |
| <b>Release Information</b>                                                                                                                                                                                | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                                        | <p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p> |
| <div>  <b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                            | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                           | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4109</a></li> <li>• <a href="#">import on page 4183</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                |

## group (Protocols RIP)

---

**Syntax** `group group-name {`  
    `bfd-liveness-detection {`  
        `authentication {`  
            `algorithm algorithm-name;`  
            `key-chain key-chain-name;`  
            `loose-check;`  
        `}`  
        `detection-time {`  
            `threshold milliseconds;`  
        `}`  
        `minimum-interval milliseconds;`  
        `minimum-receive-interval milliseconds;`  
        `transmit-interval {`  
            `threshold milliseconds;`  
            `minimum-interval milliseconds;`  
        `}`  
        `multiplier number;`  
        `version (0 | 1 | automatic);`  
    `}`  
    `demand-circuit;`  
    `export policy;`  
    `max-retrans-time seconds;`  
    `metric-out metric;`  
    `preference number;`  
    `route-timeout seconds;`  
    `update-interval seconds;`  
    `neighbor neighbor-name {`  
        `authentication-key password;`  
        `authentication-type type;`  
        `bfd-liveness-detection {`  
            `authentication {`  
                `algorithm algorithm-name;`  
                `key-chain key-chain-name;`  
                `loose-check;`  
            `}`  
            `detection-time {`  
                `threshold milliseconds;`  
            `}`  
            `minimum-interval milliseconds;`  
            `minimum-receive-interval milliseconds;`  
            `transmit-interval {`  
                `threshold milliseconds;`  
                `minimum-interval milliseconds;`  
            `}`  
            `multiplier number;`  
            `version (0 | 1 | automatic);`  
        `}`  
    `(check-zero | no-check-zero);`  
    `demand-circuit;`  
    `import policy-name;`  
    `max-retrans-time seconds;`  
    `message-size number;`



```

 metric-in metric;
 metric-out metric;
 receive receive-options;
 route-timeout seconds;
 send send-options;
 update-interval seconds;
 }
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.                                                                                  |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4109</a></li> </ul>                                                                                                                                                                                                                                               |

## holddown (Protocols RIP)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>holddown seconds;</code>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ]                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure how long the expired route is retained in the routing table before being removed.</p> <p>When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.</p> |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 180 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 4161</a></li><li>• <a href="#">RIP Demand Circuits Overview</a></li></ul>                                                                                                                                                                                                                                    |

## import (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported by the local routing device from neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Applying Policies to RIP Routes Imported from Neighbors on page 4136</a></li> <li>• <a href="#">Routing Policy Configuration Guide</a></li> <li>• <a href="#">export on page 4179</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## message-size

---

|                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                           | <code>message-size <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                  | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                              | Statement introduced before Junos OS Release 7.4.<br>Statement for SRX Series devices introduced in Junos OS Release 9.5.<br>Statement for J Series platform introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                                                                                                                                                                                                                                                                      | Specify the number of route entries to be included in every RIP update message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <hr/>                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <div> <b>TIP:</b> To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message. Do not change the default number of route entries in a RIP update message.</div> <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                                                                                                                                                                                                                                                                          | <b><i>number</i></b> —Number of route entries per update message.<br><b>Range:</b> 25 through 255 entries<br><b>Default:</b> 25 entries                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                         | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4109</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## metric-in (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Metric Value Added to Imported RIP Routes on page 4146</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## metric-out (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.</p> <p>If you have included the <b>export</b> statement, RIP exports routes it has learned to the neighbors configured by including the <b>neighbor</b> statement.</p> <p>The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a <b>metric-in</b> value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the <b>metric-out</b> statement.</p> <p>The metric for a route can be modified with an export policy. That metric is seen when the route is exported to the next hop.</p> <p>To increase the metric for routes advertised outside a group, include the <b>metric-out</b> statement.</p> |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Examples: Controlling Traffic with Metrics in a RIP Network on page 414250</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## neighbor (Protocols RIP)

**Syntax** `neighbor neighbor-name {`  
     `authentication-key password;`  
     `authentication-type type;`  
     `bfd-liveness-detection {`  
         `authentication {`  
             `algorithm algorithm-name;`  
             `key-chain key-chain-name;`  
             `loose-check;`  
         `}`  
         `detection-time {`  
             `threshold milliseconds;`  
         `}`  
     `minimum-interval milliseconds;`  
     `minimum-receive-interval milliseconds;`  
     `transmit-interval {`  
         `threshold milliseconds;`  
         `minimum-interval milliseconds;`  
     `}`  
     `multiplier number;`  
     `version (0 | 1 | automatic);`  
     `}`  
     `(check-zero | no-check-zero);`  
     `demand-circuit;`  
     `import policy-name;`  
     `max-retrans-time seconds;`  
     `message-size number;`  
     `metric-in metric;`  
     `metric-out metric;`  
     `receive receive-options;`  
     `route-timeout seconds;`  
     `send send-options;`  
     `update-interval seconds;`  
     `}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 rip **group** *group-name*],  
 [edit protocols rip **group** *group-name*],  
 [edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

**Options** *neighbor-name*—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring RIP on page 4109](#)

---

## preference (Protocols RIP)

---

**Syntax** `preference preference;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
rip **group** *group-name*],  
[edit protocols rip **group** *group-name*],  
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

**Options** *preference*—Preference value. A lower value indicates a more preferred route.  
**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )  
**Default:** 100

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Route Preferences Overview](#)



## receive (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>receive receive-options;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure RIP receive options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>receive-options</i>—One of the following:</p> <ul style="list-style-type: none"> <li>• <b>both</b>—Accept both RIP version 1 and version 2 packets.</li> <li>• <b>none</b>—Do not receive RIP packets.</li> <li>• <b>version-1</b>—Accept only RIP version 1 packets.</li> <li>• <b>version-2</b>—Accept only RIP version 2 packets.</li> </ul> <p><b>Default:</b> <b>both</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4150</a></li> <li>• <a href="#">send on page 4192</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## rib-group (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Install RIP routes into multiple routing tables by configuring a routing table group.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Redistributing Routes Between Two RIP Instances on page 4155</a></li></ul>                                                                                                                                                                                                                 |

## rip

---

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rip {...}</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable RIP routing on the routing device.                                                                                                                                                                                                                           |
| <b>Default</b>                  | RIP is disabled on the routing device.                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4109</a></li></ul>                                                                                                                                                             |

## route-timeout (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip <a href="#">group group-name</a>],</p> <p>[edit protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 30 through 360 seconds</p> <p><b>Default:</b> 180 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Timers on page 4161</a></li> <li>• <a href="#">RIP Demand Circuits Overview</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## send (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>send <i>send-options</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure RIP send options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>send-options</i> —One of the following: <ul style="list-style-type: none"><li>• <b>broadcast</b>—Broadcast RIP version 2 packets (RIP version 1 compatible).</li><li>• <b>multicast</b>—Multicast RIP version 2 packets. This is the default.</li><li>• <b>none</b>—Do not send RIP updates.</li><li>• <b>version-1</b>—Broadcast RIP version 1 packets.</li></ul> <b>Default:</b> multicast                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4150</a></li><li>• <a href="#">receive on page 4189</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## traceoptions (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>         | Set RIP protocol-level tracing options.                                                                                                                                                                                                                                                                                                                 |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

**Default** The default RIP protocol-level trace options are inherited from the global **traceoptions** statement.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

### RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets

- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the ***files*** option.

**Syntax:** ***xk*** to specify KB, ***xm*** to specify MB, or ***xg*** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                              |                                                                                                                        |
|------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                   |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Example: Tracing RIP Protocol Traffic on page 4167</a></li> </ul> |

## update-interval (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>update-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>seconds</i></b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 60 seconds<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 4161</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



## CHAPTER 47

# Administration

- [Routine Monitoring on page 4197](#)
- [RIP Operational Commands on page 4197](#)

## Routine Monitoring

---

- [Monitoring RIP Routing Information on page 4197](#)

### Monitoring RIP Routing Information

**Purpose** Use the monitoring functionality to monitor RIP routing on routing devices.

**Action** To view RIP routing information in the CLI, enter the following CLI commands:

- **show rip statistics**
- **show rip neighbor**

**Related Documentation**

- [show rip neighbor on page 4202](#)
- [show rip statistics on page 4204](#)

## RIP Operational Commands

---

- [clear rip general-statistics](#)
- [clear rip statistics](#)
- [show rip general-statistics](#)
- [show rip neighbor](#)
- [show rip statistics](#)

## clear rip general-statistics

---

|                                                   |                                                                                                                                                                                                         |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4198</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4198</a>                                                                                          |
| <b>Syntax</b>                                     | clear rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                    |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip general-statistics                                                                                                                                                                            |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                |
| <b>Description</b>                                | Clear RIP general statistics.                                                                                                                                                                           |
| <b>Options</b>                                    | <b>none</b> —Clear RIP general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show rip general-statistics on page 4200</a></li></ul>                                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">clear rip general-statistics on page 4198</a>                                                                                                                                               |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                   |

## Sample Output

### clear rip general-statistics

```
user@host> clear rip general-statistics
```

## clear rip statistics

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4199</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4199</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                     | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><logical-system (all   <i>logical-system-name</i> )><br><neighbor><br><peer (all   <i>address</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><neighbor>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                | Clear RIP statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                    | <p><b>none</b>—Reset RIP counters for all neighbors for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Clear RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">show rip statistics on page 4204</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                      | <a href="#">clear rip statistics on page 4199</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear rip statistics

```
user@host> clear rip statistics
```

## show rip general-statistics

|                                                   |                                                                                                                                                                                          |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4200</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4200</a>                                                                           |
| <b>Syntax</b>                                     | show rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                      |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show rip general-statistics                                                                                                                                                              |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>                                | Display brief RIP statistics.                                                                                                                                                            |
| <b>Options</b>                                    | none—Display brief RIP statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                     |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><a href="#">clear rip general-statistics on page 4198</a></li> </ul>                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">show rip general-statistics on page 4200</a>                                                                                                                                 |
| <b>Output Fields</b>                              | Table 318 on page 4200 lists the output fields for the <b>show rip general-statistics</b> command. Output fields are listed in the approximate order in which they appear.               |

**Table 318: show rip general-statistics Output Fields**

| Field Name  | Field Description                                      |
|-------------|--------------------------------------------------------|
| bad msgs    | Number of invalid messages received.                   |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIP.                |
| max memory  | Most memory used by RIP.                               |

## Sample Output

### show rip general-statistics

```
user@host> show rip general-statistics
```

```
RIPv2 I/O info:
 bad msgs : 0
 no recv intf : 0
 curr memory : 0
 max memory : 0
```

## show rip neighbor

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4202</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4202</a>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                                     | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;<i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                | Display information about RIP neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                    | <p><b>none</b>—Display information about all RIP neighbors for all instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                      | <a href="#">show rip neighbor on page 4203</a><br><a href="#">show rip neighbor (With Demand Circuits Configured) on page 4203</a>                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>                              | <a href="#">Table 319 on page 4203</a> lists the output fields for the <b>show rip neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                     |

Table 319: show rip neighbor Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor</b>            | Name of the RIP neighbor.<br><br><b>NOTE:</b> Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.<br><br>If you configure demand circuits at the <b>[edit protocols rip group group-name neighbor neighbor-name]</b> hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the <b>[edit protocols rip group group-name]</b> hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits. |
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Source Address</b>      | Address of the port on the local router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Destination Address</b> | Address of the port on the remote router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Send Mode</b>           | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , or <b>version 1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Receive Mode</b>        | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>version 2</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show rip neighbor

```

user@host> show rip neighbor
Neighbor Local Source Destination Send Receive In
----- -
ge-2/3/0.0 Up 192.168.9.105 192.168.9.107 bcast both 1
at-5/1/1.42 Dn (null) (null) mcast v2 only 3
at-5/1/0.42 Dn (null) (null) mcast both 3
at-5/1/0.0 Up 20.0.0.1 224.0.0.9 mcast both 3
so-0/0/0.0 Up 192.168.9.97 224.0.0.9 mcast both 3

```

### show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor Local Source Destination Send Receive In
----- -
so-0/1/0.0(DC) Up 10.10.10.2 224.0.0.9 mcast both 1
so-0/2/0.0(DC) Up 13.13.13.2 224.0.0.9 mcast both 1

```

## show rip statistics

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4204</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4204</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                     | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>name</i>&gt; &lt;peer (all   <i>address</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;<i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                    | <p><b>none</b>—Display RIP statistics for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP statistics for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Display RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear rip statistics on page 4199</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                      | <a href="#">show rip statistics on page 4205</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                              | <a href="#">Table 320 on page 4205</a> lists the output fields for the <b>show rip statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                        |



Table 320: show rip statistics Output Fields

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIP info</b>                 | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIP.</li> <li>• <b>update interval</b>—Interval between routing table updates, in seconds.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>timeout</b>—Timeout interval, in seconds.</li> <li>• <b>restart in progress</b>—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.</li> <li>• <b>restart time</b>—Estimated time for the graceful restart to finish, in seconds.</li> <li>• <b>restart will complete in</b>—Remaining time for the graceful restart to finish, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIP.</li> <li>• <b>rts held down</b>—Number of routes held down by RIP.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> </ul>                                                                                                                                                   |
| <b><i>logical-interface</i></b> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> <li>• <b>routes learned</b>—Number of routes learned on the logical interface.</li> <li>• <b>routes advertised</b>—Number of routes advertised by the logical interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Counter</b>                  | <p>List of counter types:</p> <ul style="list-style-type: none"> <li>• <b>Updates Sent</b>—Number of update messages sent.</li> <li>• <b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li>• <b>Responses Sent</b>—Number of response messages sent.</li> <li>• <b>Bad Messages</b>—Number of invalid messages received.</li> <li>• <b>RIPv1 Updates Received</b>—Number of RIPv1 update messages received.</li> <li>• <b>RIPv1 Bad Route Entries</b>—Number of RIPv1 invalid route entry messages received.</li> <li>• <b>RIPv1 Updates Ignored</b>—Number of RIPv1 update messages ignored.</li> <li>• <b>RIPv2 Updates Received</b>—Number of RIPv2 update messages received.</li> <li>• <b>RIPv2 Bad Route Entries</b>—Number of RIPv2 invalid route entry messages received.</li> <li>• <b>RIPv2 Updates Ignored</b>—Number of RIPv2 update messages ignored.</li> <li>• <b>Authentication Failures</b>—Number of received update messages that failed authentication.</li> <li>• <b>RIP Requests Received</b>—Number of RIP request messages received.</li> <li>• <b>RIP Requests Ignored</b>—Number of RIP request messages ignored.</li> </ul> |
| <b>Total</b>                    | Total number of packets for the selected counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Last 5 min</b>               | Number of packets for the selected counter in the most recent 5-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Last minute</b>              | Number of packets for the selected counter in the most recent 1-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s  
restart in progress: restart time 60s; restart will complete in 55s  
      rts learned  rts held down  rqsts dropped  resps dropped  
                  0              0              0              0

so-0/0/0.0: 0 routes learned; 501 routes advertised

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 0     | 0          | 0           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 0     | 0          | 0           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |

## PART 16

# Multicast

- [Overview on page 4209](#)
- [Configuration on page 4245](#)
- [Administration on page 4487](#)



## CHAPTER 48

# Overview

- [Introduction to PIM Basics on page 4209](#)
- [Introduction to PIM Sparse Mode on page 4212](#)
- [Introduction to Static RP on page 4216](#)
- [Introduction to Anycast RP on page 4216](#)
- [Introduction to PIM Bootstrap Router on page 4217](#)
- [Introduction to PIM Filtering on page 4218](#)
- [Introduction to PIM RPT and SPT Cutover on page 4220](#)
- [Introduction to IGMP on page 4229](#)
- [Introduction to IGMP Snooping on page 4232](#)
- [Introduction to MSDP on page 4235](#)
- [Introduction to Source-Specific Multicast on page 4237](#)
- [Introduction to Multicast VLAN Registration on page 4241](#)

## Introduction to PIM Basics

---

- [PIM Overview on page 4209](#)
- [PIM on Aggregated Interfaces on page 4212](#)

## PIM Overview

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same router and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routers connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the router interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the router processes the PIM message, a router can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called “protocol independent” because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [\*G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routers must join and leave multicast groups explicitly. Upstream routers do not forward multicast traffic to a downstream router unless the downstream router has sent an explicit request (by means of a join message) to the rendezvous point (RP) router to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routers build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (\*G) state.
- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routers sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a router receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the router has no interested receivers for the data, and the outgoing interface list becomes empty, the router sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

### Basic PIM Network Components

PIM dense mode requires only a multicast source and series of multicast-enabled routers running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routers called *rendezvous points (RPs)* in the network core. These routers are where upstream join messages from interested receivers meet downstream traffic from the source of the

multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routers find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router (BSR)* running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic. PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable. If it is not routable, then the bootstrap router is unable to send bootstrap messages to update the RP domain members. The **show pim bootstrap** command displays only those bootstrap routers that have routable loopback addresses.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

**Related  
Documentation**

- [Supported IP Multicast Protocol Standards](#) in the *Multicast Protocols Configuration Guide*

## PIM on Aggregated Interfaces

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is “expanded” into its constituent interfaces in the next-hop database.

**Related  
Documentation**

## Introduction to PIM Sparse Mode

---

- [Understanding PIM Sparse Mode](#) on page 4213
- [Designated Router](#) on page 4215



## Understanding PIM Sparse Mode

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (\*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (\*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.



**NOTE:** State—the (\*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and \* represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into

the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).



**NOTE:** If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (\*G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).

PIM sparse mode has standard features for all of these issues.


---

### Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 146 on page 4215](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

**Figure 146: Rendezvous Point as Part of the RPT and SPT**



The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

### RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

#### Related Documentation

- [Understanding Static RP on page 4216](#)
- [Understanding RP Mapping with Anycast RP on page 4217](#)
- [Understanding the PIM Bootstrap Router on page 4217](#)
- [Understanding PIM Auto-RP](#)

### Designated Router

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On

receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.



**NOTE:** In PIM dense mode (PIM-DM), a DR is elected by the same process that PIM-SM uses. However, the only time that a DR has any effect in PIM-DM is when IGMPv1 is used on the interface. (IGMPv2 is the default.) In this case, the DR also functions as the IGMP Query Router because IGMPv1 does not have a Query Router election mechanism.

---

## Introduction to Static RP

- [Understanding Static RP on page 4216](#)

### Understanding Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

#### Related Documentation

- [Configuring Local PIM RPs on page 4264](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4266](#)

---

## Introduction to Anycast RP

- [Understanding RP Mapping with Anycast RP on page 4217](#)

## Understanding RP Mapping with Anycast RP

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in Internet draft draft-ietf-mboned-anycast-rp-08.txt, *Anycast RP Mechanism Using PIM and MSDP*. To access Internet RFCs and drafts, go to the IETF website at <http://www.ietf.org>.

### Related Documentation

- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4266](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 4273](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 4267](#)

## Introduction to PIM Bootstrap Router

- [Understanding the PIM Bootstrap Router on page 4217](#)

## Understanding the PIM Bootstrap Router

To determine which router is the rendezvous point (RP), all routers within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routers that all share the same RP router. The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

### Related Documentation

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6](#)

## Introduction to PIM Filtering

---

- [Understanding Multicast Message Filters on page 4218](#)
- [Filtering MAC Addresses on page 4219](#)
- [Filtering RP and DR Register Messages on page 4219](#)

### Understanding Multicast Message Filters

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.



**NOTE:** If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.



**NOTE:** If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

#### Related Documentation

- [Filtering MAC Addresses on page 4219](#)
- [Filtering RP and DR Register Messages on page 4219](#)
- [Filtering MSDP SA Messages on page 4236](#)

## Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

## Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address
- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and

the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

**Related  
Documentation**

- [Understanding RP Mapping with Anycast RP on page 4217](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 4282](#)

---

## Introduction to PIM RPT and SPT Cutover

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 4220](#)
- [Building an RPT Between the RP and Receivers on page 4221](#)
- [PIM Sparse Mode Source Registration on page 4222](#)
- [Multicast Shortest-Path Tree on page 4225](#)
- [SPT Cutover on page 4226](#)
- [SPT Cutover Control on page 4229](#)

## Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (\*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.



To join the shared tree, or *rendezvous-point tree (RPT)* as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.
- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (\*G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (\*G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (\*G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined
- A router along the RPT that already has a multicast forwarding state for the group that is being joined

In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to “migrate” a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (\*G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table.

#### Related Documentation

- [Understanding Multicast Reverse Path Forwarding](#)

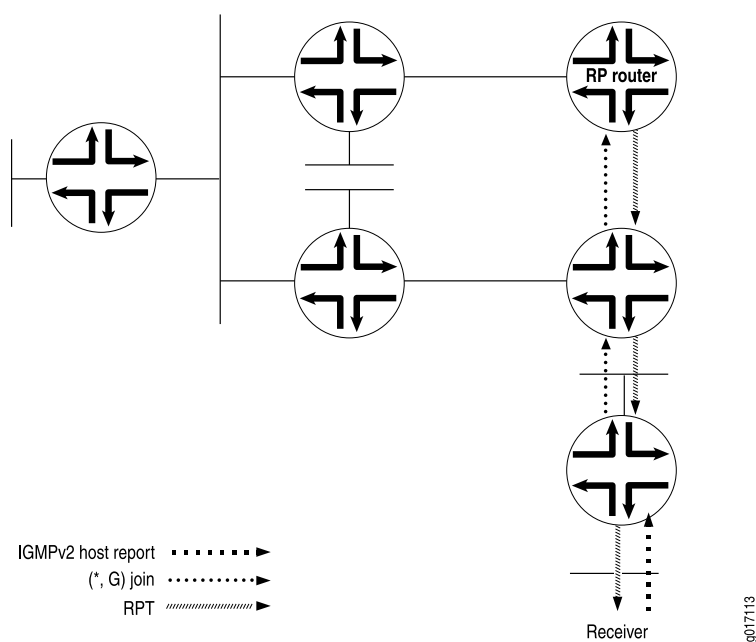
## Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see [Figure 147 on page 4222](#)). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.

3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

**Figure 147: Building an RPT Between the RP and the Receiver**




## PIM Sparse Mode Source Registration

The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:


1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see [Figure 148 on page 4223](#)).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 148: PIM Register Message and PIM Join Message Exchanged




3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see [Figure 149 on page 4224](#)).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 149: Traffic Sent from the Source to the RP Router



- The RP router sends the multicast traffic down the RPT toward the receiver (see [Figure 150 on page 4224](#)).

Figure 150: Traffic Sent from the RP Router Toward the Receiver



## Multicast Shortest-Path Tree

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.

SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point.

### Related Documentation


- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 4220](#)

## SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:


1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see [Figure 151 on page 4226](#)).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

**Figure 151: Receiver DR Sends a PIM Join Message to the Source**




4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see [Figure 152 on page 4227](#)).

Figure 152: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router



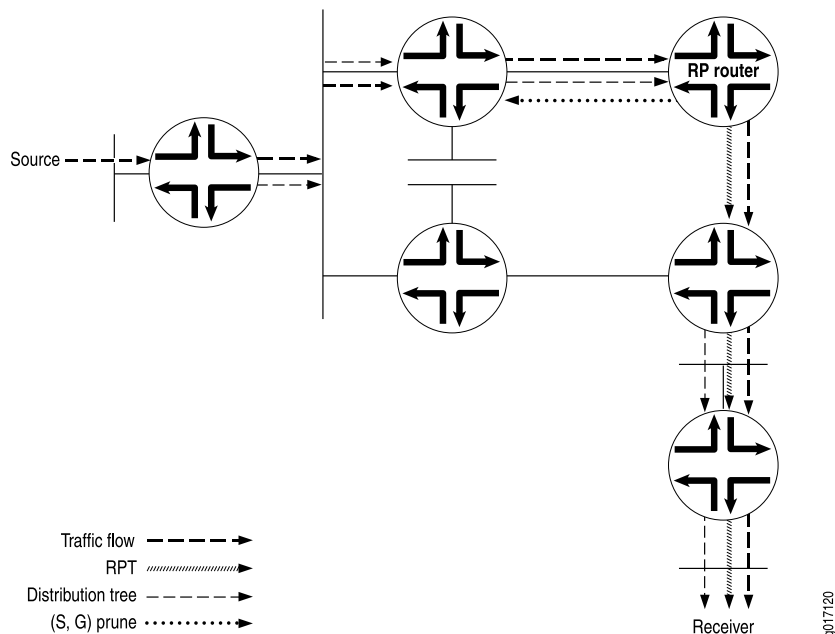
5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see [Figure 153 on page 4227](#)).

Figure 153: RP Router Receives PIM Prune Message




6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see [Figure 154 on page 4228](#)).

**Figure 154: RP Router Sends a PIM Prune Message to the Source DR**



7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see [Figure 155 on page 4228](#)).

**Figure 155: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router**





## SPT Cutover Control

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

---

## Introduction to IGMP

- [Understanding Group Membership Protocols on page 4229](#)
- [Understanding IGMP on page 4230](#)

## Understanding Group Membership Protocols

There is a big difference between the multicast protocols used between host and router and between the multicast routers themselves. Hosts on a given subnetwork need to inform their router only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routers only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routers of their participation in a multicast group. Between adjacent routers, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a router to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the router sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the router that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routers:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the router, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the router, especially on older or smaller routers.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routers can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a router to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any router attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the router.

#### Related Documentation

- *Examples: Configuring MLD*

## Understanding IGMP

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

A router receives explicit join and prune messages from those neighboring routers that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The router then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routers are automatically or statically designated as the RP, and all routers must explicitly join through the RP.

4. Each router along the path toward the RP builds a wild card (any-source) state for the group and sends join and prune messages toward the RP.

The term *route entry* is used to refer to the state maintained in a router to represent the distribution tree.

A route entry can include such fields as:

- source address
- group address
- incoming interface from which packets are accepted
- list of outgoing interfaces to which packets are sent
- timers
- flag bits

The wild card route entry's incoming interface points toward the RP.

The outgoing interfaces point to the neighboring downstream routers that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP multicast traffic.

For each attached network, a multicast router can be either a querier or a nonquerier. The querier router periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a router can specify particular routers from which it accepts or rejects traffic. With IGMPv3, a multicast router can learn which sources are of interest to neighboring routers.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routers, IGMPv3 routers must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

#### Related Documentation

- *Supported IP Multicast Protocol Standards*
- *Configuring IGMP*

## Introduction to IGMP Snooping

---

- [IGMP Snooping Overview on page 4232](#)

### IGMP Snooping Overview

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This IGMP snooping topic includes:

- [How IGMP Snooping Works on page 4232](#)
- [How IGMP Snooping Works with Routed VLAN Interfaces on page 4233](#)
- [How Hosts Join and Leave Multicast Groups on page 4233](#)
- [IGMP Snooping Support for IGMPv3 on page 4233](#)
- [IGMP Snooping and Forwarding Interfaces on page 4234](#)
- [General Forwarding Rules on page 4234](#)

### How IGMP Snooping Works

---

A switch usually learns unicast MAC addresses by checking the source address field of the frames it receives and then sends any traffic for that unicast address only to the appropriate interface. However, a multicast MAC address can never be the source address for a packet. As a result, when a switch receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, which can cause a significant amount of traffic to be sent unnecessarily.

IGMP snooping prevents this flooding. When you enable IGMP snooping, the switch monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast cache table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast packets, it uses the cache table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.



**NOTE:** IGMP snooping is enabled by default on all VLANs.

---



**NOTE:** You cannot configure IGMP snooping on a secondary (private) VLAN.

---

### How IGMP Snooping Works with Routed VLAN Interfaces

A switch can use a routed VLAN interface (RVI) to forward traffic between VLANs that connect to it. IGMP snooping works with Layer 2 interfaces and RVIs to forward multicast traffic in a switched network.

When a switch receives a multicast packet, its Packet Forwarding Engines perform a multicast lookup on the packet to determine how to forward the packet to its local interfaces. From the results of the lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces that have ports local to the Packet Forwarding Engine. If the list includes an RVI, the switch provides a bridge multicast group ID for the RVI to the Packet Forwarding Engine.

For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID, which identifies the Layer 2 interfaces in the VLAN that are interested in receiving the multicast stream. The Packet Forwarding Engine then forwards multicast traffic to bridge multicast IDs that have multicast receivers for a given multicast group.

### How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, either a host cannot respond to the periodic general IGMP queries, which results in a “silent leave” (the only leave option for IGMPv1), or a host can send a group-specific IGMPv2 leave message.

### IGMP Snooping Support for IGMPv3

IGMPv3 enables IGMP snooping to filter multicast streams based on the source address of the multicast stream. A switch supports only IGMPv3 INCLUDE reports. EXCLUDE reports are dropped.



**NOTE:** IGMPv3 and IGMPv3 snooping are not supported on QFabric systems.

When a host sends an IGMPv3 INCLUDE report through a switch interface to indicate that it wants to receive a multicast stream from a source address, the switch adds the source address to its source list. The switch requests traffic for the specified multicast group from only those IP source addresses listed in the source-list parameter. However, because a switch does not support forwarding on a per-source basis, it merges all IGMPv3 reports for a VLAN to create a (\*G,V) route. This means that hosts on the same VLAN

interested in traffic from group G can receive traffic for that group even if they do not specify the source in their INCLUDE report. For example, if Host 1 wants traffic for G from Source A and Host 2 wants traffic for G from Source B, they both receive traffic for G regardless of whether A or B sends the traffic. When IGMP snooping for IGMPv3 is used with an RVI, the same (\*G,V) route is added to the snooping information in the RVI's output interface list (olist).

### IGMP Snooping and Forwarding Interfaces

---

To determine how to forward multicast traffic, a switch with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.



**NOTE:** For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. This is often a multicast router, but if there is no multicast router on the local network, you can configure the switch itself to be an IGMP querier.

---

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

### General Forwarding Rules

---

Multicast traffic received on a switch interface in a VLAN on which IGMP snooping is enabled is forwarded according to the following rules.

IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.

- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic is forwarded as follows:

- A multicast packet with a destination address of 224.0.0.0/24 is flooded to all other interfaces on the VLAN.
- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

#### Related Documentation

- [Example: Configuring IGMP Snooping on page 4318](#)
- [Configuring IGMP Snooping on page 4317](#)
- [Changing the IGMP Snooping Group Timeout Value on page 4318](#)
- [Monitoring IGMP Snooping on page 4487](#)
- [Configuring IGMP on page 4295](#)
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments*
- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

## Introduction to MSDP

- [Understanding MSDP on page 4235](#)
- [Filtering MSDP SA Messages on page 4236](#)

### Understanding MSDP

The Multicast Source Discovery Protocol (MSDP) is used to connect multicast routing domains. It typically runs on the same router as the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

The peer with the higher IP address passively listens to a well-known port number and waits for the side with the lower IP address to establish a Transmission Control Protocol (TCP) connection. When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends source-active type, length, and values (TLVs) to its MSDP

peers. When a source-active TLV is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure that this peer is in the path that leads back to the originating RP. If not, the source-active TLV is dropped. This TLV is counted as a “rejected” source-active message.

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member. For more information about configuring MSDP mesh groups, see [“Example: Configuring MSDP with Active Source Limits and Mesh Groups” on page 4332](#).

Router R locates its MSDP peer-RPF neighbor (Router N) deterministically. A series of rules is applied in a particular order to received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected.

The six rules applied to source-active messages originating at Router S received at Router R from Router X are as follows:

1. If Router X originated the source-active message (Router X is Router S), then Router X is also the peer-RPF neighbor, and its source-active messages are accepted.
2. If Router X is a member of the Router R mesh group, or is the configured peer, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
3. If Router X is the BGP next hop of the active multicast RPF route toward Router S (Router X installed the route on Router R), then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
4. If Router X is an external BGP (EBGP) or internal BGP (IBGP) peer of Router R, and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router X's AS number, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
5. If Router X uses the same next hop as the next hop to Router S, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
6. If Router X fits none of these criteria, then Router X is not an MSDP peer-RPF neighbor, and its source-active messages are rejected.

The MSDP peers that receive source-active TLVs can be constrained by BGP reachability information. If the AS path of the network layer reachability information (NLRI) contains the receiving peer's AS number prepended second to last, the sending peer is using the receiving peer as a next hop for this source. If the split horizon information is not being received, the peer can be pruned from the source-active TLV distribution list.

**Related Documentation**

- [Configuring MSDP on page 4327](#)

## Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network,



you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.



**NOTE:** When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

#### Related Documentation

- [Understanding Multicast Administrative Scoping](#)
- [Filtering Incoming PIM Join Messages on page 4281](#)
- [Example: Configuring PIM BSR Filters on page 4278](#)

## Introduction to Source-Specific Multicast

- [Source-Specific Multicast Groups Overview on page 4237](#)
- [Understanding PIM Source-Specific Mode on page 4238](#)
- [PIM SSM on page 4239](#)

### Source-Specific Multicast Groups Overview

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (\*,G) pairs. The (\*) is a wildcard referring to any source sending to group G, and "G" refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The "S" refers to the source's unicast IP address, and the "G" refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol. IGMPv3 is available for Windows XP, Windows Vista, and most UNIX operating systems.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

## Understanding PIM Source-Specific Mode

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. The RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to

*subscribers* (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in [Table 321 on page 4239](#).

**Table 321: ASM and SSM Terminology**

| Term                | Any-Source Multicast  | Source-Specific Multicast         |
|---------------------|-----------------------|-----------------------------------|
| Address identifier  | G                     | S,G                               |
| Address designation | group                 | channel                           |
| Receiver operations | join, leave           | subscribe, unsubscribe            |
| Group address range | 224/4 excluding 232/8 | 224/4 (guaranteed only for 232/8) |

Although PIM SSM describes receiver operations as *subscribe* and *unsubscribe*, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

## PIM SSM

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.


You can also configure the Junos OS to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 156 on page 4240](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 156 on page 4240](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

**Figure 156: Receiver Announces Desire to Join Group G and Source S**




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 157 on page 4240](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

**Figure 157: Router 3 (Last-Hop Router) Joins the Source Tree**



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 158 on page 4241](#)).

Figure 158: (S,G) State Is Built Between the Source and the Receiver



To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

#### Related Documentation

- [Source-Specific Multicast Groups Overview on page 4237](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347](#)

## Introduction to Multicast VLAN Registration

- [Understanding Multicast VLAN Registration on page 4241](#)

### Understanding Multicast VLAN Registration

Multicast VLAN registration (MVR) enables you to efficiently distribute IPTV multicast streams across an Ethernet ring-based Layer 2 network and reduce the amount of bandwidth consumed by this multicast traffic.

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which IPTV multicast traffic flows throughout the Layer 2 network. The Juniper Networks EX Series Ethernet Switch or the QFX Series that is enabled for MVR selectively forwards IPTV multicast traffic from interfaces on the MVLAN (source interfaces) to hosts that are connected to interfaces that are not part of the MVLAN. These interfaces are known as *MVR receiver ports*. The MVR receiver ports can receive traffic from a port on the MVLAN but cannot send traffic onto the MVLAN, and they remain in their own VLANs for bandwidth and security reasons.

This topic includes:

- [How MVR Works on page 4241](#)

#### How MVR Works

In many ways, MVR is similar to IGMP snooping. Both MVR and IGMP snooping monitor IGMP join and leave messages and build forwarding tables based on the media access control (MAC) addresses of the hosts sending those IGMP messages. Whereas IGMP snooping operates within a given VLAN to regulate multicast traffic, MVR can operate

with hosts on different VLANs in a Layer 2 network to selectively deliver IPTV multicast traffic to requesting hosts, thereby reducing the amount of bandwidth needed to forward multicast traffic.

When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs. Interfaces that are on the MVLAN itself cannot be MVR receiver ports for that MVLAN.



**NOTE:** MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

### ***MVR Modes***

MVR operates in two modes: MVR transparent mode and MVR proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

- [MVR Transparent Mode on page 4242](#)
- [MVR Proxy Mode on page 4242](#)

### ***MVR Transparent Mode***

In MVR transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Transparent mode is the default mode.

The switch handles IGMP packets destined for both the multicast source VLAN and multicast receiver VLANs in the same way that it handles them when MVR is not being used. That is, when a host on a VLAN sends IGMP join and leave messages, the switch floods the messages to all router interfaces in the VLAN. Similarly, when a VLAN receives IGMP queries from its router interfaces, it floods the queries to all interfaces in the VLAN.

If a host on a multicast receiver port joins an MVR group on the multicast receiver VLAN, the appropriate bridging entry is added and the MVLAN forwards that group's IPTV multicast traffic on that port (even though that port is not in the MVLAN). Likewise, if a host on a multicast receiver port leaves an MVR group on the multicast receiver VLAN, the appropriate bridging entry is deleted, and the MVLAN stops forwarding that group's IPTV multicast traffic on that port. In addition, you can configure the switch to statically install the bridging entries on the multicast receiver VLAN.

### ***MVR Proxy Mode***

When you use MVR in proxy mode, the switch acts as a proxy for any MVR group in both the upstream and downstream directions. In the downstream direction, the switch acts as the querier for the groups in the MVR receiver VLANs. In the upstream direction, the switch originates the IGMP reports and leaves and answers IGMP queries from multicast routers. When the MVR receiver VLANs receive IGMP joins and leaves, the switch creates bridging entries on the MVLAN as needed, as it does in MVR transparent mode. In addition, the switch sends out IGMP joins and leaves on the MVLAN based on these bridging entries.

Configuring MVR proxy mode on the MVLAN automatically enables IGMP snooping proxy mode on all MVR receiver VLANs as well as on the MVLAN.

**Related  
Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Example: Configuring Multicast VLAN Registration on page 4322](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 4321](#)





## CHAPTER 49

# Configuration

- [Optimizing Multicast Flows on QFabric Systems on page 4245](#)
- [PIM Basics on page 4246](#)
- [PIM Designated Router on page 4253](#)
- [PIM Sparse Mode on page 4254](#)
- [Static RP on page 4264](#)
- [Anycast RP on page 4267](#)
- [PIM Bootstrap Router on page 4276](#)
- [PIM Filtering on page 4278](#)
- [PIM RPT and SPT Cutover on page 4284](#)
- [PIM and the BFD Protocol on page 4290](#)
- [IGMP on page 4295](#)
- [IGMP Snooping on page 4316](#)
- [MSDP on page 4327](#)
- [Source-Specific Multicast on page 4342](#)
- [PIM Configuration Statements on page 4354](#)
- [IGMP Configuration Statements on page 4424](#)
- [IGMP Snooping Configuration Statements on page 4448](#)
- [MSDP Configuration Statements on page 4461](#)
- [Source-Specific Multicast Configuration Statements on page 4481](#)

### Optimizing Multicast Flows on QFabric Systems

- [Optimizing the Number of Multicast Flows on QFabric Systems on page 4246](#)

## Optimizing the Number of Multicast Flows on QFabric Systems

Because of the distributed nature of QFabric systems, the default configuration does not allow the maximum number of supported Layer 3 multicast flows to be created. To allow a QFabric system to create the maximum number of supported flows, configure the following statement:

```
set fabric routing-options multicast fabric-optimized-distribution
```

After configuring this statement, you must reboot the QFabric Director group to make the change take effect.

**Related  
Documentation**

- 

---

## PIM Basics

- [Changing the PIM Version on page 4246](#)
- [Modifying the PIM Hello Interval on page 4246](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility on page 4247](#)
- [Configuring PIM Trace Options on page 4248](#)
- [Disabling PIM on page 4250](#)

## Changing the PIM Version

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

## Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
user@host# set hello-interval 255
```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-3/0/2.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

**Related Documentation** • [show pim neighbors on page 4601](#) in the [CLI Explorer](#)

## Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

```
[edit system]
user@host# set no-multicast-echo
```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

```
user@host> show system statistics icmp

icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

- Related Documentation**
- *Configuring the Junos OS to Disable the Routing Engine Response to Multicast Ping Packets* in the *Junos OS System Basics Configuration Guide*
  - *show system statistics icmp* in the [CLI Explorer](#)

## Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                             | Description                                                                                                                                                     |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>all</b>                       | Trace all operations.                                                                                                                                           |
| <b>assert</b>                    | Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN. |
| <b>autorp</b>                    | Trace bootstrap, RP, and auto-RP messages.                                                                                                                      |
| <b>bidirectional-df-election</b> | Trace bidirectional PIM designated-forwarder (DF) election events.                                                                                              |
| <b>bootstrap</b>                 | Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap router and are forwarded, hop by hop, to all routers in that domain.        |
| <b>general</b>                   | Trace general events.                                                                                                                                           |

| Flag                       | Description                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>graft</b>               | Trace graft and graft acknowledgment messages.                                                                                           |
| <b>hello</b>               | Trace hello packets, which are sent so that neighboring routers can discover one another.                                                |
| <b>join</b>                | Trace join messages, which are sent to join a branch onto the multicast distribution tree.                                               |
| <b>mdt</b>                 | Trace messages related to multicast data tunnels.                                                                                        |
| <b>normal</b>              | Trace normal events.                                                                                                                     |
| <b>nsr-synchronization</b> | Trace nonstop routing synchronization events                                                                                             |
| <b>packets</b>             | Trace all PIM packets.                                                                                                                   |
| <b>policy</b>              | Trace poison-route-reverse packets.                                                                                                      |
| <b>prune</b>               | Trace prune messages, which are sent to prune a branch off the multicast distribution tree.                                              |
| <b>register</b>            | Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group. |
| <b>route</b>               | Trace routing information.                                                                                                               |
| <b>rp</b>                  | Trace candidate RP advertisements.                                                                                                       |
| <b>state</b>               | Trace state transitions.                                                                                                                 |
| <b>task</b>                | Trace task processing.                                                                                                                   |
| <b>timer</b>               | Trace timer processing.                                                                                                                  |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the PIM trace file.

```
[edit protocols pim traceoptions]
user@host# set file pim-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags.

Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match "Rx V1 Require V2"
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

#### Related Documentation

- [PIM Overview on page 4209](#)
- *Junos OS Tracing and Logging Operations* in the *Junos OS System Basics Configuration Guide*

## Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

- [Disabling the PIM Protocol on page 4251](#)
- [Disabling PIM On an Interface on page 4251](#)
- [Disabling PIM for a Family on page 4252](#)
- [Disabling PIM for a Rendezvous Point on page 4252](#)

---

### Disabling the PIM Protocol

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
 disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.

```
user@host# set protocols pim disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

---

### Disabling PIM On an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
 interface interface-name {
 disable;
 }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

```
user@host# set protocols pim interface fe-0/1/0 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### Disabling PIM for a Family

---

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

```
user@host# set protocols pim family inet disable
user@host# set protocols pim family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### Disabling PIM for a Rendezvous Point

---

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
 rp {
 local {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
 }
 }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.

```
user@host# set protocols pim rp local family inet disable
user@host# set protocols pim rp local family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```



## PIM Designated Router

- [Configuring Interface Priority for PIM Designated Router Selection on page 4253](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 4254](#)

### Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) routing device learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has equal probability (priority 1) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
user@host# set priority 5
```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
```

```
Instance: PIM.master
Interface: ge-0/0/0.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 5
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0
```

```
Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

```
Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

- Related Documentation**
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 4254](#)
  - [Understanding PIM Sparse Mode on page 4213](#)
  - [show pim neighbors on page 4601](#) in the CLI Explorer

## Configuring PIM Designated Router Election on Point-to-Point Links

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.  

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```
2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.
3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routers on the point-to-point link. Then recheck the state.



**CAUTION:** Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

```
[edit]
user@host# run restart routing
```

- Related Documentation**
- [Understanding PIM Sparse Mode on page 4213](#)
  - [Configuring Interface Priority for PIM Designated Router Selection on page 4253](#)
  - [show pim interfaces on page 4589](#) in the CLI Explorer

---

## PIM Sparse Mode

- [Enabling PIM Sparse Mode on page 4255](#)
- [Configuring PIM Join Load Balancing on page 4256](#)
- [Modifying the Join State Timeout on page 4259](#)
- [Example: Enabling Join Suppression on page 4259](#)

## Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. The following example explicitly configures PIMv2 on the interfaces.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

```
[edit protocols pim]
user@host# set rp local family inet address 192.168.3.253
```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 198.58.3.253 version 2
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.

- `show pim interfaces`
- `show pim join`

- [show pim neighbors](#)
- [show pim rps](#)

**Related Documentation** • [Understanding PIM Sparse Mode on page 4213](#)

## Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the **show pim join extensive** or **show pim source** commands.

```

user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
 Source: *
 RP: 10.255.245.6
 Flags: sparse,rptree,wildcard
 Upstream interface: t1-0/2/3.0
 Upstream neighbor: 192.168.38.57
 Upstream state: Join to RP
 Downstream neighbors:
 Interface: t1-0/2/1.0
 192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
 Source: *
 RP: 10.255.245.6
 Flags: sparse,rptree,wildcard
 Upstream interface: so-0/3/0.0
 Upstream neighbor: 192.168.38.47
 Upstream state: Join to RP
 Downstream neighbors:
 Interface: t1-0/2/3.0
 192.168.38.16 State: JOIN Flags; SRW Timeout: 164

```

Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM join load balancing.

```

[edit protocols pim rp]
user@host# set static address 10.10.10.1
user@host# set interface all mode sparse version 2
user@host# set join-load-balance

```

The static address is the address of the RP.

3. Monitor the operation.

If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the **show pim interfaces** command.

```

user@host> show pim interfaces
Instance: PIM.master

```

| Name           | Stat | Mode   | IP V | State | NbrCnt | JoinCnt | DR address         |
|----------------|------|--------|------|-------|--------|---------|--------------------|
| lo0.0          | Up   | Sparse | 4 2  | DR    | 0      | 0       | 10.255.168.58      |
| pe-1/2/0.32769 | Up   | Sparse | 4 2  | P2P   | 0      | 0       |                    |
| so-0/3/0.0     | Up   | Sparse | 4 2  | P2P   | 1      | 1       |                    |
| t1-0/2/1.0     | Up   | Sparse | 4 2  | P2P   | 1      | 0       |                    |
| t1-0/2/3.0     | Up   | Sparse | 4 2  | P2P   | 1      | 1       |                    |
| lo0.0          | Up   | Sparse | 6 2  | DR    | 0      | 0       | fe80::2a0:a5ff:4b7 |

Note that the two equal-cost paths shown by the **show pim interfaces** command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the **show pim neighbors detail** output:

```
user@host> show pim neighbors detail
Interface: so-0/3/0.0
```

```
Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 1689116164
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 102 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 792890329
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Interface: t1-0/2/3.0
```

```
Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 678582286
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1854475503
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The **clear pim join-distribution** command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

#### Related Documentation

- *clear pim join-distribution* in the [CLI Explorer](#)
- [show pim interfaces on page 4589](#) in the [CLI Explorer](#)
- [show pim neighbors on page 4601](#) in the [CLI Explorer](#)
- [show pim source on page 4612](#) in the [CLI Explorer](#)

## Modifying the Join State Timeout

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the **join-prune-timeout** statement.

To modify the timeout, include the **join-prune-timeout** statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 240 seconds.

**Related Documentation**

- [join-prune-timeout on page 4385](#)

## Example: Enabling Join Suppression

This example describes how to enable PIM join suppression.

- [Requirements on page 4259](#)
- [Overview on page 4259](#)
- [Configuration on page 4262](#)
- [Verification on page 4263](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos® OS Network Interfaces*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Configuration Guide*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 4255](#).

### Overview

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.

This example includes the following statements:


- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.
- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

[Figure 159 on page 4261](#) shows the topology used in this example.



Figure 159: Join Suppression



The items in the figure represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

## Configuration

---

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```

5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
[edit protocols pim]
user@host# set traceoptions flag register detail
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols pim]
user@host# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
 traceoptions {
 file pim.log size 5m world-readable;
 flag join detail;
 flag prune detail;
 flag normal detail;
 flag register detail;
 }
 rp {
 static {
 address 10.255.112.160;
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 reset-tracking-bit;
 propagation-delay 500;
 override-interval 4000;
}
```

### Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- **show pim join extensive**
- **show multicast route extensive**

**Related Documentation**

- [Example: Configuring the PIM Assert Timeout on page 4284](#)
- [Example: Configuring PIM RPF Selection](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 4287](#)
- [Enabling PIM Sparse Mode on page 4255](#)
- [PIM Overview on page 4209](#)

---

## Static RP

---

- [Configuring Local PIM RPs on page 4264](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4266](#)

### Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

```
[routing-instances VPN-A protocols pim]
user@host# set rp local
```

2. Configure the IP protocol family and IP address.

IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface *interface-name*]** hierarchy level and **family inet6** at the **[edit protocols pim interface *interface-name*]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
user@host# set family inet address 10.1.2.254
```

3. (IPv4 only) Configure the routing device's RP priority.



**NOTE:** The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set group-ranges fec0::/10
user@host# set group-ranges 10.1.2.0/24
```

5. (IPv4 only) Modify the local RP hold time.

If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set hold-time 200
```

6. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set override
```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

- Related Documentation**
- [PIM Overview on page 4209](#)
  - [Understanding MLD](#)

## Configuring the Static PIM RP Address on the Non-RP Routing Device

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

```
[routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334
```



**NOTE:** Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
```



**NOTE:** The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address *address*]). PIM version 2 is the default for interface mode ([edit pim interface *interface-name*]). Explicitly configured versions override the defaults.

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp static address
 2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

**Related  
Documentation**

- [PIM Overview on page 4209](#)
- [Understanding MLD](#)

## Anycast RP

- [Example: Configuring PIM Anycast With or Without MSDP on page 4267](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 4271](#)
- [Configuring a PIM Anycast RP Router Using Only PIM on page 4272](#)
- [Configuring All PIM Anycast Non-RP Routers on page 4273](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 4273](#)

### Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32;
 primary;
 address 198.58.3.253/32;
 }
 }
 }
}
```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
```



```

 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
}
}

```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```

protocols {
 msdp {
 peer 198.58.3.250 {
 local-address address 198.58.3.254;
 }
 }
}

```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
 }
 }
}

```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface**

**all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
}
```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
 pim {
 rp {
 static {
 address 198.58.3.253 {
 version 2;
 }
 }
 }
 }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}
```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
 }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
```

```
peer 198.58.3.250 {
 local-address 198.58.3.254;
}
}
}
```

## Configuring a PIM Anycast RP Router Using Only PIM

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
 }
 }
}
```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```
protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
}
```

```

 }
 }
}
interface all {
 mode sparse;
 version 2;
}
interface fxp0.0 {
 disable;
}
}
}

```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

## Configuring All PIM Anycast Non-RP Routers

Use the **mode** statement at the **[edit protocols pim rp interface all]** hierarchy level to specify sparse mode on all interfaces. Then add the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```

protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}

```

## Example: Configuring Multiple RPs in a Domain with Anycast RP

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

- [Requirements on page 4273](#)
- [Overview on page 4274](#)
- [Configuration on page 4274](#)
- [Verification on page 4276](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos® OS Network Interfaces*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Configuration Guide*.

- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 4255](#).

---

## Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

---

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**RP Routers**

```
set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set protocols msdp local-address 192.168.132.1
set protocols msdp peer 192.168.12.1
set protocols pim rp local address 10.1.1.2
set routing-options router-id 192.168.132.1
```

**Non-RP Routers**      `set protocols pim rp static address 10.1.1.2`

### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.1.1.2/32
```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary
```

```
[edit routing-options]
user@host# set router-id 192.168.132.1
```

3. On each RP router in the domain, configure the local RP address, using the shared address.

```
[edit protocols pim]
user@host# set rp local address 10.1.1.2
```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

```
[edit protocols msdp]
user@host# set local-address 192.168.132.1
user@host# set peer 192.168.12.1
```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

```
[edit protocols pim]
user@host# set rp static address 10.1.1.2
```

6. If you are done configuring the devices, commit the configuration.

```
user@host# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
lo0 {
 unit 0 {
 family inet {
 address 192.168.132.1/32 {
 primary;
 }
 address 10.1.1.2/32;
 }
 }
}
```

*On the RP routers:*

```
user@host# show protocols
msdp {
 local-address 192.168.132.1;
 peer 192.168.12.1;
}
pim {
 rp {
 local {
 address 10.1.1.2;
 }
 }
}
```

*On the non-RP routers:*

```
user@host# show protocols
pim {
 rp {
 static {
 address 10.1.1.2;
 }
 }
}
```

```
 }
 }
}
```

```
user@host# show routing-options
router-id 192.168.132.1;
```

---

### Verification

To verify the configuration, run the `show pim rps extensive inet` command.

#### Related Documentation

- [Example: Configuring PIM Anycast With or Without MSDP on page 4267](#)
- [Understanding PIM Sparse Mode on page 4213](#)
- [Understanding RP Mapping with Anycast RP on page 4217](#)

---

## PIM Bootstrap Router

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4276](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 4277](#)
- [Example: Configuring PIM BSR Filters on page 4278](#)

### Configuring PIM Bootstrap Properties for IPv4 or IPv6

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.





**NOTE:** In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the combined IPv4 and IPv6 configuration, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the `bootstrap` statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routers in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
user@host# exit
```

3. Configure the policies.

```
user@host# edit policy-options policy-statement pim-bootstrap-import
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
```

4. Monitor the operation of PIM bootstrap routers by running the `show pim bootstrap` command.

#### Related Documentation

- [Understanding PIM Sparse Mode on page 4213](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 4277](#)
- [show pim bootstrap on page 4587](#) in the CLI Explorer

### Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```
protocols {
 pim {
```

```
rp {
 bootstrap {
 family inet {
 priority 1;
 import pim-import;
 export pim-export;
 }
 family inet6 {
 priority 1;
 import pim-import;
 export pim-export;
 }
 }
}
}
policy-options {
 policy-statement pim-import {
 from interface so-0/1/0;
 then reject;
 }
 policy-statement pim-export {
 to interface so-0/1/0;
 then reject;
 }
}
```

### Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
 pim {
 rp {
 bootstrap-import no-bsr;
 bootstrap-export no-bsr;
 }
 }
}
policy-options {
 policy-statement no-bsr {
 then reject;
 }
}
```

---

## PIM Filtering

- [Configuring Interface-Level PIM Neighbor Policies on page 4279](#)
- [Filtering Outgoing PIM Join Messages on page 4280](#)
- [Filtering Incoming PIM Join Messages on page 4281](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 4282](#)

## Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

```
[edit policy-options]
user@host# set prefix-list nbrGroup 1 20.20.20.1/32
user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
user@host# set policy-statement nbr-policy then reject
```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set neighbor-policy nbr-policy
```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

### Related Documentation

- [Understanding PIM Sparse Mode on page 4213](#)
- *Defining Routing Policies* in the *Routing Policy Configuration Guide*
- [show pim statistics on page 4615](#) in the *CLI Explorer*

## Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
225.1.1.1/32 exact
user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept
```

2. Verify the policy configuration by running the **show policy-options** command.

```
user@host# show policy-options
policy-statement block-groups {
 term t1 {
 from {
 route-filter 224.0.1.2/32 exact;
 route-filter 225.1.1.1/32 exact;
 then reject;
 }
 term last {
 then accept;
 }
 }
}
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```
user@host> show pim statistics | grep filtered
```

|                          |     |
|--------------------------|-----|
| RP Filtered Source       | 0   |
| Rx Joins/Prunes filtered | 0   |
| Tx Joins/Prunes filtered | 254 |

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

**Related Documentation**

- [Filtering Incoming PIM Join Messages on page 4281](#)

## Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (\*G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See [Table 322 on page 4281](#) for a list of match conditions.

**Table 322: PIM Join Filter Match Conditions**

| Match Condition              | Matches On                                                                           |
|------------------------------|--------------------------------------------------------------------------------------|
| <b>interface</b>             | Router interface or interfaces specified by name or IP address                       |
| <b>neighbor</b>              | Neighbor address (the source address in the IP header of the join and prune message) |
| <b>route-filter</b>          | Multicast group address embedded in the join and prune message                       |
| <b>source-address-filter</b> | Multicast source address embedded in the join and prune message                      |

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (\*G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

[edit [policy-statement](#) pim-join-filter term bad-groups]

```
user@host# set from route-filter 224.0.1.2/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term bad-sources]
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term last]
user@host# set then accept
```

2. Apply one or more policies to routes being imported into the routing table from PIM.

```
[edit protocols pim]
user@host# set import pim-join-filter
```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

#### Related Documentation

- *Understanding Multicast Administrative Scoping*
- [Filtering Outgoing PIM Join Messages on page 4280](#)
- [show pim join on page 4592](#) in the [CLI Explorer](#)
- *show policy* in the [CLI Explorer](#)

## Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP router.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

```
[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [reject_224_1_1_1 | accept_224_1_1_5]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- RP Filtered Source
- Rx Joins/Prunes filtered
- Tx Joins/Prunes filtered
- Rx Register msgs filtering drop
- Tx Register msgs filtering drop

**Related  
Documentation**

- [PIM Sparse Mode Source Registration on page 4222](#)
- [Filtering RP and DR Register Messages on page 4219](#)
- [show pim statistics on page 4615](#) in the CLI Explorer

---

## PIM RPT and SPT Cutover

- [Example: Configuring the PIM Assert Timeout on page 4284](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 4287](#)

### Example: Configuring the PIM Assert Timeout

This example shows how to configure the timeout period for a PIM assert forwarder.

- [Requirements on page 4284](#)
- [Overview on page 4285](#)
- [Configuration on page 4286](#)

---

#### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos® OS Network Interfaces*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Configuration Guide*.
- Configure PIM Sparse Mode on the interfaces. See “[Enabling PIM Sparse Mode](#)” on [page 4255](#).



---

## Overview

---

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In [Figure 160 on page 4286](#), both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the source, and a preference metric for the unicast cost. The preference metric expresses a preference between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.


When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

[Figure 160 on page 4286](#) shows the topology for this example.

Figure 160: PIM Assert Topology



### Configuration

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an assert timeout:

1. Configure the timeout period, in seconds.  

```
[edit protocols pim]
user@host# set assert-timeout 60
```
2. (Optional) Trace assert messages.  

```
[edit protocols pim]
user@host# set traceoptions file PIM.log
user@host# set traceoptions flag assert detail
```
3. If you are done configuring the device, commit the configuration.  

```
user@host# commit
```
4. To verify the configuration, run the following commands:
  - `show pim join`
  - `show pim statistics`

#### Related Documentation

- [Configuring PIM Trace Options on page 4248](#)
- [SPT Cutover on page 4226](#)
- [SPT Cutover Control on page 4229](#)

## Example: Configuring the PIM SPT Threshold Policy

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

- [Requirements on page 4287](#)
- [Overview on page 4287](#)
- [Configuration on page 4288](#)
- [Verification on page 4290](#)

### Requirements

---

Before you begin:

- Configure the router interfaces. See the *Junos® OS Network Interfaces*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Configuration Guide*.
- Configure PIM Sparse Mode on the interfaces. See “[Enabling PIM Sparse Mode](#)” on [page 4255](#).

### Overview

---

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the SPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.
- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

### Configuration

---

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter
 224.1.1.1/32 exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
 10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

3. If you are done configuring the device, commit the configuration.

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

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

### Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
 term one {
 from {
 route-filter 224.1.1.1/32 exact;
 source-address-filter 10.10.10.1/32 exact;
 }
 then accept;
 }
 term two {
 then reject;
 }
}
```

```
 }
 }

user@host# show protocols
pim {
 spt-threshold {
 infinity spt-infinity-policy;
 }
}
```

---

### Verification

To verify the configuration, run the `show pim join` command.

**Related Documentation**

- [SPT Cutover Control on page 4229](#)

---

## PIM and the BFD Protocol

- [Configuring BFD for PIM on page 4290](#)
- [Configuring BFD Authentication for PIM on page 4292](#)

### Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

**Related Documentation**

- *show bfd session* in the [CLI Explorer](#)

## Configuring BFD Authentication for PIM

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- [Configuring BFD Authentication Parameters on page 4292](#)
- [Viewing Authentication Information for BFD Sessions on page 4293](#)

---

### Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication algorithm
keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication keychain
bfd-pim
```





**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

[edit security]

```
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret
9ggaJDmPQ6/tJgF/AtREVsyPsnCtUhm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

[edit protocols pim]

```
user@host# set interface if3-pim bfd-liveness-detection authentication loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat these steps to configure the other end of the BFD session.

### Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if3-pim** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUhm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface if3-pim {
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-pim;
 }
 }
}
[edit security]
authentication key-chains {
```

```

key-chain bfd-pim {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
 start-time "2009-6-1.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9l.reP38ny.TszF2/9";
 start-time "2009-6-1.15:29:20 -0700";
 }
}

```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session detail

```
user@host# show bfd session detail
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**  
 Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic NbrSignal  
 Remote state Up, version 1  
 Replicated

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**  
**keychain bfd-pim, algo keyed-sha-1, mode strict**  
 Session up time 00:04:42  
 Local diagnostic None, remote diagnostic NbrSignal  
 Remote state Up, version 1  
 Replicated  
 Min async interval 0.300, min slow interval 1.000  
 Adaptive async TX interval 0.300, RX interval 0.300  
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3  
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3  
 Local discriminator 2, remote discriminator 2  
 Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict**

#### Related Documentation

- [Understanding Bidirectional Forwarding Detection Authentication for PIM](#)
- [Configuring BFD for PIM on page 4290](#)
- [authentication-key-chains on page 4920](#)

- [bfd-liveness-detection on page 4363](#)
- [show bfd session in the CLI Explorer](#)

---

## IGMP

---

- [Configuring IGMP on page 4295](#)
- [Enabling IGMP on page 4297](#)
- [Changing the IGMP Version on page 4298](#)
- [Modifying the IGMP Host-Query Message Interval on page 4299](#)
- [Modifying the IGMP Last-Member Query Interval on page 4299](#)
- [Specifying Immediate-Leave Host Removal for IGMP on page 4300](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4301](#)
- [Accepting IGMP Messages from Remote Subnetworks on page 4302](#)
- [Modifying the IGMP Query Response Interval on page 4303](#)
- [Modifying the IGMP Robustness Variable on page 4304](#)
- [Limiting the Maximum IGMP Message Rate on page 4305](#)
- [Enabling IGMP Static Group Membership on page 4305](#)
- [Recording IGMP Join and Leave Events on page 4312](#)
- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4313](#)
- [Tracing IGMP Protocol Traffic on page 4314](#)
- [Disabling IGMP on page 4316](#)

## Configuring IGMP

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
 accounting;
 interface interface-name {
 disable;
 (accounting | no-accounting);
 group-policy [policy-names];
 immediate-leave;
 oif-map map-name;
 promiscuous-mode;
 ssm-map ssm-map-name;
 static {
 group multicast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
 }
 version version;
 }
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).



**NOTE:** You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one “pseudo PIM interface” is created to represent all IGMP downstream (IGMP-only) interfaces on the router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

## Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

```
[edit protocols igmp]
user@host# set interface fe-0/0/0.0
```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0 {
 disable;
}
```

3. Enable IGMP on the interface by deleting the **disable** statement.

```
[edit protocols igmp]
delete interface ge-1/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
```

```
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

**Related  
Documentation**

- [Understanding IGMP on page 4230](#)
- [Disabling IGMP on page 4316](#)
- [show igmp interface on page 4529](#)

## Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.

**Related  
Documentation**

- [Understanding IGMP on page 4230](#)
- [show pim interfaces on page 4589](#)
- [show igmp statistics on page 4533](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

## Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

### Related Documentation

- [Understanding IGMP on page 4230](#)
- [Modifying the IGMP Query Response Interval on page 4303](#)
- [Modifying the IGMP Robustness Variable on page 4304](#)
- [show igmp interface on page 4529](#)
- [show igmp statistics on page 4533](#)

## Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols igmp]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

---

**Related  
Documentation**

- [Modifying the IGMP Robustness Variable on page 4304](#)
- [show pim interfaces on page 4589](#)

## Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.



When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

```
[edit protocols IGMP]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the Immediate Leave field in the output of the `show igmp interface` command.

#### Related Documentation

- [Understanding IGMP on page 4230](#)
- [show igmp interface on page 4529](#)

## Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
```

```
user@host# set then reject
```

2. Configure an IGMPv3 policy.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

#### Related Documentation

- [Understanding IGMP on page 4230](#)
- [Defining Routing Policies](#)
- [show igmp statistics on page 4533](#)

## Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.



**NOTE:** When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.



**NOTE:** When enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.

3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

- Related Documentation**
- [Understanding IGMP on page 4230](#)
  - [Configuring the Loopback Interface in the Junos® OS Network Interfaces](#)
  - [show igmp interface on page 4529](#)
  - [show igmp statistics on page 4533](#)

## Modifying the IGMP Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.
 

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```
2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

- Related Documentation**
- [Understanding IGMP on page 4230](#)
  - [Modifying the IGMP Host-Query Message Interval on page 4299](#)
  - [Modifying the IGMP Robustness Variable on page 4304](#)
  - [show igmp interface on page 4529](#)
  - [show igmp statistics on page 4533](#)

## Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

- Related Documentation**
- [Modifying the IGMP Host-Query Message Interval on page 4299](#)
  - [Modifying the IGMP Query Response Interval on page 4303](#)
  - [Modifying the IGMP Last-Member Query Interval on page 4299](#)
  - [show pim interfaces on page 4589](#)
  - RFC 2236, *Internet Group Management Protocol, Version 2*
  - RFC 3376, *Internet Group Management Protocol, Version 3*

## Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the router.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a router with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

- Related Documentation**
- [maximum-transmit-rate \(Protocols IGMP\) on page 4436](#)

## Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
 static {
 group 225.1.1.1;
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 static {
 group 225.1.1.1 {
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```

```

Group: 225.1.1.2
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
Group: 225.1.1.3
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment
0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 group-increment 0.0.0.2;
 group-count 3;
 }
 }
}

```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```
user@host> show igmp group
```

```

Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.3
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.5
 Source: 10.0.0.2

```

```
Last reported by: Local
Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```



When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 }
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.3
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.4
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count 3 source-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 source-increment 0.0.0.2;
 }
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```
user@host> show igmp group
```

```
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.4
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.6
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 exclude;
 source 10.0.0.2;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
Interface: fe-0/1/2
Group: 225.1.1.1
Group mode: Exclude
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```

#### Related Documentation

- [Enabling MLD Static Group Membership](#)
- [group \(Protocols IGMP\) on page 4428](#)
- [group-count \(Protocols IGMP\) on page 4429](#)
- [group-increment \(Protocols IGMP\) on page 4429](#)
- [source-count \(Protocols IGMP\) on page 4443](#)

- [source-increment \(Protocols IGMP\) on page 4443](#)
- [static \(Protocols IGMP\) on page 4445](#)

## Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

[Table 323 on page 4312](#) describes the recordable IGMP events.

**Table 323: IGMP Event Messages**

| ERRMSG Tag                  | Definition                                                     |
|-----------------------------|----------------------------------------------------------------|
| RPD_IGMP_JOIN               | Records IGMP join events.                                      |
| RPD_IGMP_LEAVE              | Records IGMP leave events.                                     |
| RPD_IGMP_ACCOUNTING_ON      | Records when IGMP accounting is enabled on an IGMP interface.  |
| RPD_IGMP_ACCOUNTING_OFF     | Records when IGMP accounting is disabled on an IGMP interface. |
| RPD_IGMP_MEMBERSHIP_TIMEOUT | Records IGMP membership timeout events.                        |

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
.*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events
```

```
*** igmp-events ***
```

```
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

**Related  
Documentation**

- [Understanding IGMP on page 4230](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6220](#)

## Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure

a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

#### Related Documentation

- [Enabling IGMP Static Group Membership on page 4305](#)

## Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                       | Description           |
|----------------------------|-----------------------|
| <b>all</b>                 | Trace all operations. |
| <b>client-notification</b> | Trace notifications.  |

| Flag                     | Description                                                                         |
|--------------------------|-------------------------------------------------------------------------------------|
| <b>general</b>           | Trace general flow.                                                                 |
| <b>group</b>             | Trace group operations.                                                             |
| <b>host-notification</b> | Trace host notifications.                                                           |
| <b>leave</b>             | Trace leave group messages (IGMPv2 only).                                           |
| <b>mtrace</b>            | Trace mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.   |
| <b>normal</b>            | Trace normal events.                                                                |
| <b>packets</b>           | Trace all IGMP packets.                                                             |
| <b>policy</b>            | Trace policy processing.                                                            |
| <b>query</b>             | Trace IGMP membership query messages, including general and group-specific queries. |
| <b>report</b>            | Trace membership report messages.                                                   |
| <b>route</b>             | Trace routing information.                                                          |
| <b>state</b>             | Trace state transitions.                                                            |
| <b>task</b>              | Trace task processing.                                                              |
| <b>timer</b>             | Trace timer processing.                                                             |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

#### Related Documentation

- [Understanding IGMP on page 4230](#)
- *Junos OS Tracing and Logging Operations* in the *Junos OS System Basics Configuration Guide*
- [mtrace on page 4512](#) in the *CLI Explorer*

## Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols igmp interface interface-name]`
- `[edit logical-systems logical-system-name protocols igmp interface interface-name]`

#### Related Documentation

- [Enabling IGMP on page 4297](#)

## IGMP Snooping

---

- [Configuring IGMP Snooping on page 4317](#)
- [Changing the IGMP Snooping Group Timeout Value on page 4318](#)
- [Example: Configuring IGMP Snooping on page 4318](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 4321](#)
- [Example: Configuring Multicast VLAN Registration on page 4322](#)



## Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).



**NOTE:** You cannot configure IGMP snooping on a secondary VLAN.

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Configure the switch to immediately remove group membership from interfaces on a VLAN when it receives a leave message through that VLAN, and have it not forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

3. Configure an interface to belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name static group
group-address
```

4. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

5. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count 4
```

6. If you want the switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

The switch uses the address that you configure as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.



**NOTE:** The `igmp-querier` statement is not supported on QFabric systems.

- Related Documentation**
- [IGMP Snooping Overview on page 4232](#)
  - [Example: Configuring IGMP Snooping on page 4318](#)
  - [Changing the IGMP Snooping Group Timeout Value on page 4318](#)
  - [Monitoring IGMP Snooping on page 4487](#)

## Changing the IGMP Snooping Group Timeout Value

The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. A switch calculates the timeout value by using the **query-interval** and **query-response-interval** values.

When you enable IGMP snooping, the **query-interval** and **query-response-interval** values are applied to all VLANs on the switch. The values are:

- **query-interval**—125 seconds
- **query-response-interval**—10 seconds

The switch automatically calculates the group timeout value for an IGMP snooping-enabled switch by multiplying the **query-interval** value by 2 (the default **robust-count** value) and then adding the **query-response-interval** value. By default, the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table:  $(125 \times 2) + 10 = 260$ .

You can modify the group timeout value by changing the **robust-count** value. For example, if you want the system to wait 510 seconds before timing groups out— $(125 \times 4) + 10 = 510$ —enter this command:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count (IGMP Snooping) 4
```

- Related Documentation**
- [Verifying the IGMP Snooping Group Timeout Value on page 4488](#)
  - [Example: Configuring IGMP Snooping on page 4318](#)
  - [Configuring IGMP Snooping on page 4317](#)

## Example: Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This example describes how to configure IGMP snooping:

- [Requirements on page 4319](#)
- [Overview and Topology on page 4319](#)
- [Configuration on page 4319](#)

Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

Overview and Topology

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

[Table 324 on page 4319](#) shows the components of the topology for this example.

Table 324: Components of the IGMP Snooping Topology

| Components                             | Settings                     |
|----------------------------------------|------------------------------|
| VLAN name                              | employee-vlan, tag 20        |
| Interfaces in employee-vlan            | ge-0/0/1, ge-0/0/2, ge-0/0/3 |
| Multicast IP address for employee-vlan | 225.100.100.100              |

Configuration

To configure basic IGMP snooping on a switch:

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into a terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:  
[edit protocols]  
user@switch# **set igmp-snooping vlan employee-vlan**
2. Configure a interface to belong to a multicast group:  
[edit protocols]

```
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 255.100.100.100
```

3. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
```

4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

**Results** Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
```

```
vlan employee-vlan {
 robust-count 4;
}
interface ge-0/0/2 {
 multicast-router-interface;
}
interface ge-0/0/3 {
 static {
 group 255.100.100.100;
 }
}
```

**Related  
Documentation**

- [IGMP Snooping Overview on page 4232](#)
- [Configuring IGMP Snooping on page 4317](#)
- [Changing the IGMP Snooping Group Timeout Value on page 4318](#)
- [Monitoring IGMP Snooping on page 4487](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1971.](#)

## Configuring Multicast VLAN Registration (CLI Procedure)

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast source VLAN (MVLAN) to still receive multicast streams from the MVLAN, allowing an MVLAN to be shared across a Layer 2 network. Hosts remain in their own VLANs for bandwidth and security reasons but are able to receive multicast streams from the MVLAN.

You can configure one or more VLANs on a switch to be MVLANs or MVR receiver VLANs. By default, MVR is not configured on EX Series switches and the QFX Series.



**NOTE:** MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.



**NOTE:** When you configure MVR, the following restrictions apply:

- You cannot enable multicast protocols on VLAN interfaces that are members of MVLANs.
- If you configure an MVLAN in proxy mode, IGMP snooping proxy mode is automatically enabled on all MVR receiver VLANs of this MVLAN. If a VLAN is an MVR receiver VLAN for multiple MVLANs, all of the MVLANs must have proxy mode enabled or all must have proxy mode disabled. You can enable proxy mode only on VLANs that are configured as MVR source VLANs and that are not configured for Q-in-Q tunneling.
- After you configure a VLAN as an MVLAN, that VLAN is no longer available for other uses.

To configure MVR:

1. Configure the VLAN named mv0 to be an MVLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding source groups
225.10.0.0/16
```

2. Configure the MVLAN mv0 to be a proxy VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 proxy source-address 10.0.0.1
```

3. Configure the VLAN named v2 to be an MVR receiver VLAN with mv0 as its source:

```
[edit protocols]
user@switch# set igmp-snooping vlan v2 data-forwarding receiver source-vlans mv0
```

4. Install forwarding entries in the MVR receiver VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding receiver install
```

### Related Documentation

- [Example: Configuring Multicast VLAN Registration on page 4322](#)

- [Understanding Multicast VLAN Registration on page 4241](#)

## Example: Configuring Multicast VLAN Registration

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, which enable the MVLAN to be shared across the Layer 2 network and eliminate the need to send duplicate multicast streams to each requesting VLAN in the network. Hosts remain in their own VLANs for bandwidth and security reasons.

This example describes how to configure MVR on EX Series switches and the QFX Series.

- [Requirements on page 4322](#)
- [Overview and Topology on page 4322](#)
- [Configuration on page 4325](#)

---

### Requirements

This example uses the following hardware and software components:

- One EX Series switch or the QFX Series
- Junos OS Release 9.6 or later for EX Series switches or Junos OS Release 12.3 or later for the QFX Series

Before you configure MVR, be sure you have:

- Configured two or more VLANs on the switch. See the task for your platform:
  - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches](#)
  - [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#) for the QFX Series
- Connected the switch to a network that can transmit IPTV multicast streams from a video server.
- Connected a host that is capable of receiving IPTV multicast streams to an interface in one of the VLANs.

---

### Overview and Topology

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast traffic can then be selectively forwarded from interfaces on the MVLAN (source ports) to hosts that are connected to interfaces (multicast receiver ports) that are not part of the multicast source VLAN. When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs

to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

You can configure MVR to operate in one of two modes: transparent mode (the default mode) or proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

In transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. [Figure 161 on page 4324](#) shows how MVR operates in transparent mode.

In proxy mode, the switch acts as a proxy for the IGMP multicast router in the MVLAN for MVR group memberships established in the MVR receiver VLANs and generates and sends IGMP packets into the MVLAN as needed. [Figure 162 on page 4325](#) shows how MVR operates in proxy mode.

This example shows how to configure MVR in both transparent mode and proxy mode on an EX Series switch or the QFX Series. The topology includes a video server that is connected to a multicast router, which in turn forwards the IPTV multicast traffic in the MVLAN to the Layer 2 network.

[Figure 161 on page 4324](#) shows the MVR topology in transparent mode. Interfaces P1 and P2 on Switch C belong to service VLAN s0 and MVLAN mv0. Interface P4 of Switch C also belongs to service VLAN s0. In the upstream direction of the network, only non-IPTV traffic is being carried in individual customer VLANs of service VLAN s0. VLAN c0 is an example of this type of customer VLAN. IPTV traffic is being carried on MVLAN mv0. If any host on any customer VLAN connected to port P4 requests an MVR stream, Switch C takes the stream from VLAN mv0 and replicates that stream onto port P4 with tag mv0. IPTV traffic, along with other network traffic, flows from port P4 out to the Digital Subscriber Line Access Multiplexer (DSLAM) D1.

Figure 161: MVR Topology in Transparent Mode





Figure 162 on page 4325 shows the MVR topology in proxy mode. Interfaces P1 and P2 on Switch C belong to MVLAN mv0 and customer VLAN c0. Interface P4 on Switch C is an access port of customer VLAN c0. In the upstream direction of the network, only non-IPTV traffic is being carried on customer VLAN c0. Any IPTV traffic requested by hosts on VLAN c0 is replicated untagged to port P4 based on streams received in MVLAN mv0. IPTV traffic flows from port P4 out to an IPTV-enabled device in Host H1. Other traffic, such as data and voice traffic, also flows from port P4 to other network devices in Host H1.



Figure 162: MVR Topology in Proxy Mode



For information on VLAN tagging, see the topic for your platform:

- [Understanding Bridging and VLANs on EX Series Switches](#)
- [“Understanding VLANs” on page 1873 on the QFX Series](#)

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit protocols igmp-snooping]** hierarchy level.

```
set vlan mv0 data-forwarding source groups 225.10.0.0/16
set vlan v2 data-forwarding receiver source-vlans mv0
set vlan v2 data-forwarding receiver install
set vlan mv0 proxy source-address 10.1.1.1
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MVR:

1. Configure VLAN mv0 to be an MVLAN:  

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 data-forwarding source groups 225.10.0.0/16
```
2. Configure VLAN v2 to be a multicast receiver VLAN with mv0 as its source:  

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver source-vlans mv0
```
3. (Optional) Install forwarding entries in the multicast receiver VLAN v2:  

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver install
```
4. (Optional) Configure MVR in proxy mode:  

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 proxy source-address 10.1.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show** command at the **[edit protocols igmp-snooping]** hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit protocols igmp-snooping]
user@switch# show
vlan mv0 {
 proxy {
 source-address 10.1.1.1;
 }
 data-forwarding {
 source {
 groups 225.10.0.0/16;
 }
 }
}
vlan v2 {
 data-forwarding {
 receiver {
 source-vlans mv0;
 install;
 }
 }
}
```

**Related Documentation**

- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 4321](#)
- [Understanding Multicast VLAN Registration on page 4241](#)

## MSDP

- [Configuring MSDP on page 4327](#)
- [Tracing MSDP Protocol Traffic on page 4328](#)
- [Configuring the Interface to Accept Traffic from a Remote Source on page 4330](#)
- [Example: Configuring MSDP on page 4331](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 4338](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 4341](#)

## Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 data-encapsulation (disable | enable);
 export [policy-names];
 group group-name {
 ... group-configuration ...
 }
 hold-time seconds;
 import [policy-names];
 local-address address;
 keep-alive seconds;
 peer address {
 ... peer-configuration ...
 }
 rib-group group-name;
 source ip-prefix </prefix-length> {
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 sa-hold-time seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 }
}
```

```

peer address {
 ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
 just following ...
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
}
peer address {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

By default, MSDP is disabled.

#### Related Documentation

- [Example: Configuring MSDP in a Routing Instance](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332](#)

## Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag | Description           |
|------|-----------------------|
| all  | Trace all operations. |

| Flag                          | Description                              |
|-------------------------------|------------------------------------------|
| <b>general</b>                | Trace general events.                    |
| <b>keepalive</b>              | Trace keepalive messages.                |
| <b>normal</b>                 | Trace normal events.                     |
| <b>packets</b>                | Trace all MSDP packets.                  |
| <b>policy</b>                 | Trace policy processing.                 |
| <b>route</b>                  | Trace MSDP changes to the routing table. |
| <b>source-active</b>          | Trace source-active packets.             |
| <b>source-active-request</b>  | Trace source-active request packets.     |
| <b>source-active-response</b> | Trace source-active response packets.    |
| <b>state</b>                  | Trace state transitions.                 |
| <b>task</b>                   | Trace task processing.                   |
| <b>timer</b>                  | Trace timer processing.                  |

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the MSDP trace file.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file msdp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set flag source-active | match 230.0.0.3
```

7. View the trace file.

```
user@host> file list /var/log
```

```
user@host> file show /var/log/msdp-trace
```

#### Related Documentation

- [Understanding MSDP on page 4235](#)
- *Junos OS Tracing and Logging Operations* in the *Junos OS System Basics Configuration Guide*

## Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.

In this sample configuration, the incoming interface (**ge-1/3/0**) is on a provider edge (PE) router on the receiver side of a multicast VPN.

To accept traffic from a remote source:

1. Edit the incoming interface.

```
[edit protocols pim interface ge-1/3/0.0]
```

```
user@host# set accept-remote-source
```

2. If the incoming interface is not the only way to reach the remote source, ensure that the best path to reach the remote source is through the incoming interface. One way to do this is to use AS path prepending on the other possible routes.

```
[edit policy-options policy-statement as-path-prepend term prepend]
```

```
user@host# set from route-filter 192.168.0.0/16 orlonger
```

```
user@host# set from route-filter 172.16.0.0/16 orlonger
```

```
user@host# set then as-path-prepend "1111"
```

Another way to do this might be to configure a static route on the receiver side PE router to the source.

4. After the configuration is committed, use the **show pim statistics** and **show msdp source** commands to verify that the interface is accepting traffic from the remote source.

- Related Documentation**
- *Example: Allowing MBGP MVPN Remote Sources*
  - *Prepending AS Numbers to BGP AS Paths* in the *Routing Policy Configuration Guide*
  - [show msdp source on page 4547](#) in the *CLI Explorer*
  - [show pim statistics on page 4615](#) in the *CLI Explorer*

## Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
 interface-routes {
 rib-group ifrg;
 }
 rib-groups {
 ifrg {
 import-rib [inet.0 inet.2];
 }
 mcrg {
 export-rib inet.2;
 import-rib inet.2;
 }
 }
}
protocols {
 bgp {
 group lab {
 type internal;
 family any;
 neighbor 192.168.6.18 {
 local-address 192.168.6.17;
 }
 }
 }
}
pim {
 dense-groups {
 224.0.1.39/32;
 224.0.1.40/32;
 }
 rib-group mcrg;
 rp {
 local {
 address 192.168.1.1;
 }
 }
 interface all {
 mode sparse-dense;
 version 1;
 }
}
msdp {
 rib-group mcrg;
```

```
group lab {
 peer 192.168.6.18 {
 local-address 192.168.6.17;
 }
}
}
```

## Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

- [Requirements on page 4332](#)
- [Overview on page 4332](#)
- [Configuration on page 4336](#)
- [Verification on page 4338](#)

### Requirements

---

Before you begin:

- Configure the router interfaces. See the *Junos® OS Network Interfaces*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Configuration Guide*.
- Enable PIM sparse mode. See “[PIM Overview](#)” on page 4209.
- Configure the router as a PIM sparse-mode RP. See “[Configuring Local PIM RPs](#)” on page 4264.

### Overview

---

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early discard (RED) to drop some but not all MSDP active source messages. Beginning with Junos OS 12.2, you can optionally configure a warning threshold so the device can log warning messages in the system log when a certain number of source-active messages have been received. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of source-active messages have been received. These log messages convey when the configured message limit has been exceeded, when the configured warning threshold has been exceeded, and when the number of messages drop below the configured warning threshold.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.



By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.



**NOTE:** The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

The warning threshold is a percentage of maximum number of MSDP source-active messages received, so you must configure the source-active message limit to configure a warning threshold. The range for the warning threshold is 1 through 100 percent. You can further specify the amount of time (in seconds) between the log messages. The range is 6 through 32,767 seconds.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the

source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.



**CAUTION:** When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.




**NOTE:** An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the **default-peer** statement.

Table 325 on page 4334 explains how flooding is handled by peers in this example. Figure 163 on page 4335 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

**Table 325: Source-Active Message Flooding Explanation**

| Source-Active Message Received From | Source-Active Message Flooded To                     | Source-Active Message Not Flooded To |
|-------------------------------------|------------------------------------------------------|--------------------------------------|
| Peer 21                             | Peer 11, Peer 12, Peer 13, Peer 31, Peer 32          | Peer 22                              |
| Peer 11                             | Peer 21, Peer 22, Peer 31, Peer 32                   | Peer 12, Peer 13                     |
| Peer 31                             | Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32 | –                                    |

Figure 163: Source-Active Message Flooding



This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **active-source-limit log-warning 80**—(Optional) Applies a warning threshold of 80 percent. In this example, the active source maximum is 10,000, so the device will start logging warning messages once it receives 8,000 active source messages.
- **active-source-limit log-interval 20**—(Optional) Applies a 20 second waiting period between system log messages.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation,

multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the MSDP-group group are mesh group members.
- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group MSDP-group.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

### Configuration

---

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp active-source-limit log-warning 80
set protocols msdp active-source-limit log-interval 20
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
```

#### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
```

```

user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500

```

3. (Optional) Configure the threshold at which warning messages are logged and the amount of time between log messages.

```

[edit protocols msdp]
user@host# set active-source-limit log-warning 80
user@host# set active-source-limit log-interval 20

```

4. Configure the mesh group.

```

[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3

```

5. If you are done configuring the device, commit the configuration.

```

[edit routing-instances]
user@host# commit

```

## Results

Confirm your configuration by entering the **show protocols** command.

```

user@host# show protocols
msdp {
 data-encapsulation disable;
 active-source-limit {
 maximum 10000;
 log-warning 80;
 log-interval 20;
 }
 peer 10.0.0.1 {
 active-source-limit {
 maximum 5000;
 threshold 4000;
 }
 }
 source 10.1.0.0/16 {
 active-source-limit {
 maximum 500;
 }
 }
 group MSDP-group {
 mode mesh-group;
 local-address 10.1.2.3;
 peer 10.10.10.10 {
 active-source-limit {
 maximum 7500;
 }
 }
 }
}

```

## Verification

---

To verify the configuration, run the following commands:

- `show msdp source-active`
- `show msdp statistics`

### Related Documentation

- *Example: Configuring MSDP in a Routing Instance*
- [Filtering MSDP SA Messages on page 4236](#)
- *Configuring RED Drop Profiles in the Junos OS Class of Service Configuration Guide*
- [Configuring Local PIM RPs on page 4264](#)

## Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the

**primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32;
 primary;
 address 198.58.3.253/32;
 }
 }
 }
}

```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```

protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
 }
}

```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```

protocols {
 msdp {
 peer 198.58.3.250 {
 local-address address 198.58.3.254;
 }
 }
}

```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
 lo0 {
 description "PIM RP";
 unit 0 {
 family inet {
 address 198.58.3.254/32 {
 primary;
 }
 address 198.58.3.253/32;
 }
 }
 }
}
```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
 pim {
 rp {
 local {
 family inet {
 address 198.58.3.253;
 anycast-pim {
 rp-set {
 address 198.58.3.240;
 address 198.58.3.241 forward-msdp-sa;
 }
 local-address 198.58.3.254; #If not configured, use lo0 primary
 }
 }
 }
 }
 }
}
```



```

 }
 }
}
interface all {
 mode sparse;
 version 2;
}
interface fxp0.0 {
 disable;
}
}
}

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```

protocols {
 pim {
 rp {
 static {
 address 198.58.3.253 {
 version 2;
 }
 }
 }
 }
}

```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```

protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
}

```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
 pim {
 rp {
 local {
 family inet;
 address 198.58.3.253;
 }
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
 }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
 msdp {
 peer 198.58.3.250 {
 local-address 198.58.3.254;
 }
 }
}
```

---


## Source-Specific Multicast

- [Example: Configuring PIM SSM on a Network on page 4342](#)
- [Example: Configuring an SSM-Only Domain on page 4344](#)
- [Example: Configuring SSM Mapping on page 4344](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347](#)
- [Example: Configuring SSM Maps for Different Groups to Different Sources on page 4350](#)

### Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 164 on page 4343](#).

Figure 164: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```

user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable

```



**NOTE:** When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```

user@router1> show configuration protocol igmp

[edit protocols igmp]
interface all {
 version 3;
}
interface fxp0.0 {
 disable;
}

```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

```

user@router1> show igmp interface
Interface State Querier Timeout Version Groups
fe-0/0/0.0 Up 198.58.3.245 213 3 0
fe-0/0/1.0 Up 198.58.3.241 220 3 0
fe-0/0/2.0 Up 198.58.3.237 218 3 0
Configured Parameters:
IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550

```

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```

user@router2> show pim join extensive
232.1.1.1 10.4.1.2 sparse
Upstream interface: fe-1/1/3.0

```

```

Upstream State: Local Source
Keepalive timeout: 209
Downstream Neighbors:
 Interface: so-1/0/2.0
 10.10.71.1 State: Join Flags: S Timeout: 209

```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```

user@router1> show pim join extensive
232.1.1.1 10.4.1.2 sparse
Upstream interface: so-1/0/2.0
Upstream State: Join to Source
Keepalive timeout: 209
Downstream Neighbors:
 Interface: fe-0/2/3.0
 10.3.1.1 State: Join Flags: S Timeout: Infinity

```

## Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface interface-name]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```

[edit]
protocols {
 pim {
 interface all {
 mode sparse;
 version 2;
 }
 interface fxp0.0 {
 disable;
 }
 }
 igmp {
 interface fe-0/1/2 {
 version 3;
 }
 }
}

```

## Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```
user@router1# set policy-options policy-statement ssm-policy-example term A from
route-filter 232.1.1.1/32 exact
user@router1# set policy-options policy-statement ssm-policy-example term A then
accept
user@router1# set policy-options policy-statement ssm-policy-example term B from
route-filter ff35::1/128 exact
user@router1# set policy-options policy-statement ssm-policy-example term B then
accept
```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options

[edit policy-options]
policy-statement ssm-policy-example {
 term A {
 from {
 route-filter 232.1.1.1/32 exact;
 }
 then accept;
 }
 term B {
 from {
 route-filter ff35::1/128 exact;
 }
 then accept;
 }
 then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options

[edit routing-options]
multicast {
 ssm-map ssm-map-ipv6-example {
 policy ssm-policy-example;
 source [fec0::1 fec0::12];
 }
 ssm-map ssm-map-ipv4-example {
 policy ssm-policy-example;
 source [10.10.10.4 192.168.43.66];
 }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol

[edit protocols]
igmp {
 interface fe-0/1/0.0 {
 ssm-map ssm-map-ipv4-example;
 }
}
mld {
 interface fe-0/1/1.0 {
 ssm-map ssm-map-ipv6-example;
 }
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```

user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
 Querier: 192.168.224.28
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map: ssm-map-ipv4-example

user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
 Querier: fec0:0:0:1::12
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map: ssm-map-ipv6-example

```

## Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

- [Requirements on page 4347](#)
- [Overview on page 4347](#)
- [Configuration on page 4349](#)
- [Verification on page 4350](#)

### Requirements


Before you begin, configure the router interfaces. See the *Junos® OS Network Interfaces*.

### Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 165 on page 4348](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 165 on page 4348](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 165: Receiver Sends Messages to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 166 on page 4348](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 166: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 167 on page 4348](#)).


Figure 167: (S,G) State Is Built Between the Source and the Receiver



SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in [Figure 168 on page 4348](#).

Figure 168: Simple RPF Topology





## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [232.0.0.0/8 239.0.0.0/8]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

## Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
 area 0.0.0.0 {
 interface fxp0.0 {
 disable;
 }
 interface all;
 }
}
pim {
 rp {
 local {
 address 10.255.72.46;
 group-ranges {
 239.0.0.0/24;
 }
 }
 }
 interface fe-1/0/0.0 {
 mode sparse;
 }
 interface lo0.0 {
 mode sparse;
 }
}

user@host# show routing-options
multicast {
 ssm-groups [232.0.0.0/8 239.0.0.0/8];
 asm-override-ssm;
}
```

---

### Verification

To verify the configuration, run the following commands:

- [show igmp group](#)
- [show igmp statistics](#)
- [show pim join](#)

### Related Documentation

- [Source-Specific Multicast Groups Overview on page 4237](#)

## Example: Configuring SSM Maps for Different Groups to Different Sources

- [Multiple SSM Maps and Groups for Interfaces on page 4350](#)
- [Example: Configuring Multiple SSM Maps Per Interface on page 4351](#)

---

### Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

- [Requirements on page 4351](#)
- [Overview on page 4351](#)
- [Configuration on page 4351](#)
- [Verification on page 4353](#)

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, POLICY-ipv4-example1, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

| Routing Policy Name         | Multicast Group Join Messages for a Route Filter at This Destination Address | Multicast Source Addresses   |
|-----------------------------|------------------------------------------------------------------------------|------------------------------|
| POLICY-ipv4-example1 term 1 | 232.1.1.1                                                                    | 10.10.10.4,<br>192.168.43.66 |
| POLICY-ipv4-example1 term 2 | 232.1.1.2                                                                    | 10.10.10.5,<br>192.168.43.67 |

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure this example, perform the following task:

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the [edit] hierarchy level.

set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter 232.1.1.1/32 exact  
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source 10.10.10.4  
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source 192.168.43.66  
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept  
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter 232.1.1.2/32 exact  
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source 10.10.10.5

```

set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1

```

### Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```

[edit policy-options policy-statement POLICY-ipv4-example1 term 1]
user@host# set from route-filter 232.1.1.1/32 exact
user@host# set then ssm-source 10.10.10.4
user@host# set then ssm-source 192.168.43.66
user@host# set then accept

```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```

[edit policy-options policy-statement POLICY-ipv4-example1 term 2]
user@host# set from route-filter 232.1.1.2/32 exact
user@host# set then ssm-source 10.10.10.5
user@host# set then ssm-source 192.168.43.67
user@host# set then accept

```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```

[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1

```

### Results

After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

user@host#> show policy-options
policy-statement POLICY-ipv4-example1 {
 term 1 {
 from {
 route-filter 232.1.1.1/32 exact;
 }
 then {
 ssm-source [10.10.10.4 192.168.43.66];
 accept;
 }
 }
 term 2 {
 from {
 route-filter 232.1.1.2/32 exact;
 }
 then {
 ssm-source [10.10.10.5 192.168.43.67];
 accept;
 }
 }
}

```

```

}

user@host# show protocols
igmp {
 interface fe-0/1/0.0 {
 ssm-map-policy POLICY-ipv4-example1;
 }
}

```

### Verification

Confirm that the configuration is working properly.

- [Displaying Information About IGMP-Enabled Interfaces on page 4353](#)
- [Displaying the PIM Groups on page 4353](#)
- [Displaying the Entries in the IP Multicast Forwarding Table on page 4353](#)

### Displaying Information About IGMP-Enabled Interfaces

**Purpose** Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

**Action** Use the [show igmp interface](#) operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```

user@host> show igmp interface
Interface: fe-0/1/0.0
 Querier: 10.111.30.1
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map Policy: POLICY-ipv4-example1;

```

```

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

```

```

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

The command output displays the name of IGMP logical interface (fe-0/1/0.0), the address of the routing device that has been elected to send membership queries and group information.

### Displaying the PIM Groups

**Purpose** Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

**Action** Use the [show pim join extensive 232.1.1.1](#) operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

### Displaying the Entries in the IP Multicast Forwarding Table

**Purpose** Verify that the IP multicast forwarding table displays the mroute state.

**Action** Use the [show multicast route extensive](#) operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

**Related Documentation**

- [Example: Configuring Source-Specific Multicast](#)
- [Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs](#)

---

## PIM Configuration Statements

---

- [address \(Anycast RPs\) on page 4356](#)
- [address \(Local RPs\) on page 4357](#)
- [address \(Static RPs\) on page 4358](#)
- [algorithm on page 4359](#)
- [anycast-pim on page 4360](#)
- [assert-timeout on page 4361](#)
- [authentication \(Protocols PIM\) on page 4362](#)
- [bfd-liveness-detection \(Protocols PIM\) on page 4363](#)
- [bootstrap on page 4364](#)
- [bootstrap-export on page 4365](#)
- [bootstrap-import on page 4366](#)
- [bootstrap-priority on page 4367](#)
- [detection-time \(BFD for PIM\) on page 4368](#)
- [disable \(PIM\) on page 4369](#)
- [dr-election-on-p2p on page 4370](#)
- [dr-register-policy on page 4370](#)
- [embedded-rp on page 4371](#)
- [export \(Protocols PIM Bootstrap\) on page 4372](#)
- [export \(Protocols PIM\) on page 4372](#)
- [family \(Bootstrap\) on page 4373](#)
- [family \(Protocols PIM\) on page 4374](#)
- [family \(Local RP\) on page 4375](#)
- [group \(RPF Selection\) on page 4376](#)
- [group-ranges on page 4377](#)
- [hello-interval \(Protocols PIM\) on page 4378](#)
- [hold-time \(Protocols PIM\) on page 4379](#)
- [import \(Protocols PIM Bootstrap\) on page 4380](#)
- [import \(Protocols PIM\) on page 4381](#)
- [infinity on page 4382](#)

- [interface](#) on page 4383
- [join-load-balance](#) on page 4384
- [join-prune-timeout](#) on page 4385
- [key-chain \(Protocols PIM\)](#) on page 4386
- [local](#) on page 4387
- [local-address \(Protocols PIM\)](#) on page 4388
- [loose-check](#) on page 4389
- [maximum-rps](#) on page 4390
- [minimum-interval \(PIM BFD Liveness Detection\)](#) on page 4391
- [minimum-interval \(PIM BFD Transmit Interval\)](#) on page 4392
- [minimum-receive-interval](#) on page 4393
- [mode \(Protocols PIM\)](#) on page 4394
- [multiplier](#) on page 4394
- [neighbor-policy](#) on page 4395
- [next-hop \(PIM RPF Selection\)](#) on page 4395
- [no-adaptation \(PIM BFD Liveness Detection\)](#) on page 4396
- [override-interval](#) on page 4397
- [pim](#) on page 4398
- [prefix-list \(PIM RPF Selection\)](#) on page 4401
- [priority \(Bootstrap\)](#) on page 4402
- [priority \(PIM Interfaces\)](#) on page 4403
- [priority \(PIM RPs\)](#) on page 4404
- [propagation-delay](#) on page 4405
- [reset-tracking-bit](#) on page 4406
- [rib-group \(Protocols PIM\)](#) on page 4407
- [rp](#) on page 4408
- [rp-register-policy](#) on page 4410
- [rp-set](#) on page 4411
- [rpf-selection](#) on page 4412
- [source \(PIM RPF Selection\)](#) on page 4413
- [spt-threshold](#) on page 4414
- [static \(Protocols PIM\)](#) on page 4415
- [threshold \(PIM BFD Detection Time\)](#) on page 4416
- [threshold \(PIM BFD Transmit Interval\)](#) on page 4417
- [transmit-interval \(PIM BFD Liveness Detection\)](#) on page 4418
- [traceoptions \(Protocols PIM\)](#) on page 4419
- [version \(BFD\)](#) on page 4422

- [version \(PIM\) on page 4423](#)
- [wildcard-source \(PIM RPF Selection\) on page 4424](#)

---

## address (Anycast RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address <i>address</i> &lt;forward-msdp-sa&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim</b></code><br><code><b>rp-set</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>address</i></b> —RP address in an RP set.<br><br><b><i>forward-msdp-sa</i></b> —(Optional) Forward MSDP SAs to this address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



## address (Local RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address <i>address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the local rendezvous point (RP) address.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>address</i></b> —Local RP address.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4264</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                     |

## address (Static RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>address address {<br/>  group-ranges {<br/>    destination-ip-prefix &lt;/prefix-length&gt;;<br/>  }<br/>  override;<br/>  version version;<br/>}</pre>                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name protocols pim rp static],<br/>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols<br/>  pim rp static],<br/>[edit protocols pim static],<br/>[edit routing-instances routing-instance-name protocols pim rp static]</pre>                                      |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                         |
| <b>Description</b>              | <p>Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> |
| <b>Options</b>                  | <p><b>address</b>—Static RP address.</p> <p><b>Default:</b> 224.0.0.0/4</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4266</a></li></ul>                                                                                                                                                                                                     |

## algorithm

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>algorithm <i>algorithm-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify the algorithm to use for BFD authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b><i>algorithm-name</i></b>—Name of algorithm to use for BFD authentication:</p> <ul style="list-style-type: none"> <li>• <b>simple-password</b>—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.</li> <li>• <b>keyed-md5</b>—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.</li> <li>• <b>meticulous-keyed-md5</b>—Meticulous keyed Message Digest 5 hash algorithm.</li> <li>• <b>keyed-sha-1</b>—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.</li> <li>• <b>meticulous-keyed-sha-1</b>—Meticulous keyed Secure Hash Algorithm I.</li> </ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional Forwarding Detection Authentication for PIM</i></li> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4292</a></li> <li>• <a href="#">authentication on page 4362</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## anycast-pim

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>anycast-pim {<br/>  rp-set {<br/>    address address &lt;forward-msdp-sa&gt;;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp local family</b> (inet   inet6)],<br>[edit protocols pim <b>rp local family</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure properties for anycast RP using PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4267</a></li></ul>                                                                                                                                                                                                                                                                                                           |

## assert-timeout

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>assert-timeout <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time for routing device to wait before another assert message cycle.<br><b>Range:</b> 5 through 210 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the PIM Assert Timeout on page 4284</a></li> </ul>                                                                                                                                                                                                                                                                                                                                             |

## authentication (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>authentication {<br/>  algorithm <i>algorithm-name</i>;<br/>  key-chain <i>key-chain-name</i>;<br/>  loose-check;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 4292</a></li><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li><li>• <a href="#">key-chain (Protocols PIM) on page 4386</a></li><li>• <a href="#">loose-check on page 4389</a></li></ul> |

## bfd-liveness-detection (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> bfd-liveness-detection {   authentication {     algorithm <i>algorithm-name</i>;     key-chain <i>key-chain-name</i>;     loose-check;   }   detection-time {     threshold <i>milliseconds</i>;   }   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   version (0   1   automatic); } </pre> |
| <b>Hierarchy Level</b>          | <pre> [edit protocols pim interface <i>interface-name</i>], [edit protocols pim interface <i>interface-name</i> <i>family</i> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <i>family</i> (inet   inet6)] </pre>                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p><b>authentication</b> option introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4290</a></li> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4292</a></li> </ul>                                                                                                                                                                                                                                                                                                        |

## bootstrap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bootstrap {<br/>    family (inet   inet6) {<br/>        export [ <i>policy-names</i> ];<br/>        import [ <i>policy-names</i> ];<br/>        priority <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                                                |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    pim <i>rp</i>],<br/>[edit protocols pim <i>rp</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                         |
| <b>Description</b>              | <p>Configure parameters to control bootstrap routers and messages.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li></ul>                                                                                                                                                                     |



## bootstrap-export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit protocols pim <a href="#">rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-import on page 4366</a></li> </ul>                                                                                                                                          |

## bootstrap-import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a> ],<br>[edit protocols pim <a href="#">rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-export on page 4365</a></li></ul>                                                                                                                               |

## bootstrap-priority

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>],</p> <p>[edit protocols pim <i>rp</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                  |
| <b>Description</b>              | Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.                                                                                                                                                        |
| <b>Options</b>                  | <p><i>number</i>—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 0</p>                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> </ul>                                                                                                                                                                                                                                       |

## detection-time (BFD for PIM)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>detection-time {<br/>    threshold milliseconds;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Release Information      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description              | <p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the <b>clear bfd adaptation</b> command to return BFD interval timers to their configured values. The <b>clear bfd adaptation</b> command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li><li>• <a href="#">threshold on page 4416</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## disable (PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> protocols pim <b>family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim <b>rp local family</b> (inet   inet6)], [edit protocols pim], [edit protocols pim <b>family</b> (inet   inet6)], [edit protocols pim interface <i>interface-name</i>], [edit protocols pim <b>rp local family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> statement extended to the <b>[family]</b> hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Explicitly disable PIM at the protocol, interface or family hierarchy levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Disabling PIM on page 4250</li> <li>disable (PIM Graceful Restart)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## dr-election-on-p2p

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-election-on-p2p;                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                          |
| <b>Description</b>              | Enable PIM designated router (DR) election on point-to-point (P2P) links.                                                                                                                                                                                                           |
| <b>Default</b>                  | No PIM DR election is performed on point-to-point links.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Designated Router Election on Point-to-Point Links on page 4254</a></li></ul>                                                                                                                                   |

## dr-register-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-register-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ],<br>[edit protocols pim <i>rp</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to control outgoing PIM register messages.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 4282</a></li><li>• <a href="#">rp-register-policy on page 4410</a></li></ul>                                                                                                                                |

## embedded-rp

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> embedded-rp {   group-ranges {     destination-ip-prefix &lt;/prefix-length&gt;;   }   maximum-rps limit; } </pre>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                  |
| <b>Description</b>              | <p>Configure properties for embedded IP version 6 (IPv6) RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Embedded RP for IPv6</i></li> </ul>                                                                                                                                                                                                                                                |

## export (Protocols PIM Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],<br>[edit protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4</a></li><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6</a></li><li>• <a href="#">import (Protocols PIM Bootstrap) on page 4380</a></li></ul>                                                                                                                                                                                                                         |

## export (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.                                                       |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Outgoing PIM Join Messages on page 4280</a></li></ul>                                                                                                                                                                 |



## family (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {     export [ <i>policy-names</i> ];     import [ <i>policy-names</i> ];     priority <i>number</i>; }</pre>                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit protocols pim <b>rp bootstrap</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                              |
| <b>Description</b>              | Configure which IP protocol type bootstrap properties to apply.                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> </ul>                                                                                                                                                                                                       |

## family (Protocols PIM)

---

|                              |                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | family (inet   inet6) {<br>disable;<br>}                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>       | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                   |
| <b>Description</b>           | Enable the PIM protocol for the specified family.                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>               | <b>inet</b> —Enable the PIM protocol for the IP version 4 (IPv4) address family.<br><br><b>inet6</b> —Enable the PIM protocol for the IP version 6 (IPv6) address family.<br><br>The remaining statement is explained separately.                                                                                                                                                            |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Disabling PIM on page 4250</a></li><li>• <i>disable (PIM Graceful Restart)</i></li><li>• <a href="#">disable (PIM) on page 4369</a></li></ul>                                                                                                                                                                                            |

## family (Local RP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> family (inet   inet6) {     disable;     address address;     anycast-pim {         local-address address;         rp-set {             address address &lt;forward-msdp-sa&gt;;         }     }     group-ranges {         destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number; } </pre>               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>],</p> <p>[edit protocols pim <b>rp local</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | Configure which IP protocol type local RP properties to apply.                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 4264</a></li> </ul>                                                                                                                                                                                                                                                               |

## group (RPF Selection)

---

|                                 |                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group group-address{   source source-address{     next-hop next-hop-address;   }   wildcard-source {     next-hop next-hop-address;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]                                                               |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                     |
| <b>Description</b>              | Configure the PIM group address for which you configure RPF selection <a href="#">group (RPF Selection)</a> .                                        |
| <b>Default</b>                  | By default, PIM RPF selection is not configured.                                                                                                     |
| <b>Options</b>                  | <b>group-address</b> —PIM group address for which you configure RPF selection.                                                                       |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                      |

## group-ranges

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp embedded-rp</b>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                  | The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>destination-ip-prefix&lt;/prefix-length&gt;</i></b> —Addresses or address ranges for which this routing device can be an RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4264</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Configuring PIM Embedded RP for IPv6</i> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## hello-interval (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hello-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify how often the routing device sends PIM hello packets out of an interface.                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time between PIM hello packets.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hold-time on page 4379</a></li><li>• <a href="#">Modifying the PIM Hello Interval on page 4246</a></li></ul>                                                                                                                                                                                                                                                           |

## hold-time (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 150 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4264</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |

## import (Protocols PIM Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit protocols pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">export (Protocols PIM Bootstrap) on page 4372</a></li></ul>                                                                                                                                                                              |



## import (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from PIM. Use the <b>import</b> statement to filter PIM join messages and prevent them from entering the network.                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Incoming PIM Join Messages on page 4281</a></li> </ul>                                                                                                                                                                                  |

## infinity

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>infinity [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">spt-threshold</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">spt-threshold</a> ],<br>[edit protocols pim <a href="#">spt-threshold</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">spt-threshold</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the <b>infinity</b> statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 4287</a></li></ul>                                                                                                                                                                                                                                                                                |

## interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> <b>interface</b> (all   <i>interface-name</i>) {     <b>disable</b>;     family (inet   inet6) {         <b>disable</b>;     }     <b>hello-interval</b> <i>seconds</i>;     mode (dense   sparse   sparse-dense);     <b>neighbor-policy</b> [ <i>policy-names</i> ];     <b>override-interval</b> <i>milliseconds</i>;     <b>priority</b> <i>number</i>;     <b>propagation-delay</b> <i>milliseconds</i>;     <b>reset-tracking-bit</b>;     <b>version</b> <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Enable PIM on an interface and configure interface-specific properties.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">PIM on Aggregated Interfaces on page 4212</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                  |

## join-load-balance

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>join-load-balance {<br/>    automatic;<br/>}</pre>                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                          |
| <b>Description</b>              | Enable load balancing of PIM join messages across interfaces and routing devices.                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>automatic</b> —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <a href="#">Configuring PIM Join Load Balancing on page 4256</a></li><li>• <i>clear pim join-distribution</i> in the <a href="#">CLI Explorer</a></li></ul> |

## join-prune-timeout

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | join-prune-timeout <i>seconds</i> ;                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>              | Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Number of seconds to wait for the periodic join message to arrive.<br><b>Range:</b> 210 through 240 seconds<br><b>Default:</b> 210 seconds                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the Join State Timeout on page 4259</a></li> </ul>                                                                                                                                                                   |

## key-chain (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>key-chain <i>key-chain-name</i>;</code>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication]    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement modified in Junos OS Release 12.2 to include <b>family</b> in the hierarchy level.                                                                      |
| <b>Description</b>              | Specify the security keychain to use for BFD authentication.                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Name of the security keychain to use for BFD authentication. The name is a unique integer between <b>0</b> and <b>63</b> . This must match one of the keychains in the <b>authentication-key-chains</b> statement at the [edit security] hierarchy level. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 4292</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM authentication on page 4362</a></li></ul>                                     |

## local

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> local {   disable;   address address;   family (inet   inet6) {     disable;     address address;     anycast-pim {       local-address address;       rp-set {         address address &lt;forward-msdp-sa&gt;;       }     }     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number;   }   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   hold-time seconds;   override;   priority number; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the routing device's RP properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 4264</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                          |

## local-address (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>address</b> —Anycast RP IPv4 or IPv6 address, depending on <b>family</b> configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4267</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## loose-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loose-check;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4292</a></li> <li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM authentication on page 4362</a></li> </ul>                                                                                                                                                                                                                                                                                                                               |

## maximum-rps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-rps <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit protocols pim <a href="#">rp embedded-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Limit the number of RPs that the routing device acknowledges.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>limit</i> —Number of RPs.<br><b>Range:</b> 1 through 500<br><b>Default:</b> 100                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Embedded RP for IPv6</i></li></ul>                                                                                                                                                                                                                                                                                                                       |

## minimum-interval (PIM BFD Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <b>transmit-interval</b> <b>minimum-interval</b> and <b>minimum-receive-interval</b> statements. |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Minimum transmit and receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4290</a></li> </ul>                                                                                                                                                                                                                                                                                                                   |

## minimum-interval (PIM BFD Transmit Interval)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                    |
| <b>Description</b>         | Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection] hierarchy level. |
| <b>Options</b>             | <i>milliseconds</i> —Minimum transmit interval value.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                                                                                            |



**NOTE:** The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li><li>• <a href="#">minimum-interval on page 4391</a></li><li>• <a href="#">threshold on page 4417</a></li></ul> |

## minimum-receive-interval

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-receive-interval <i>milliseconds</i> ;                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                |
| <b>Description</b>              | Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ] hierarchy level. |
| <b>Options</b>                  | <i>milliseconds</i> —Minimum receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4290</a></li> </ul>                                                                                                                                                                                                                                                                                                  |

## mode (Protocols PIM)

---

|                            |                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | mode (dense   sparse   sparse-dense);                                                                                                                          |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                              |
| <b>Description</b>         | Configure PIM to operate in sparse, dense, or sparse-dense mode.                                                                                               |



**NOTE:** The QFX Series does not support dense or sparse-dense mode.

---

|                                 |                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>dense</b> —Operate in dense mode.<br><b>sparse</b> —Operate in sparse mode.<br><b>sparse-dense</b> —Operate in sparse-dense mode.<br><b>Default:</b> sparse |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                            |

## multiplier

---

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multiplier <i>number</i> ;                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                 |
| <b>Description</b>              | Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.                                                                                                |
| <b>Options</b>                  | <b>number</b> —Number of hello packets.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li></ul>                                                                                                                     |

## neighbor-policy

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>neighbor-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply a PIM interface-level policy to filter neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>policy-name</i> —Name of the policy that filters neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Interface-Level PIM Neighbor Policies on page 4279</a></li> </ul>                                                                                                                                                                                                                                                                                      |

## next-hop (PIM RPF Selection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>next-hop <i>next-hop-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source] |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>next-hop-address</i> —Specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM RPF Selection</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## no-adaptation (PIM BFD Liveness Detection)

---

|                                 |                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-adaptation;                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li></ul>                                                                                            |



## override-interval

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>override-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],<br/>         [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>           pim interface <i>interface-name</i>],<br/>         [edit protocols pim],<br/>         [edit protocols pim interface <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols pim]<br/>         [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p>This is a random timer with a value in milliseconds.</p> <p><b>Range:</b> 0 through maximum override value</p> <p><b>Default:</b> 2000 milliseconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Enabling Join Suppression on page 4259</a></li> <li>• <a href="#">propagation-delay on page 4405</a></li> <li>• <a href="#">reset-tracking-bit on page 4406</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## pim

---

```
Syntax pim {
 disable;
 assert-timeout seconds;
 dense-groups {
 addresses;
 }
 dr-election-on-p2p;
 export;
 family (inet | inet6) {
 disable;
 }
 graceful-restart {
 disable;
 restart-duration seconds;
 }
 import [policy-names];
 interface interface-name {
 accept-remote-source;
 disable;
 family (inet | inet6) {
 disable;
 }
 hello-interval seconds;
 mode (dense | sparse | sparse-dense);
 neighbor-policy [policy-names];
 override-interval milliseconds;
 priority number;
 propagation-delay milliseconds;
 reset-tracking-bit;
 version version;
 }
 join-load-balance;
 join-prune-timeout;
 nonstop-routing;
 override-interval milliseconds;
 propagation-delay milliseconds;
 reset-tracking-bit;
 rib-group group-name;
 rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 }
 bootstrap-import [policy-names];
 bootstrap-export [policy-names];
```

```

bootstrap-priority number;
dr-register-policy [policy-names];
embedded-rp {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 maximum-rps limit;
}
local {
 family (inet | inet6) {
 address address;
 anycast-pim {
 disable;
 rp-set {
 address address <forward-msdp-sa>;
 }
 local-address address;
 }
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
}
rp-register-policy [policy-names];
spt-threshold {
 infinity [policy-names];
}
static {
 address address {
 group-ranges {
 version version;
 destination-ip-prefix </prefix-length>;
 }
 }
}
rpf-selection {
 group group-address {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
 prefix-list prefix-list-addresses {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
}
traceoptions {

```

```
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 tunnel-devices [mt-fpc/pic/port];
}
```

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>family</b> statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.       |
| <b>Description</b>              | Enable PIM on the routing device.<br><br>The statements are explained separately.                                                                                                                                                                                   |
| <b>Default</b>                  | PIM is disabled on the routing device.                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |

## prefix-list (PIM RPF Selection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> prefix-list <i>prefix-list-addresses</i> {   source <i>source-address</i> {     next-hop <i>next-hop-address</i>;   }   wildcard-source {     next-hop <i>next-hop-address</i>;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | (Optional) Configure a list of prefixes (addresses) for multiple PIM groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>prefix-list-addresses</i></b>—List of prefixes (addresses) for multiple PIM groups.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## priority (Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit protocols pim <b>rp bootstrap</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the bootstrap router.                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>number</b> —Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.<br><b>Range:</b> 0 through a 32-bit number<br><b>Default:</b> 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-priority on page 4367</a></li></ul>                                                                                                                                                                                            |

## priority (PIM Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the designated router.                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>number</i></b> —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.<br><b>Range:</b> 0 through a 32-bit number<br><b>Default:</b> 1                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Interface Priority for PIM Designated Router Selection on page 4253</a></li> </ul>                                                                                                                                                                                                                                                                        |

## priority (PIM RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim</code><br><code>rp bidirectional address <i>address</i>],</code><br><code>[edit protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit protocols pim <b>rp local family</b> (inet   inet6)],</code><br><code>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | For PIM-SM, configure this routing device's priority for becoming an RP.<br><br>For bidirectional PIM, configure this RP address' priority for becoming an RP.<br><br>The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.                                                                                                                                                                                       |
| <b>Options</b>                  | <b><i>number</i></b> —Priority for becoming an RP. A lower value corresponds to a higher priority.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Local PIM RPs on page 4264</a> in the <i>Multicast Protocols Configuration Guide</i></li><li>• <i>Example: Configuring Bidirectional PIM</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



## propagation-delay

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>propagation-delay <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit protocols pim],</code><br><code>[edit protocols pim interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols pim],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>pim interface <i>interface-name</i>]</code> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Set a delay for implementing a PIM prune message on the upstream router on a multicast network for which join suppression has been enabled. The router waits for the prune pending period to detect whether a join message is currently being suppressed by another router.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>milliseconds</i></b>—Interval for the prune pending timer, which is the sum of the <b>propagation-delay</b> value and the <b>override-interval</b> value.</p> <p><b>Range:</b> 250 through 2000 milliseconds</p> <p><b>Default:</b> 500 milliseconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Enabling Join Suppression on page 4259</a></li> <li>• <a href="#">override-interval on page 4397</a></li> <li>• <a href="#">reset-tracking-bit on page 4406</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## reset-tracking-bit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | reset-tracking-bit;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols pim],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ]                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds ( $1.1 \times$ periodic through $1.4 \times$ periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression on page 4259</a></li><li>• <a href="#">override-interval on page 4397</a></li><li>• <a href="#">propagation-delay on page 4405</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                |

## rib-group (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> rib-group {     inet <i>group-name</i>;     inet6 <i>group-name</i>; } </pre>                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Associate a routing table group with PIM.                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>table-name</i> —Name of the routing table. The name must be one that you defined with the <b>rib-groups</b> statement at the <b>[edit routing-options]</b> hierarchy level.                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Example: Configuring a Dedicated PIM RPF Routing Table</li> </ul>                                                                                                                                                                            |

## rp

---

```
Syntax rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 bootstrap-export [policy-names];
 bootstrap-import [policy-names];
 bootstrap-priority number;
 dr-register-policy [policy-names];
 embedded-rp {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 maximum-rps limit;
 }
 group-rp-mapping {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
 }
 log-interval seconds;
 maximum limit;
 threshold value;
}
local {
 family (inet | inet6) {
 disable;
 address address;
 anycast-pim {
 local-address address;
 address address <forward-msdp-sa>;
 rp-set {
 }
 }
 }
}
```

```

 }
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
}
}
register-limit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [policy-names];
static {
 address address {
 override;
 version version;
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 }
}
}
}

```

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | <p>Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.</p> <p>The remaining statements are explained separately.</p>                                                                                                  |
| <b>Default</b>                  | If you do not include the <b>rp</b> statement, the routing device can never become the RP.                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |

- Related Documentation**
- [Understanding PIM Sparse Mode on page 4213](#)

---

## rp-register-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | rp-register-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ],<br>[edit protocols pim <b>rp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to control incoming PIM register messages.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 4282</a></li><li>• <a href="#">dr-register-policy on page 4370</a></li></ul>                                                                                                                                |

## rp-set

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rp-set {   address address &lt;forward-msdp-sa&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4267</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                         |

## rpf-selection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rpf-selection {<br/>  group group-address {<br/>    source source-address {<br/>      next-hop next-hop-address;<br/>    }<br/>    wildcard-source {<br/>      next-hop next-hop-address;<br/>    }<br/>  }<br/>  prefix-list prefix-list-addresses {<br/>    source source-address {<br/>      next-hop next-hop-address;<br/>    }<br/>    wildcard-source {<br/>      next-hop next-hop-address;<br/>    }<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | If you omit the <b>rpf-selection</b> statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>source-address</b> —Specific source address for the PIM group.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                                                                                                                                                                                                                        |



## source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source source-address {     next-hop next-hop-address; }</pre>                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ] |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                              |
| <b>Description</b>              | Configure the source address for the PIM group.                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>source-address</b>—Specific source address for the PIM group.</p> <p>The remaining statements are explained separately.</p>                                                                                                             |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>                                                                                                                                             |

## spt-threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>spt-threshold {<br/>    infinity [ <i>policy-names</i> ];<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    pim],<br/>[edit protocols pim],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim]</pre>                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 4287</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## static (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   address address {     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     override;     version version;   } }</pre>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],<br/> [edit protocols pim <b>rp</b>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more <b>address</b> statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4266</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                     |

## threshold (PIM BFD Detection Time)

---

|                            |                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time]                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>         | Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.                                            |



**NOTE:** The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the [minimum-interval](#) or the [minimum-receive-interval](#) statement.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>milliseconds</i> —Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li><li>• <a href="#">detection-time on page 4368</a></li><li>• <a href="#">minimum-interval on page 4391</a></li><li>• <a href="#">minimum-receive-interval on page 4393</a></li></ul> |

## threshold (PIM BFD Transmit Interval)

|                            |                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                     |
| <b>Description</b>         | Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.                                    |
| <b>Options</b>             | <i>milliseconds</i> —Value for the transmit interval adaptation threshold.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )                                                                                                           |



**NOTE:** The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

|                                 |                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4290</a></li> <li>• <a href="#">bfd-liveness-detection on page 4363</a></li> </ul> |

## transmit-interval (PIM BFD Liveness Detection)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>transmit-interval {<br/>    minimum-interval milliseconds;<br/>    threshold milliseconds;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Hierarchy Level          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                 |
| Description              | <p>Specify the transmit interval for the <b>bfd-liveness-detection</b> statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li><li>• <a href="#">bfd-liveness-detection on page 4363</a></li><li>• <a href="#">threshold on page 4417</a></li><li>• <a href="#">minimum-interval on page 4392</a></li><li>• <a href="#">minimum-receive-interval on page 4393</a></li></ul>                                                                                                                                                                                                                                                                                                    |

## traceoptions (Protocols PIM)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br/>         [edit protocols pim],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols pim]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | <p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>             | The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>pim-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files<br/> <b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Tracing Flags</b></p> <ul style="list-style-type: none"> <li>• <b>assert</b>—Assert messages</li> <li>• <b>bidirectional-df-election</b>—Bidirectional PIM designated-forwarder (DF) election events</li> </ul> |

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.



**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 0 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                               |
|---------------------------------|-------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                |
|                                 | routing-control and trace-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Trace Options on page 4248</a></li> <li>• <a href="#">Tracing DVMRP Protocol Traffic</a></li> <li>• <a href="#">Tracing MSDP Protocol Traffic on page 4328</a></li> <li>• <a href="#">Configuring PIM Trace Options on page 4248</a></li> </ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## version (BFD)

---

|                                 |                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | version (0   1   automatic);                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols piminterface <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                  |
| <b>Description</b>              | Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.                                                                                                                                              |
| <b>Options</b>                  | Configure the BFD version to detect: <b>1</b> (BFD version 1) or <b>automatic</b> (autodetect the BFD version)<br><b>Default:</b> automatic                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4290</a></li></ul>                                                                                                                                      |

## version (PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>version</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the version of PIM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>version</b>—PIM version number.</p> <p><b>Range:</b> 1 or 2</p> <p><b>Default:</b> PIMv1 for rendezvous point (RP) mode (at the [edit protocols pim rp static address <i>address</i>] hierarchy level). PIMv2 for interface mode (at the [edit protocols pim interface <i>interface-name</i>] hierarchy level).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling PIM Sparse Mode on page 4255</a></li> <li>• <a href="#">Configuring PIM Dense Mode Properties</a></li> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## wildcard-source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | wildcard-source {<br>next-hop next-hop-address;<br>}                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit routing-instances routing-instance-name protocols pim rpf-selection group group-address],<br>[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                             |
| <b>Description</b>              | Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.<br><br>The remaining statements are explained separately.                                                   |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                 |

## IGMP Configuration Statements

---

- [accounting \(Protocols IGMP\) on page 4425](#)
- [accounting \(Protocols IGMP Interface\) on page 4426](#)
- [asm-override-ssm on page 4426](#)
- [disable \(Protocols IGMP\) on page 4427](#)
- [exclude \(Protocols IGMP\) on page 4427](#)
- [group \(Protocols IGMP\) on page 4428](#)
- [group-count \(Protocols IGMP\) on page 4429](#)
- [group-increment \(Protocols IGMP\) on page 4429](#)
- [group-limit on page 4430](#)
- [group-policy \(Protocols IGMP\) on page 4431](#)
- [igmp on page 4432](#)
- [immediate-leave \(Protocols IGMP\) on page 4434](#)
- [interface \(Protocols IGMP\) on page 4435](#)
- [maximum-transmit-rate \(Protocols IGMP\) on page 4436](#)
- [oif-map on page 4436](#)
- [passive \(IGMP\) on page 4437](#)
- [promiscuous-mode \(Protocols IGMP\) on page 4438](#)

- [query-interval \(Protocols IGMP\) on page 4438](#)
- [query-last-member-interval \(Protocols IGMP\) on page 4439](#)
- [query-response-interval \(Protocols IGMP\) on page 4440](#)
- [robust-count \(Protocols IGMP\) on page 4441](#)
- [source \(Protocols IGMP\) on page 4442](#)
- [source-count \(Protocols IGMP\) on page 4443](#)
- [source-increment \(Protocols IGMP\) on page 4443](#)
- [ssm-map \(Protocols IGMP\) on page 4444](#)
- [ssm-map-policy \(IGMP\) on page 4444](#)
- [static \(Protocols IGMP\) on page 4445](#)
- [traceoptions \(Protocols IGMP\) on page 4446](#)
- [version \(Protocols IGMP\) on page 4448](#)

## accounting (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | accounting;                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable the collection of IGMP join and leave event statistics on the system.                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Recording IGMP Join and Leave Events on page 4312</a></li> </ul>                                                                      |

## accounting (Protocols IGMP Interface)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (accounting   no-accounting);                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable or disable the collection of IGMP join and leave event statistics for an interface.                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Recording IGMP Join and Leave Events on page 4312</a></li></ul>                                                                        |

## asm-override-ssm

---

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | asm-override-ssm;                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                 |
| <b>Description</b>              | Enable the routing device to accept any-source multicast join messages (*;G) for group addresses that are within the default or configured range of source-specific multicast groups.                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347</a></li></ul>                                                                                                                                                                       |

## disable (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Disable IGMP on the system.                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling IGMP on page 4316</a></li> </ul>                                                                                                |


## exclude (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | exclude;                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.                                                      |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li> </ul>                                                                                                                                                   |

## group (Protocols IGMP)

---

|                                                                                                                                                                     |                                                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                                                                                                                                              | <pre>group multicast-group-address {<br/>  exclude;<br/>  group-count number;<br/>  group-increment increment;<br/>  source ip-address {<br/>    source-count number;<br/>    source-increment increment;<br/>  }<br/>}</pre> |
| Hierarchy Level                                                                                                                                                     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name static</a> ],<br>[edit protocols <a href="#">igmp interface interface-name static</a> ]                                  |
| Release Information                                                                                                                                                 | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                |
| Description                                                                                                                                                         | Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.                                                                             |
| <hr/>                                                                                                                                                               |                                                                                                                                                                                                                               |
| <div> <b>NOTE:</b> You must specify a unique address for each group.</div> <hr/> |                                                                                                                                                                                                                               |
| The remaining statements are explained separately.                                                                                                                  |                                                                                                                                                                                                                               |
| Required Privilege Level                                                                                                                                            | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                           |
| Related Documentation                                                                                                                                               | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li></ul>                                                                                                          |



## group-count (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-count <i>number</i>;</code>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                       |
| <b>Description</b>              | Specify the number of static groups to be created.                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>number</i> —Number of static groups.<br><b>Default:</b><br><b>Range:</b> 1 through 512                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li> </ul>                                                                                                                                                   |

## group-increment (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-increment <i>increment</i>;</code>                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                       |
| <b>Description</b>              | Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                         |
| <b>Options</b>                  | <i>increment</i> —Number of times the address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li> </ul>                                                                                                                                                   |

## group-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.</p> <p>To confirm the configured group limit on the interface, use the <b>show igmp interface</b> command.</p> |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups that can join the interface.                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>limit</i> —group limit value for the interface.<br><b>Range:</b> 1 through 32767                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4313</a></li><li>• <i>group-threshold</i></li><li>• <i>log-interval</i></li></ul>                                                                                                                      |

---

## group-policy (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> ]                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                |
| <b>Description</b>              | When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4301</a></li></ul>                                                                                                                                                        |

## igmp

---

**Syntax**    `igmp {  
    accounting;  
    interface interface-name {  
        disable;  
        (accounting | no-accounting);  
        group-limit limit;  
        group-policy [ policy-names ];  
        group-threshold  
        immediate-leave;  
        log-interval  
        oif-map map-name;  
        passive;  
        promiscuous-mode;  
        ssm-map ssm-map-name;  
        ssm-map-policy ssm-map-policy-name;  
        static {  
            group multicast-group-address {  
                exclude;  
                group-count number;  
                group-increment increment;  
                source ip-address {  
                    source-count number;  
                    source-increment increment;  
                }  
            }  
        }  
        version version;  
    }  
    query-interval seconds;  
    query-last-member-interval seconds;  
    query-response-interval seconds;  
    robust-count number;  
    traceoptions {  
        file filename <files number> <size size> <world-readable | no-world-readable>;  
        flag flag <flag-modifier> <disable>;  
    }  
}`

**Hierarchy Level**    `[edit logical-systems logical-system-name protocols],`  
                          `[edit protocols]`


**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP on page 4297</a></li></ul>                                                                                                             |

## immediate-leave (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the <b>immediate-leave</b> statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |
|                                 | <p> <b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- Related Documentation**
- [Specifying Immediate-Leave Host Removal for IGMP on page 4300](#)

## interface (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {     disable;     (accounting   no-accounting);     group-limit <i>limit</i>;     group-policy [ <i>policy-names</i> ];     immediate-leave;     oif-map <i>map-name</i>;     passive;     promiscuous-mode;     ssm-map <i>ssm-map-name</i>;     ssm-map-policy <i>ssm-map-policy-name</i>;     static {         group <i>mcast-group-address</i> {             exclude;             group-count <i>number</i>;             group-increment <i>increment</i>;             source <i>ip-address</i> {                 source-count <i>number</i>;                 source-increment <i>increment</i>;             }         }     }     version <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b>],</p> <p>[edit protocols <b>igmp</b>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Enable IGMP on an interface and configure interface-specific properties.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP on page 4297</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## maximum-transmit-rate (Protocols IGMP)

---

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | maximum-transmit-rate <i>packets-per-second</i> ;                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                 |
| <b>Description</b>              | Limit the transmission rate of IGMP packets                                                                                                                        |
| <b>Options</b>                  | <b>packets-per-second</b> —Maximum number of IGMP packets transmitted in one second by the router.<br><b>Range:</b> 1 through 10000<br><b>Default:</b> 500 packets |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Limiting the Maximum IGMP Message Rate on page 4305</a></li></ul>                                              |


## oif-map

---

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | oif-map <i>map-name</i> ;                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                         |
| <b>Description</b>              | Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast with Subscriber VLANs</a></li></ul>                                                     |



## passive (IGMP)

|                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                   | <code>passive &lt;allow-receive&gt; &lt;send-general-query&gt; &lt;send-group-query&gt;;</code>                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                      | Statement introduced in Junos OS Release 9.6.<br><b>allow-receive</b> , <b>send-general-query</b> , and <b>send-group-query</b> options were added in Junos OS Release 10.0.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                    |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                              | Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.                                                                                                                    |
| <div>  <p><b>NOTE:</b> You can selectively activate up to two out of the three available options for the <b>passive</b> statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the <b>passive</b> statement.</p> </div> |                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                  | <p><b>allow-receive</b>—Enables IGMP to receive control traffic on the interface.</p> <p><b>send-general-query</b>—Enables IGMP to send general queries on the interface.</p> <p><b>send-group-query</b>—Enables IGMP to send group-specific and group-source-specific queries on the interface.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                 | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li> <li>• <a href="#">Enabling IGMP on page 4297</a></li> </ul>                                                                                                                                |

## promiscuous-mode (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>promiscuous-mode;</code>                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit dynamic-profiles <i>profile-name</i> protocols igmp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 9.2 for dynamic profiles.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                              |
| <b>Description</b>              | Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a Dynamic Profile for Client Access</a></li><li>• <a href="#">Accepting IGMP Messages from Remote Subnetworks on page 4302</a></li></ul>                                                                                                                                      |

## query-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                   |
| <b>Description</b>              | Specify how often the querier router sends general host-query messages.                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>seconds</b> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Host-Query Message Interval on page 4299</a></li><li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 4439</a></li><li>• <a href="#">query-response-interval (Protocols IGMP) on page 4440</a></li></ul> |

## query-last-member-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-last-member-interval <i>seconds</i> ;                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                          |
| <b>Description</b>              | Specify how often the querier router sends group-specific query messages.                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 999999<br><b>Default:</b> 1 second                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Last-Member Query Interval on page 4299</a></li> <li>• <a href="#">query-interval (Protocols IGMP) on page 4438</a></li> <li>• <a href="#">query-response-interval (Protocols IGMP) on page 4440</a></li> </ul> |

## query-response-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-response-interval <i>seconds</i> ;                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify how long the querier router waits to receive a response to a host-query message from a host.                                                                                                                                                                                |
| <b>Options</b>                  | <b>seconds</b> —The query response interval must be less than the query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Query Response Interval on page 4303</a></li><li>• <a href="#">query-interval (Protocols IGMP) on page 4438</a></li><li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 4439</a></li></ul> |

---

## robust-count (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.                            |
| <b>Options</b>                  | <i>number</i> —Robustness variable.<br><b>Range:</b> 2 through 10<br><b>Default:</b> 2                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Robustness Variable on page 4304</a></li></ul>                                                                          |

## source (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source <i>ip-address</i> {<br/>    <i>source-count</i> <i>number</i>;<br/>    <i>source-increment</i> <i>increment</i>;<br/>}</code>                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ],<br>[edit protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                 |
| <b>Description</b>              | Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.                                                                                                                                                                    |
| <b>Options</b>                  | <i>ip-address</i> —IPv4 unicast address.<br><br>The remaining statements are explained separately.                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li></ul>                                                                                                                                                                           |

## source-count (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-count <i>number</i>;</code>                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>              | Configure the number of multicast source addresses that should be accepted for each static group created.                                                                                                                                                                                      |
| <b>Options</b>                  | <i>number</i> —Number of source addresses.<br><b>Default:</b> 1<br><b>Range:</b> 1 through 1024                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li> </ul>                                                                                                                                                                         |

## source-increment (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-increment <i>number</i>;</code>                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>              | Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                              |
| <b>Options</b>                  | <i>increment</i> —Number of times the source address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4305</a></li> </ul>                                                                                                                                                                         |

## ssm-map (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map <i>ssm-map-name</i>;</code>                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Apply an SSM map to an IGMP interface.                                                                                                                                                     |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of SSM map.                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 4344</a></li></ul>                                                                            |

## ssm-map-policy (IGMP)

---

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                            |
| <b>Description</b>              | Apply an SSM map policy to an IGMP interface.                                                                                                                                  |
| <b>Options</b>                  | <i>ssm-map-policy-name</i> —Name of SSM map policy.                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Maps for Different Groups to Different Sources on page 4350</a></li></ul>                         |



## static (Protocols IGMP)

```
Syntax static {
 group multicast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
 }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **igmp interface** *interface-name*],  
[edit protocols **igmp interface** *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



**NOTE:** To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Enabling IGMP Static Group Membership on page 4305](#)

## traceoptions (Protocols IGMP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | <p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>             | The default IGMP trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>igmp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>IGMP Tracing Flags</b></p> <ul style="list-style-type: none"><li>• <b>leave</b>—Leave group messages (for IGMP version 2 only).</li><li>• <b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li><li>• <b>packets</b>—All IGMP packets.</li></ul> |

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Tracing IGMP Protocol Traffic on page 4314](#)

---

## version (Protocols IGMP)

---

**Syntax** `version version;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **igmp interface** *interface-name*],  
[edit protocols **igmp interface** *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Specify the version of IGMP.

**Options** **version**—IGMP version number.

**Range:** 1, 2, or 3

**Default:** IGMP version 2

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Changing the IGMP Version on page 4298](#)

---

## IGMP Snooping Configuration Statements

---

- [data-forwarding on page 4449](#)
- [disable \(IGMP Snooping\) on page 4450](#)
- [group \(IGMP Snooping\) on page 4450](#)
- [groups \(Multicast VLAN Registration\) on page 4451](#)
- [igmp-snooping on page 4452](#)
- [install \(Multicast VLAN Registration\) on page 4453](#)
- [interface \(IGMP Snooping\) on page 4453](#)
- [multicast-router-interface \(IGMP Snooping\) on page 4454](#)

- [proxy \(Multicast VLAN Registration\) on page 4454](#)
- [receiver on page 4455](#)
- [robust-count \(IGMP Snooping\) on page 4455](#)
- [source \(Multicast VLAN Registration\) on page 4456](#)
- [source-vlans on page 4456](#)
- [static \(IGMP Snooping\) on page 4457](#)
- [traceoptions \(IGMP Snooping\) on page 4458](#)
- [vlan \(IGMP Snooping\) on page 4460](#)
- [version \(IGMP Snooping\) on page 4461](#)

## data-forwarding

**Syntax**

```
data-forwarding {
 receiver {
 source-vlans vlan-list;
 install;
 }
 source {
 groups group-prefix;
 }
}
```

**Hierarchy Level** [edit protocols igmp-snooping vlan (all | *vlan-name*)]

**Release Information** Statement introduced in Junos OS Release 9.6 for EX Series switches.  
Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.

The remaining statements are explained separately.

**Default** Disabled

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Multicast VLAN Registration on page 4322](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 4321](#)

## disable (IGMP Snooping)

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable {<br/>    interface <i>interface-name</i>;<br/>}</code>                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> <i>vlan-name</i> ]                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                  |
| <b>Description</b>              | Disable IGMP snooping on all interfaces in a VLAN or on a specific VLAN interface.                                                                                                 |
| <b>Default</b>                  | If you do not specify an interface, all interfaces in the given VLAN are disabled (because IGMP snooping is disabled by default).                                                  |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li></ul> |

## group (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group <i>ip-address</i>;</code>                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> <i>vlan-name</i> <a href="#">interface</a> <i>interface-name</i> <a href="#">static</a> ]                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure a static multicast group using a valid IP multicast address.                                                                                                                                                                             |
| <b>Default</b>                  | None.                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>ip-address</i> —IP address of the multicast group receiving data on an interface.                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 4543</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li></ul> |

---

## groups (Multicast VLAN Registration)

---

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>groups <i>group-prefix</i>;</code>                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding source]                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify the IP address range of the multicast VLAN (MVLAN) source interfaces.                                                                                                                                                  |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>group-prefix</i> —IP address range of the source group. Each MVLAN must have exactly one <b>groups</b> statement. If there are multiple MVLANs on the switch, their group ranges must be unique.                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li></ul> |

## igmp-snooping

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>igmp-snooping {<br/>  traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt; &lt;match<br/>      <i>regex</i>&gt;;<br/>    flag <i>flag</i> (detail   disable   receive   send);<br/>  }<br/>  vlan <i>vlan-name</i> {<br/>    data-forwarding {<br/>      source {<br/>        groups <i>group-prefix</i>;<br/>      }<br/>      receiver {<br/>        source-vlans <i>vlan-list</i>;<br/>        install;<br/>      }<br/>    }<br/>    disable {<br/>      interface <i>interface-name</i>;<br/>    }<br/>    immediate-leave;<br/>    interface <i>interface-name</i> {<br/>      multicast-router-interface;<br/>      static {<br/>        group <i>ip-address</i>;<br/>      }<br/>    }<br/>    robust-count <i>number</i>;<br/>    version <i>number</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br><b>version</b> statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Enable and configure IGMP snooping.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | IGMP snooping is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



## install (Multicast VLAN Registration)

|                                 |                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | install;                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                         |
| <b>Description</b>              | Install forwarding entries in the multicast receiver VLAN. By default, the multicast VLAN (MVLAN) installs forwarding entries for MVLAN groups only.                                                                              |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li> <li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li> </ul> |

## interface (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     multicast-router-interface;     static {         group <i>ip-address</i>;     } }</pre>                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols <b>igmp-snooping</b> vlan <i>vlan-name</i> ]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | Enable IGMP snooping on an interface and configure interface-specific properties.<br><br>The remaining statements are explained separately.                                                                                                            |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping vlans on page 4543</a></li> </ul> |

## multicast-router-interface (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>multicast-router-interface;</code>                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> <i>vlan-name</i> <a href="#">interface</a> <i>interface-name</i> ]                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Configure an interface to forward IGMP messages to multicast routers.                                                                                                                                                                              |
| <b>Default</b>                  | Disabled. If this statement is disabled, the interface drops IGMP messages it receives.                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 4543</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li></ul> |

## proxy (Multicast VLAN Registration)

---

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>proxy source-address <i>ip-address</i>;</code>                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> (all   <i>vlan-name</i> )]                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify that the VLAN operate in proxy mode. The proxy option is supported only for a VLAN acting as a data-forwarding source.                                                                                                 |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                       |
| <b>Options</b>                  | <code>source-address <i>ip-address</i></code> —IP address of the source VLAN to act as proxy.                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li></ul> |

## receiver

---

|                                 |                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> receiver {   source-vlans <i>vlan-list</i>;   install; }</pre>                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                         |
| <b>Description</b>              | Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).<br><br>The remaining statements are explained separately.                                                                                            |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li> <li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li> </ul> |

## robust-count (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | robust-count <i>number</i> ;                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <b>igmp-snooping</b> vlan <i>vlan-name</i> ]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the <b>query-interval</b> statement.                                             |
| <b>Default</b>                  | 2 intervals                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>number</i> —Number of intervals the switch waits before timing out a multicast group.<br><b>Range:</b> 2 through 10                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping vlans on page 4543</a></li> </ul> |

## source (Multicast VLAN Registration)

---

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source {<br/>    <b>groups</b> <i>group-prefix</i>;<br/>}</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Configure a VLAN to be a multicast source VLAN (MVLAN).<br><br>The remaining statement is explained separately.                                                                                                                |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li></ul> |

## source-vlans

---

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-vlans <i>vlan-list</i>;</code>                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify a list of multicast VLANs (MVLANs) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANs must be in proxy mode or none of them can be in proxy mode.                          |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>vlan-list</i> —Names of the MVLANs.                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration on page 4322</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure) on page 4321</a></li></ul> |

---

## static (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {<br/>    group ip-address;<br/>}</pre>                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> <i>vlan-name</i> <a href="#">interface</a> <i>interface-name</i> ]                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | <p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                 |
| <b>Default</b>                  | No multicast groups are statically defined.                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li><li>• <a href="#">show igmp-snooping vlans on page 4543</a></li></ul> |

## traceoptions (IGMP Snooping)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;size <i>size</i>&gt; &lt;replace&gt; &lt;world-readable  <br/>    no-world-readable&gt;;<br/>    flag <i>flag</i> (detail   disable   receive   send);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Hierarchy Level     | [edit protocols <a href="#">igmp-snooping</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Description         | Define tracing operations for IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Default             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>general</b>—Trace general IGMP snooping protocol events.</li><li>• <b>krt</b>—Trace communication over routing sockets.</li><li>• <b>nexthop</b>— Trace next-hop related events.</li><li>• <b>normal</b>—Trace normal IGMP snooping protocol events.</li><li>• <b>packets</b>—Trace all IGMP packets.</li><li>• <b>policy</b>—Trace policy processing.</li><li>• <b>query</b>—Trace IGMP membership query messages.</li><li>• <b>report</b>—Trace membership report messages.</li><li>• <b>route</b>—Trace routing information.</li><li>• <b>state</b>—Trace IGMP state transitions.</li><li>• <b>task</b>—Trace routing protocol task processing.</li></ul> |

- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN related events.

**no-stamp**—(Optional) Do not time stamp trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**size size** —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option. Use **xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes.

**Range:** 10 KB through 1 gigabytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |                                                             |
|---------------------------------|-------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.        |
|                                 | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                       |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 4318</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> </ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## vlan (IGMP Snooping)

---

**Syntax**    `vlan vlan-name {  
              disable {  
                  interface interface-name;  
              }  
              immediate-leave;  
              interface interface-name {  
                  multicast-router-interface;  
                  static {  
                      group ip-address;  
                  }  
              }  
              version number;  
          }`

**Hierarchy Level**    [edit protocols [igmp-snooping](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              **version** statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Configure IGMP snooping parameters for a VLAN.

The remaining statements are explained separately.



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlangs` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range. For IGMP snooping, secondary private VLANs are not listed.

---

**Default**    IGMP snooping options apply to the specified VLAN.

**Options**    *vlan-name*—Name of a VLAN.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring IGMP Snooping on page 4318](#)
- [Configuring IGMP Snooping on page 4317](#)
- [show igmp-snooping vlans on page 4543](#)



## version (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages. |
| <b>Default</b>                  | If you do not configure the <b>version</b> statement, the default is IGMPv2.                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1 through 3                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> </ul>                                                                                                                                                                                                                        |

## MSDP Configuration Statements

- [active-source-limit on page 4462](#)
- [authentication-key on page 4463](#)
- [data-encapsulation on page 4464](#)
- [default-peer on page 4465](#)
- [disable \(Protocols MSDP\) on page 4466](#)
- [export \(Protocols MSDP\) on page 4467](#)
- [group on page 4468](#)
- [import \(Protocols MSDP\) on page 4469](#)
- [local-address on page 4470](#)
- [maximum on page 4471](#)
- [mode \(Protocols MSDP\) on page 4472](#)
- [msdp on page 4473](#)
- [peer \(Protocols MSDP\) on page 4475](#)
- [rib-group \(Protocols MSDP\) on page 4476](#)

- [source on page 4477](#)
- [threshold on page 4478](#)
- [traceoptions \(Protocols MSDP\) on page 4479](#)

## active-source-limit

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>active-source-limit {<br/>  log-interval <i>seconds</i>;<br/>  log-warning <i>value</i>;<br/>  maximum <i>number</i>;<br/>  threshold <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Hierarchy Level          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> <b>source</b> <i>ip-address/prefix-length</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  <b>msdp</b> <b>source</b> <i>ip-address/prefix-length</i>],<br/>[edit protocols <b>msdp</b>],<br/>[edit protocols <b>msdp</b> group <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit protocols <b>msdp</b> <b>source</b> <i>ip-address/prefix-length</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> group <i>group-name</i><br/>  <b>peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> <b>peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> <b>source</b><br/>  <i>ip-address/prefix-length</i>]</pre> |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Description              | Limit the number of active source messages the routing device accepts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Default                  | If you do not include this statement, the router accepts any number of MSDP active source messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Options                  | The options are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## authentication-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key peer-key;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit protocols <code>msdp group group-name peer address</code>],</p> <p>[edit protocols <code>msdp peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | If you do not include this statement, the router accepts any valid MSDP messages from the peer address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>peer-key</b> —MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (, ), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## data-encapsulation

---

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | data-encapsulation (disable   enable);                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ],<br>[edit protocols <a href="#">msdp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.                                                                                                                                                                                                 |
| <b>Default</b>                  | If you do not include this statement, the RP encapsulates multicast data.                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>disable</b> —(Optional) Do not use MSDP data encapsulation.<br><b>enable</b> —Use MSDP data encapsulation.<br><b>Default:</b> enable                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li></ul>                                                                                                                                                                                                          |

## default-peer

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | default-peer;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## disable (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],<br/>[edit protocols <b>msdp</b>],<br/>[edit protocols <b>msdp group</b> <i>group-name</i>],<br/>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit protocols <b>msdp peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Explicitly disable MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling MSDP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## export (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more policies to routes being exported from the routing table into MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> <li>• <a href="#">import on page 4469</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## group

---

**Syntax**    `group group-name {  
          disable;  
          export [ policy-names ];  
          import [ policy-names ];  
          local-address address;  
          mode (mesh-group | standard);  
          traceoptions {  
            file filename <files number> <size size> <world-readable | no-world-readable>;  
            flag flag <flag-modifier> <disable>;  
          }  
          peer address; {  
            disable;  
            active-source-limit {  
              maximum number;  
              threshold number;  
            }  
            authentication-key peer-key;  
            default-peer;  
            export [ policy-names ];  
            import [ policy-names ];  
            local-address address;  
            traceoptions {  
              file filename <files number> <size size> <world-readable | no-world-readable>;  
              flag flag <flag-modifier> <disable>;  
            }  
          }  
          }  
          }`

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols [msdp](#)],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                          [msdp](#)],  
                          [edit protocols [msdp](#)],  
                          [edit routing-instances *routing-instance-name* protocols [msdp](#)]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the [peer](#) statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

**Options**    *group-name*—Name of the MSDP group.

The remaining statements are explained separately.



|                              |                                                                                                            |
|------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                       |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul> |

## import (Protocols MSDP)

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>       | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b>   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>           | Apply one or more policies to routes being imported into the routing table from MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>               | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> <li>• <a href="#">export on page 4467</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## local-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp</b>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>address</b> —IP address of the local end of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Example: Configuring MSDP in a Routing Instance</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## maximum

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the maximum number of MSDP active source messages the router accepts.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>number</i> —Maximum number of active source messages.<br><b>Range:</b> 1 through 1,000,000<br><b>Default:</b> 25,000                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li> <li>• <a href="#">threshold on page 4478</a></li> </ul>                                                                                                                                                                                                                                      |

## mode (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode (mesh-group   standard);                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is <b>standard</b> .                                                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not include this statement, default flooding is applied.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>mesh-group</b> —Group of peers that are mesh group members.<br><br><b>standard</b> —Use standard MSDP source-active flooding rules.<br><b>Default:</b> standard                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li></ul>                                                                                                                                                                                                                                                                      |

## msdp

```

Syntax msdp {
 disable;
 active-source-limit {
 log-interval seconds;
 log-warning value;
 maximum number;
 threshold number;
 }
 data-encapsulation (disable | enable);
 export [policy-names];
 group group-name {
 ...group-configuration ...
 }
 hold-time seconds;
 import [policy-names];
 local-address address;
 keep-alive seconds;
 peer address {
 ...peer-configuration ...
 }
 rib-group group-name;
 source ip-prefix</prefix-length> {
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 sa-hold-time seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 peer address {
 ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
 just following ...
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
 peer address {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 }

```

```
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],  
[edit protocols],  
[edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.4 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.

**Default** MSDP is disabled on the router or switch.

**Options** The statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

## peer (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> peer address {     disable;     active-source-limit {         maximum number;         threshold number;     }     authentication-key peer-key;     default-peer;     export [ policy-names ];     import [ policy-names ];     local-address address;     traceoptions {         file filename &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;         flag flag &lt;flag-modifier&gt; &lt;disable&gt;;     } } </pre>                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>]</p>               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Define an MSDP peering relationship. An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple <b>peer</b> statements.</p> <p>By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the <b>peer (Protocols MSDP)</b> statement.</p> <p>At least one peer must be configured for MSDP to function. You must configure <b>address</b> and <b>local-address</b>.</p> |
| <b>Options</b>             | <p><b>address</b>—Name of the MSDP peer.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege</b>  | routing—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Level</b>               | routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Related Documentation** • *Example: Configuring MSDP in a Routing Instance*

---

## rib-group (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ],<br>[edit protocols <a href="#">msdp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                      |
| <b>Description</b>              | Associate a routing table group with MSDP.                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group. The name must be one that you defined with the <b>rib-groups</b> statement at the [edit routing-options] hierarchy level.                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | • <i>Example: Configuring MSDP in a Routing Instance</i>                                                                                                                                                                                                                                                                                                    |



## source

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source ip-address &lt;/prefix-length&gt; {     active-source-limit {         maximum number;         threshold number;     } }</pre>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                          |
| <b>Description</b>              | Limit the number of active source messages the routing device accepts from sources in this address range.                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | If you do not include this statement, the routing device accepts any number of MSDP active source messages.                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | The other statements are explained separately.                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li> </ul>                                                                                                                                                                                                                       |

## threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>threshold <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>number</i></b> —RED threshold for active source messages.<br><b>Range:</b> 1 through 1,000,000<br><b>Default:</b> 24,000                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4332</a></li><li>• <a href="#">maximum on page 4471</a></li></ul>                                                                                                                                                                                                                                           |

## traceoptions (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure MSDP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>             | <p>The default MSDP trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the [edit <b>routing-options</b>] hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>msdp-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

**Range:** 2 through 1000 files

**Default:** 2 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information

- **receive**—Packets being received

- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow any user to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                           |                                                                               |
|---------------------------|-------------------------------------------------------------------------------|
| <b>Required Privilege</b> | routing and trace—To view this statement in the configuration.                |
| <b>Level</b>              | routing-control and trace-control—To add this statement to the configuration. |

|                              |                                                                                                                |
|------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Tracing MSDP Protocol Traffic on page 4328</a></li> </ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------|

## Source-Specific Multicast Configuration Statements

- [asm-override-ssm on page 4482](#)
- [policy \(SSM Maps\) on page 4483](#)
- [ssm-groups on page 4484](#)
- [ssm-map \(Protocols IGMP\) on page 4485](#)
- [ssm-map \(Routing Options Multicast\) on page 4485](#)
- [ssm-map-policy \(IGMP\) on page 4486](#)

## asm-override-ssm

---

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | asm-override-ssm;                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                 |
| <b>Description</b>              | Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347</a></li></ul>                                                                                                                                                                       |

## policy (SSM Maps)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>],</p> <p>[edit routing-options multicast <a href="#">ssm-map</a> <i>ssm-map-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Apply one or more policies to an SSM map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies for SSM mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To view this statement in the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SSM Mapping on page 4344</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                      |

## ssm-groups

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-groups [ <i>ip-addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit routing-options multicast]</code>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the <b>ssm-groups</b> statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the <b>ssm-groups</b> statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p> |
| <b>Options</b>                  | <i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4347</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



## ssm-map (Protocols IGMP)

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map ssm-map-name;</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Apply an SSM map to an IGMP interface.                                                                                                                                                     |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of SSM map.                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SSM Mapping on page 4344</a></li> </ul>                                                                          |

## ssm-map (Routing Options Multicast)

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map ssm-map-name {<br/>    <b>policy</b> [ <i>policy-names</i> ];<br/>    source [ <i>addresses</i> ];<br/>}</code>                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                 |
| <b>Description</b>              | Configure SSM mapping.                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of the SSM map.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SSM Mapping on page 4344</a></li> </ul>                                                                                                                                                                                                                   |

## ssm-map-policy (IGMP)

---

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                            |
| <b>Description</b>              | Apply an SSM map policy to an IGMP interface.                                                                                                                                  |
| <b>Options</b>                  | <i>ssm-map-policy-name</i> —Name of SSM map policy.                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Maps for Different Groups to Different Sources on page 4350</a></li></ul>                         |

## CHAPTER 50

# Administration

- [Routine Monitoring on page 4487](#)
- [Monitoring Commands for Multicast Protocols on page 4488](#)

### Routine Monitoring

---

- [Monitoring IGMP Snooping on page 4487](#)
- [Verifying the IGMP Snooping Group Timeout Value on page 4488](#)

### Monitoring IGMP Snooping

**Purpose** Use the monitoring feature to view status and information about the IGMP snooping configuration.

**Action** To display IGMP snooping details in the CLI, enter the following commands:

- **show igmp-snooping vlans**
- **show igmp-snooping statistics**
- **show igmp-snooping route**
- **show igmp-snooping membership**

**Meaning** [Table 326 on page 4487](#) summarizes the IGMP snooping details displayed.

**Table 326: Summary of IGMP Snooping Output Fields**

| Field                 | Values                                              |
|-----------------------|-----------------------------------------------------|
| IGMP Snooping Monitor |                                                     |
| VLAN                  | VLAN for which IGMP snooping is enabled.            |
| Interfaces            | Interface connected to a multicast router.          |
| Groups                | Number of the multicast groups learned by the VLAN. |
| MRouters              | Multicast router.                                   |
| Receivers             | Multicast receiver.                                 |

Table 326: Summary of IGMP Snooping Output Fields (*continued*)

| Field                  | Values                                                             |
|------------------------|--------------------------------------------------------------------|
| IGMP Route Information |                                                                    |
| VLAN                   | VLAN for which IGMP snooping is enabled.                           |
| Next-Hop               | Next hop assigned by the switch after performing the route lookup. |
| Group                  | Multicast groups learned by the VLAN.                              |

- Related Documentation**
- [IGMP Snooping Overview on page 4232](#)
  - [Example: Configuring IGMP Snooping on page 4318](#)
  - [Configuring IGMP Snooping on page 4317](#)
  - [Changing the IGMP Snooping Group Timeout Value on page 4318](#)

## Verifying the IGMP Snooping Group Timeout Value

**Purpose** Verify that the IGMP snooping group timeout value has been changed correctly from its default value.

**Action** Display the IGMP snooping membership information, which contains the group timeout value that was derived from the IGMP configuration:

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.1
Receiver count: 1, Flags: <v2-hosts>
ge-0/0/15.0 Uptime: 00:00:05 timeout: 510
```

**Meaning** The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. When you enable IGMP snooping, the default IGMP snooping group timeout value of 260 seconds is applied to all VLANs, which means that the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table. You can change the timeout value by using the **robust-count** option.

- Related Documentation**
- [Changing the IGMP Snooping Group Timeout Value on page 4318](#)

## Monitoring Commands for Multicast Protocols

- [clear igmp membership](#)
- [clear igmp-snooping membership](#)
- [clear igmp statistics](#)
- [clear igmp-snooping statistics](#)

- `clear msdp cache`
- `clear msdp statistics`
- `clear multicast bandwidth-admission`
- `clear multicast scope`
- `clear multicast sessions`
- `clear multicast statistics`
- `clear pim join`
- `clear pim register`
- `clear pim statistics`
- `mtrace`
- `mtrace from-source`
- `mtrace monitor`
- `mtrace to-gateway`
- `show configuration protocols igmp`
- `show igmp group`
- `show igmp interface`
- `show igmp statistics`
- `show igmp-snooping membership`
- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show igmp-snooping vlans`
- `show msdp`
- `show msdp source`
- `show msdp source-active`
- `show msdp statistics`
- `show multicast flow-map`
- `show multicast interface`
- `show multicast mrinfo`
- `show multicast next-hops`
- `show multicast pim-to-igmp-proxy`
- `show multicast pim-to-mld-proxy`
- `show multicast route`
- `show multicast rpf`
- `show multicast scope`
- `show multicast sessions`
- `show multicast usage`
- `show pim bootstrap`

- [show pim interfaces](#)
- [show pim join](#)
- [show pim neighbors](#)
- [show pim rps](#)
- [show pim source](#)
- [show pim statistics](#)
- [show system statistics igmp](#)
- [test msdp](#)

## clear igmp membership

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4491</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4491</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                       | <pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear Internet Group Management Protocol (IGMP) group members.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Clear all IGMP members on all interfaces and for all address ranges.</p> <p><b>group address-range</b>—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is <b>224.2/16</b>. If you omit the destination prefix length, the default is <b>/32</b>.</p> <p><b>interface interface-name</b>—(Optional) Clear all IGMP group members on an interface.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show igmp group on page 4525</a></li> <li>• <a href="#">show igmp interface on page 4529</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                        | <a href="#">clear igmp membership on page 4491</a><br><a href="#">clear igmp membership interface on page 4492</a><br><a href="#">clear igmp membership group on page 4493</a>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>                                | See <a href="#">show igmp group</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Sample Output

### clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 186     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 186     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 187     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 188     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

```

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

### clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 210     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 210     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 215     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 216     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

```

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |



## clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 210     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 210     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 215     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 216     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```
user@host> clear igmp membership group 239.225/16
```

```
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 231     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 233     |
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 236     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

## clear igmp-snooping membership

---

|                                 |                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear igmp-snooping membership</code><br><code>&lt;vlan <i>vlan-name</i>&gt;</code>                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Clear IGMP snooping membership information.                                                                  |
| <b>Options</b>                  | <code>vlan <i>vlan-name</i></code> —(Optional) Name of the VLAN.                                             |
| <b>Required Privilege Level</b> | view                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping membership on page 4536</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping membership on page 4494</a>                                                  |

### Sample Output

clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

## clear igmp statistics

|                                    |                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 4495</a><br><a href="#">Syntax (EX Series Switches) on page 4495</a>                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | clear igmp statistics<br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches)</b> | clear igmp statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                             |
| <b>Description</b>                 | Clear Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <b>none</b> —Clear IGMP statistics on all interfaces.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear IGMP statistics for the specified interface only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">show igmp statistics on page 4533</a></li> </ul>                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>       | <a href="#">clear igmp statistics on page 4495</a>                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | See <a href="#">show igmp statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                        |

## Sample Output

### clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type Received Sent Rx errors
Membership Query 8883 459 0
V1 Membership Report 0 0 0
DVMRP 19784 35476 0
PIM V1 18310 0 0
Cisco Trace 0 0 0
V2 Membership Report 0 0 0
Group Leave 0 0 0
Mtrace Response 0 0 0

```

|                                     |   |   |   |
|-------------------------------------|---|---|---|
| Mtrace Request                      | 0 | 0 | 0 |
| Domain Wide Report                  | 0 | 0 | 0 |
| V3 Membership Report                | 0 | 0 | 0 |
| Other Unknown types                 |   |   | 0 |
| IGMP v3 unsupported type            |   |   | 0 |
| IGMP v3 source required for SSM     |   |   | 0 |
| IGMP v3 mode not applicable for SSM |   |   | 0 |

IGMP Global Statistics

|                |      |
|----------------|------|
| Bad Length     | 0    |
| Bad Checksum   | 0    |
| Bad Receive If | 0    |
| Rx non-local   | 1227 |

user@host> clear igmp statistics

user@host> show igmp statistics

IGMP packet statistics for all interfaces

| IGMP Message type                   | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query                    | 0        | 0    | 0         |
| V1 Membership Report                | 0        | 0    | 0         |
| DVMRP                               | 0        | 0    | 0         |
| PIM V1                              | 0        | 0    | 0         |
| Cisco Trace                         | 0        | 0    | 0         |
| V2 Membership Report                | 0        | 0    | 0         |
| Group Leave                         | 0        | 0    | 0         |
| Mtrace Response                     | 0        | 0    | 0         |
| Mtrace Request                      | 0        | 0    | 0         |
| Domain Wide Report                  | 0        | 0    | 0         |
| V3 Membership Report                | 0        | 0    | 0         |
| Other Unknown types                 |          |      | 0         |
| IGMP v3 unsupported type            |          |      | 0         |
| IGMP v3 source required for SSM     |          |      | 0         |
| IGMP v3 mode not applicable for SSM |          |      | 0         |
| IGMP Global Statistics              |          |      |           |
| Bad Length                          | 0        |      |           |
| Bad Checksum                        | 0        |      |           |
| Bad Receive If                      | 0        |      |           |
| Rx non-local                        | 0        |      |           |

## clear igmp-snooping statistics

---

|                                 |                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear igmp-snooping statistics                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Clear IGMP snooping statistics.                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping statistics on page 4541</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping statistics on page 4497</a>                                                  |

### Sample Output

#### clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

## clear msdp cache

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear msdp cache</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;peer <i>peer-address</i>&gt;</code>                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>none</b> —Clear entries in the MSDP source-active cache for all instances, logical systems, and peers.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear entries for a specific MSDP instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>peer <i>peer-address</i></b> —(Optional) Clear the MSDP source-active cache entries learned from a specific peer. |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp source-active on page 4549</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">clear msdp cache on page 4498</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear msdp cache

```
user@host> clear msdp cache
```

## clear msdp statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear msdp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Clear Multicast Source Discovery Protocol (MSDP) peer statistics.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Clear MSDP statistics for all peers.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Clear the statistics for the specified peer.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp statistics on page 4552</a></li> </ul>                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear msdp statistics on page 4499</a>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear msdp statistics

```
user@host> clear msdp statistics
```

## clear multicast bandwidth-admission

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear multicast bandwidth-admission &lt;group <i>group-address</i>&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;source <i>source-address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Reapply IP multicast bandwidth admissions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p><b>group <i>group-address</i></b>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p><b>inet</b>—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p><b>inet6</b>—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"><li>• If the interface is congested, and it was admitted previously, it is removed.</li><li>• If the interface was rejected previously, the <b>clear multicast bandwidth-admission</b> command enables the interface to be admitted as long as enough bandwidth exists on the interface.</li><li>• If you do not specify an interface, issuing the <b>clear multicast bandwidth-admission</b> command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.</li></ul> <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p><b>source <i>source-address</i></b>—(Optional) Use with the <b>group</b> option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



**Related Documentation** • [show multicast interface on page 4558](#)

**List of Sample Output** [clear multicast bandwidth-admission on page 4501](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

## clear multicast scope

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4502</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4502</a>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                  | Clear IP multicast scope statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast scope statistics.</p> <p><b>inet</b>—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast scope statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show multicast scope on page 4579</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast scope on page 4502</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### clear multicast scope

```
user@host> clear multicast scope
```

## clear multicast sessions

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4503</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4503</a>                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | clear multicast sessions<br><logical-system (all   <i>logical-system-name</i> )><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear multicast sessions<br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                     |
| <b>Description</b>                                  | Clear IP multicast sessions.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast sessions.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>regular-expression</i></b>—(Optional) Clear only multicast sessions that contain the specified regular expression.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show multicast sessions on page 4581</a></li> </ul>                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast sessions on page 4503</a>                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                        |

## Sample Output

### clear multicast sessions

```
user@host> clear multicast sessions
```

## clear multicast statistics

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4504</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4504</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                       | <pre>clear multicast statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear multicast statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                  | Clear IP multicast statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—Clear multicast statistics for all supported address families on all interfaces.</p> <p><b>inet</b>—(Optional) Clear multicast statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast statistics for IPv6 family addresses.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear multicast statistics for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show multicast statistics</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast statistics on page 4504</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### clear multicast statistics

```
user@host> clear multicast statistics
```

## clear pim join

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4505</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4505</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>clear pim join &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim join &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear the Protocol Independent Multicast (PIM) join and prune states.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                      | <p><b>none</b>—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p><b>group-address</b>—(Optional) Clear the PIM join and prune states for a group address.</p> <p><b>inet   inet6</b>—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim join</b> command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show pim join on page 4592</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">clear pim join on page 4506</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

clear pim join

```
user@host> clear pim join
```

## clear pim register

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4507</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4507</a><br><a href="#">Syntax (PTX Series) on page 4507</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (PTX Series)</b>                          | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) register message counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <p><b>none</b>—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM register message counters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim register</b> command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Related Documentation** • [show pim statistics on page 4615](#)

**List of Sample Output** [clear pim register on page 4508](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear pim register](#)

```
user@host> clear pim register
```



## clear pim statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4509</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4509</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                       | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM statistics for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim statistics</b> command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show pim statistics on page 4615</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">clear pim statistics on page 4510</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                                | See <a href="#">show pim statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 2111 4222 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 14200 13115 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
PIM statistics summary for all interfaces:
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Intf disabled 2007
Rx V1 Require V2 0
Rx Register not RP 0
RP Filtered Source 0
Unknown Reg Stop 0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 1 0 0
V1 Register 0 0 0
...
```



## mtrace

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mtrace source</code><br><code>&lt;logical-system logical-system-name&gt;</code><br><code>&lt;routing-instance routing-instance-name&gt;</code>                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 12.3 for the PTX Series. |
| <b>Description</b>              | Display trace information about an IP multicast path.                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>source</b> —Source hostname or address.<br><br><b>logical-system (logical-system-name)</b> —(Optional) Perform this operation on a logical system.<br><br><b>routing-instance routing-instance-name</b> —(Optional) Trace a particular routing instance.                                                                                                            |
| <b>Additional Information</b>   | The <b>mtrace</b> command for multicast traffic is similar to the <b>traceroute</b> command used for unicast traffic. Unlike <b>traceroute</b> , <b>mtrace</b> traces traffic backwards, from the receiver to the source.                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">mtrace source on page 4514</a>                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 327 on page 4512</a> describes the output fields for the <b>mtrace</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                              |

**Table 327: mtrace Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b>number-of-hops</b>             | Number of hops from the source to the named router or switch. |
| <b>router-name</b>                | Name of the router or switch for this hop.                    |
| <b>address</b>                    | Address of the router or switch for this hop.                 |

Table 327: mtrace Output Fields (*continued*)

| Field Name      | Field Description                              |
|-----------------|------------------------------------------------|
| <i>protocol</i> | Protocol used (for example, PIM).              |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of    | Time-to-live (TTL) threshold.                  |

## Sample Output

### mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
 -1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
 -2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
 -3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

## mtrace from-source

**Syntax** `mtrace from-source source source`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<ttl ttl>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

**Options** **brief | detail**—(Optional) Display the specified level of output.

**extra-hops *extra-hops***—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

**group *group***—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

**interval *interval***—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

**loop**—(Optional) Loop indefinitely, displaying rate and loss statistics.

**max-hops *max-hops***—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

**max-queries *max-queries***—(Optional) Maximum number of query attempts for any hop. The range of values is 1 through **32**. The default is **3**.

**multicast-response**—(Optional) Always request the response using multicast.

**no-resolve**—(Optional) Do not attempt to display addresses symbolically.

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance** *routing-instance-name*—(Optional) Trace a particular routing instance.

**source** *source*—Source hostname or address.

**ttl** *tll*—(Optional) IP time-to-live (TTL) value. You can specify a number between 0 and 255. Local queries to the multicast group use a value of 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time** *wait-time*—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level** view

**List of Sample Output** [mtrace from-source on page 4517](#)

**Output Fields** [Table 328 on page 4516](#) describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

**Table 328: mtrace from-source Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b>number-of-hops</b>             | Number of hops from the source to the named router or switch. |
| <b>router-name</b>                | Name of the router or switch for this hop.                    |
| <b>address</b>                    | Address of the router or switch for this hop.                 |
| <b>protocol</b>                   | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |
| <b>source</b>                     | Source address.                                               |
| <b>Response Dest</b>              | Response destination address.                                 |
| <b>Overall</b>                    | Average packet rate for all traffic at each hop.              |



Table 328: mtrace from-source Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Packet Statistics for Traffic From</b> | Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop. |
| <b>Receiver</b>                           | IP address receiving the multicast.                                                                              |
| <b>Query source</b>                       | IP address sending the mtrace query.                                                                             |

## Sample Output

### mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source Response Dest Overall Packet Statistics For Traffic From
192.1.4.2 192.1.1.2 Packet 192.1.4.2 To 225.1.1.1
 v ___/ rtt 2 ms Rate Lost/Sent = Pct Rate
192.1.2.1
192.1.3.2 routerC.lab.mycompany.net
 v ^ ttl 2 0/0 = -- 0 pps
192.1.4.1
192.1.2.2 routerB.lab.mycompany.net
 v __ ttl 3 ?/0 0 pps
192.1.1.2 192.1.1.2
Receiver Query Source

```

## mtrace monitor

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mtrace monitor                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series. |
| <b>Description</b>              | Listen passively for IP multicast responses. To exit the <b>mtrace monitor</b> command, type Ctrl+c.                                                                                     |
| <b>Options</b>                  | <b>none</b> —Trace the master instance.                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">mtrace monitor on page 4519</a>                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 329 on page 4518</a> describes the output fields for the <b>mtrace monitor</b> command. Output fields are listed in the approximate order in which they appear.        |

**Table 329: mtrace monitor Output Fields**

| Field Name              | Field Description                                             |
|-------------------------|---------------------------------------------------------------|
| <b>Mtrace query at</b>  | Date and time of the query.                                   |
| <b>by</b>               | Address of the host issuing the query.                        |
| <b>resp to</b>          | Response destination.                                         |
| <b>qid</b>              | Query ID number.                                              |
| <b>packet from...to</b> | IP address of the query source and default group destination. |
| <b>from...to</b>        | IP address of the multicast source and the response address.  |
| <b>via group</b>        | IP address of the group to trace.                             |
| <b>mxhop</b>            | Maximum hop setting.                                          |

## Sample Output

### mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

## mtrace to-gateway

---

**Syntax** `mtrace to-gateway gateway gateway`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interface interface-name>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<ttl ttl>`  
`<unicast-response>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about a multicast path from this router or switch to a gateway router or switch.

**Options** `gateway gateway`—Send the trace query to a gateway multicast address.

`brief | detail`—(Optional) Display the specified level of output.

`extra-hops extra-hops`—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.

`group group`—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

`interface interface-name`—(Optional) Source address for sending the trace query.

`interval interval`—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.

`loop`—(Optional) Loop indefinitely, displaying rate and loss statistics.

`max-hops max-hops`—(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**. The default value is **32**.

`max-queries max-queries`—(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**. The default value is **3**.

`multicast-response`—(Optional) Always request the response using multicast.

`no-resolve`—(Optional) Do not attempt to display addresses symbolically.

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance *routing-instance-name***—(Optional) Trace a particular routing instance.

**ttl *tll***—(Optional) IP time-to-live value. You can specify a number between 0 and 225.

Local queries to the multicast group use TTL 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time *wait-time***—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level** view

**List of Sample Output** [mtrace to-gateway on page 4521](#)

**Output Fields** [Table 330 on page 4521](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

**Table 330: mtrace to-gateway Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b><i>number-of-hops</i></b>      | Number of hops from the source to the named router or switch. |
| <b><i>router-name</i></b>         | Name of the router or switch for this hop.                    |
| <b><i>address</i></b>             | Address of the router or switch for this hop.                 |
| <b><i>protocol</i></b>            | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |

## Sample Output

### mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief
```

```
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
```

```
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerA.lab.mycompany.net (192.1.1.2) PIM thresh^ 1
-2 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-3 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

## show configuration protocols igmp

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show configuration protocols igmp                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                           |
| <b>Description</b>              | Display Internet Group Management Protocol (IGMP) information.                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">IGMP Snooping Overview on page 4232</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show configuration protocols igmp on page 4523</a>                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 331 on page 4523</a> describes the output fields for the <b>show configuration protocols igmp</b> command that relate to IGMP querying.                 |

**Table 331: show igmp group Output Fields**

| Field Name              | Field Description                                                                                        | Level of Output |
|-------------------------|----------------------------------------------------------------------------------------------------------|-----------------|
| accounting              | Enables notification for join and leave events.                                                          | All levels      |
| igmp-querier            | Configured source address for the IGMP querier.                                                          | All levels      |
| interface               | Name of the interface that receives IGMP membership reports.                                             | All levels      |
| query-interval          | Interval at which the IGMP querier sends general host-query messages to solicit membership information.  | All levels      |
| query-response-interval | How long the IGMP querier waits to receive a response from a query message before sending another query. | All levels      |
| src-address             | Source address of IGMP queries.                                                                          |                 |
| version                 | IGMP version.                                                                                            | All levels      |

## Sample Output

### show configuration protocols igmp

```

user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
 version 2;
}
igmp-querier {

```

```
src-address 10.0.0.2;
}
```



## show igmp group

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4525</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4525</a>                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) group membership information.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about membership for all IGMP groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group membership for the specified IP address only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 4491</a></li> </ul>                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <a href="#">show igmp group (Include Mode) on page 4526</a><br><a href="#">show igmp group (Exclude Mode) on page 4527</a><br><a href="#">show igmp group brief on page 4527</a><br><a href="#">show igmp group detail on page 4527</a>                                                                                                                                                                                   |
| <b>Output Fields</b>                                | <p><a href="#">Table 331 on page 4523</a> describes the output fields for the <b>show igmp group</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                 |

**Table 332: show igmp group Output Fields**

| Field Name        | Field Description                                                                                                                                       | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>  | Name of the interface that received the IGMP membership report. A name of <b>local</b> indicates that the local routing device joined the group itself. | All levels      |
| <b>Group</b>      | Group address.                                                                                                                                          | All levels      |
| <b>Group Mode</b> | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                                                  | All levels      |

Table 332: show igmp group Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                         | Level of Output |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Source           | Source address.                                                                                                                                                                                                           | All levels      |
| Source timeout   | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.      | detail          |
| Last reported by | Address of the host that last reported membership in this group.                                                                                                                                                          | All levels      |
| Timeout          | Time remaining until the group membership is removed.                                                                                                                                                                     | brief none      |
| Group timeout    | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail          |
| Type             | Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>                                             | All levels      |

## Sample Output

### show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

```

Group: 224.0.0.22
Source: 0.0.0.0
Last reported by: Local
Timeout: 0 Type: Dynamic

```

### show igmp group (Exclude Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: 224.0.0.22
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

### show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

### show igmp group detail

```

user@host> show igmp group detail
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Group mode: Exclude

```

```
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout: 0 Type: Dynamic
Group: 224.0.0.22
Group mode: Exclude
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout: 0 Type: Dynamic
```

## show igmp interface

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                 | <a href="#">Syntax on page 4529</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 4529</a>                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                         | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches and the QFX Series)</b> | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                            | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                             |
| <b>Description</b>                                    | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                        | <p><b>none</b>—Display standard information about all IGMP-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                       | view                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                          | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 4491</a></li> </ul>                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                          | <a href="#">show igmp interface on page 4531</a><br><a href="#">show igmp interface brief on page 4532</a><br><a href="#">show igmp interface detail on page 4532</a><br><a href="#">show igmp interface &lt;interface-name&gt; on page 4532</a>                                                                                                                                                                                    |
| <b>Output Fields</b>                                  | <p><a href="#">Table 333 on page 4529</a> describes the output fields for the <b>show igmp interface</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                       |

**Table 333: show igmp interface Output Fields**

| Field Name | Field Description                                                               | Level of Output |
|------------|---------------------------------------------------------------------------------|-----------------|
| Interface  | Name of the interface.                                                          | All levels      |
| Querier    | Address of the routing device that has been elected to send membership queries. | All levels      |

Table 333: show igmp interface Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>              | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>SSM Map Policy</b>     | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Timeout</b>            | How long until the IGMP querier is declared to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Version</b>            | IGMP version being used on the interface: <b>1</b> , <b>2</b> , or <b>3</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Groups</b>             | Number of groups on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Group limit</b>        | Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Group threshold</b>    | Configured threshold at which a warning message is generated.<br><br>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Group log-interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Immediate Leave</b>    | State of the immediate leave option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.</li> <li>• <b>Off</b>—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>Promiscuous Mode</b>   | State of the promiscuous mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces.</li> <li>• <b>Off</b>—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Passive</b>            | State of the passive mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves.</li> <li>• <b>Off</b>—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li>• <b>send-general-query</b>—The interface sends general queries.</li> <li>• <b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li>• <b>allow-receive</b>—The interface receives control traffic.</li> </ul> | All levels      |

Table 333: show igmp interface Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| OIF map               | Name of the OIF map (if configured) associated with the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| SSM map               | Name of the source-specific multicast (SSM) map (if configured) used on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| Configured Parameters | Information configured by the user: <ul style="list-style-type: none"> <li>• <b>IGMP Query Interval</b>—Interval (in seconds) at which this router sends membership queries when it is the querier.</li> <li>• <b>IGMP Query Response Interval</b>—Time (in seconds) that the router waits for a report in response to a general query.</li> <li>• <b>IGMP Last Member Query Interval</b>—Time (in seconds) that the router waits for a report in response to a group-specific query.</li> <li>• <b>IGMP Robustness Count</b>—Number of times the router retries a query.</li> </ul> | All levels      |
| Derived Parameters    | Derived information: <ul style="list-style-type: none"> <li>• <b>IGMP Membership Timeout</b>—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.</li> <li>• <b>IGMP Other Querier Present Timeout</b>—Time (in seconds) that the router waits for the IGMP querier to send a query.</li> </ul>                                                                                                                                                                                  | All levels      |

## Sample Output

### show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
 Querier: 10.111.30.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
 Querier: 10.111.10.1
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
 Querier: 10.111.20.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

### show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 4531](#).

### show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 4531](#).

### show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout: None Version: 3 Groups: 1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```



## show igmp statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4533</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4533</a>                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                        |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p><b>none</b>—Display IGMP statistics for all interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP statistics about the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp statistics on page 4495</a></li> </ul>                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show igmp statistics on page 4534</a><br><a href="#">show igmp statistics interface on page 4535</a>                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                                | <p><a href="#">Table 334 on page 4533</a> describes the output fields for the <b>show igmp statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                 |

**Table 334: show igmp statistics Output Fields**

| Field Name             | Field Description                                                                          |
|------------------------|--------------------------------------------------------------------------------------------|
| IGMP packet statistics | Heading for IGMP packet statistics for all interfaces or for the specified interface name. |

Table 334: show igmp statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IGMP Message type</b>      | <p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> <li>• <b>Membership Query</b>—Number of membership queries sent and received.</li> <li>• <b>V1 Membership Report</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>DVMRP</b>—Number of DVMRP messages sent or received.</li> <li>• <b>PIM V1</b>—Number of PIM version 1 messages sent or received.</li> <li>• <b>Cisco Trace</b>—Number of Cisco trace messages sent or received.</li> <li>• <b>V2 Membership Report</b>—Number of version 2 membership reports sent or received.</li> <li>• <b>Group Leave</b>—Number of group leave messages sent or received.</li> <li>• <b>Mtrace Response</b>—Number of Mtrace response messages sent or received.</li> <li>• <b>Mtrace Request</b>—Number of Mtrace request messages sent or received.</li> <li>• <b>Domain Wide Report</b>—Number of domain-wide reports sent or received.</li> <li>• <b>V3 Membership Report</b>—Number of version 3 membership reports sent or received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>IGMP v3 unsupported type</b>—Number of messages received with unknown and unsupported IGMP version 3 message types.</li> <li>• <b>IGMP v3 source required for SSM</b>—Number of IGMP version 3 messages received that contained no source.</li> <li>• <b>IGMP v3 mode not applicable for SSM</b>—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul> |
| <b>Received</b>               | Number of messages received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Sent</b>                   | Number of messages sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Rx errors</b>              | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>IGMP Global Statistics</b> | <p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with a bad IP checksum. No further classification was performed.</li> <li>• <b>Bad Receive If</b>—Number of messages received on an interface not enabled for IGMP.</li> <li>• <b>Rx non-local</b>—Number of messages received from senders that are not local.</li> <li>• <b>Timed out</b>—Number of groups that timed out as a result of not receiving an explicit leave message.</li> <li>• <b>Rejected Report</b>—Number of reports dropped because of the IGMP group policy.</li> <li>• <b>Total Interfaces</b>—Number of interfaces configured to support IGMP.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type Received Sent Rx errors
Membership Query 8883 459 0
V1 Membership Report 0 0 0

```

|                                     |      |   |   |
|-------------------------------------|------|---|---|
| DVMRP                               | 0    | 0 | 0 |
| PIM V1                              | 0    | 0 | 0 |
| Cisco Trace                         | 0    | 0 | 0 |
| V2 Membership Report                | 0    | 0 | 0 |
| Group Leave                         | 0    | 0 | 0 |
| Mtrace Response                     | 0    | 0 | 0 |
| Mtrace Request                      | 0    | 0 | 0 |
| Domain Wide Report                  | 0    | 0 | 0 |
| V3 Membership Report                | 0    | 0 | 0 |
| Other Unknown types                 |      |   | 0 |
| IGMP v3 unsupported type            |      |   | 0 |
| IGMP v3 source required for SSM     |      |   | 0 |
| IGMP v3 mode not applicable for SSM |      |   | 0 |
| IGMP Global Statistics              |      |   |   |
| Bad Length                          | 0    |   |   |
| Bad Checksum                        | 0    |   |   |
| Bad Receive If                      | 0    |   |   |
| Rx non-local                        | 1227 |   |   |
| Timed out                           | 0    |   |   |
| Rejected Report                     | 0    |   |   |
| Total Interfaces                    | 2    |   |   |

#### show igmp statistics interface

```

user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type Received Sent Rx errors
Membership Query 0 230 0
V1 Membership Report 0 0 0

```

## show igmp-snooping membership

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping membership &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;vlan <i>vlan-id</i>   <i>vlan-name</i>&gt;</pre>                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display IGMP snooping membership information.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP snooping information for the specified interface.</p> <p><b>vlan <i>vlan-id</i>   <i>vlan-name</i></b>—(Optional) Display IGMP snooping information for the specified VLAN.</p>     |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 4487</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping route on page 4539</a></li> <li>• <a href="#">show igmp-snooping statistics on page 4541</a></li> <li>• <a href="#">show igmp-snooping vlans on page 4543</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping membership on page 4537</a></p> <p><a href="#">show igmp-snooping membership detail on page 4538</a></p>                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <p><a href="#">Table 335 on page 4536</a> lists the output fields for the <b>show igmp-snooping membership</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                 |

**Table 335: show igmp-snooping membership Output Fields**

| Field Name | Field Description                 | Level of Output |
|------------|-----------------------------------|-----------------|
| VLAN       | Name of the VLAN.                 | All             |
| Interfaces | Interfaces assigned to the VLAN.  | All             |
| Tag        | Numerical identifier of the VLAN. | <b>detail</b>   |

Table 335: show igmp-snooping membership Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                     | Level of Output |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Router interfaces   | Names of multicast router interfaces.                                                                                                                                                                 | <b>detail</b>   |
| • static or dynamic | Whether the multicast router interface is <b>static</b> or <b>dynamic</b> .                                                                                                                           | <b>detail</b>   |
| • Uptime            | For static interfaces, length of time since the interface was configured as a multicast router interface; for dynamic interfaces, length of time since the first query was received on the interface. | <b>detail</b>   |
| • timeout           | Query timeout in seconds.                                                                                                                                                                             | <b>detail</b>   |
| Group               | IP multicast address of the multicast group.                                                                                                                                                          | <b>detail</b>   |
| Receiver count      | Number of interfaces that have membership in a multicast group.                                                                                                                                       | <b>detail</b>   |
| Flags               | IGMP version of the host sending a join message.                                                                                                                                                      | <b>detail</b>   |
| Uptime              | Length of time a multicast group has been active on the interface.                                                                                                                                    | <b>detail</b>   |
| timeout             | Time (in seconds) left until the entry for the multicast group is removed.                                                                                                                            | All             |
| Last reporter       | Last host to report membership for the multicast group.                                                                                                                                               | <b>detail</b>   |
| Include source      | Source addresses from which multicast streams are allowed based on IGMPv3 reports.                                                                                                                    | <b>detail</b>   |

## Sample Output

### show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: v1
 224.1.1.1 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.3 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.5 * 258 secs
 Interfaces: ge-0/0/0.0
 224.1.1.7 * 258 secs

```

```
Interfaces: ge-0/0/0.0
224.1.1.9 * 258 secs
Interfaces: ge-0/0/0.0
224.1.1.11 * 258 secs
Interfaces: ge-0/0/0.0
```

### show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.2
Receiver count: 1, Flags: <V3-hosts>
 ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
 Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
Group: 225.0.0.1
Receiver count: 1, Flags: <V2-hosts>
 ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
Router interfaces:
 ge-0/0/3.0 static Uptime: 00:08:45
 ge-0/0/2.0 static Uptime: 00:08:45
 ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
Group: 225.0.0.3
Receiver count: 1, Flags: <V3-hosts>
 ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
Group: 225.1.1.1
Receiver count: 1, Flags: <V2-hosts>
 ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
Group: 225.2.2.2
Receiver count: 1, Flags: <V2-hosts Static>
 ge-0/0/5.0 Uptime: 00:23:13
```

## show igmp-snooping route

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping route &lt;brief   detail&gt; &lt;ethernet-switching &lt;brief   detail   vlan (vlan-id   vlan-name )&gt;&gt; &lt;inet &lt;brief   detail   vlan vlan-name&gt;&gt; &lt;vlan vlan-name&gt;</pre>                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display IGMP snooping route information.                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>ethernet-switching</b>—(Optional) Display Ethernet switching information.</p> <p><b>inet</b>—(Optional) Display <b>inet</b> information.</p> <p><b>vlan vlan-name</b>—(Optional) Display route information for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 4487</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping statistics on page 4541</a></li> <li>• <a href="#">show igmp-snooping vlans on page 4543</a></li> </ul>                                                         |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping route on page 4540</a></p> <p><a href="#">show igmp-snooping route vlan v1 on page 4540</a></p>                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <p><a href="#">Table 336 on page 4539</a> lists the output fields for the <b>show igmp-snooping route</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                             |

**Table 336: show igmp-snooping route Output Fields**

| Field Name | Field Description                              |
|------------|------------------------------------------------|
| Table      | (For internal use only. Value is always 0.)    |
| VLAN       | Name of the VLAN.                              |
| Group      | Multicast group address.                       |
| Interfaces | Interfaces on which IGMP packets were snooped. |
| Next-hop   | ID associated with the next-hop device.        |

## Sample Output

### show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN Group Next-hop
V11 224.1.1.1, * 533
 Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN Group Next-hop
v12 224.1.1.3, * 534
 Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

### show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN Group Next-hop
v1 224.1.1.1, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.3, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.5, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.7, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.9, * 1266
 Interfaces: ge-0/0/0.0
v1 224.1.1.11, * 1266
 Interfaces: ge-0/0/0.0
```



## show igmp-snooping statistics

|                                 |                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show igmp-snooping statistics</b>                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display IGMP snooping statistics.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 4487</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping route on page 4539</a></li> <li>• <a href="#">show igmp-snooping vlans on page 4543</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show igmp-snooping statistics on page 4542</a>                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 337 on page 4541 lists the output fields for the <b>show igmp-snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                  |

**Table 337: show igmp-snooping statistics Output Fields**

| Field Name        | Field Description                                                                           |
|-------------------|---------------------------------------------------------------------------------------------|
| Bad length        | IGMP packet has illegal or bad length.                                                      |
| Bad checksum      | IGMP or IP checksum is incorrect.                                                           |
| Invalid interface | Packet was received through an invalid interface.                                           |
| Not local         | Number of packets received from senders that are not local.                                 |
| Receive unknown   | Unknown IGMP type.                                                                          |
| Timed out         | Number of timeouts for all multicast groups.                                                |
| IGMP Type         | Type of IGMP message ( <b>Queries</b> , <b>Reports</b> , <b>Leaves</b> , or <b>Other</b> ). |
| Received          | Number of IGMP packets received.                                                            |
| Transmitted       | Number of IGMP packets transmitted.                                                         |
| Recv Errors       | Number of general receive errors.                                                           |

## Sample Output

### show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
```

```
Bad length: 0 Bad checksum: 0 Invalid interface: 0
```

```
Not local: 0 Receive unknown: 0 Timed out: 58
```

| IGMP Type | Received | Transmitted | Recv Errors |
|-----------|----------|-------------|-------------|
| Queries:  | 74295    | 0           | 0           |
| Reports:  | 18148423 | 0           | 16333523    |
| Leaves:   | 0        | 0           | 0           |
| Other:    | 0        | 0           | 0           |

## show igmp-snooping vlans

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show igmp-snooping vlans</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;vlan <i>vlan-id</i>   <i>vlan-name</i>&gt;</code>                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display IGMP snooping VLAN information.                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>vlan <i>vlan-id</i>   vlan <i>vlan-number</i></b>—(Optional) Display VLAN information for the specified VLAN.</p>                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 4487</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4317</a></li> <li>• <a href="#">show igmp-snooping route on page 4539</a></li> <li>• <a href="#">show igmp-snooping statistics on page 4541</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping vlans on page 4544</a></p> <p><a href="#">show igmp-snooping vlans vlan on page 4544</a></p> <p><a href="#">show igmp-snooping vlans vlan detail on page 4544</a></p>                                                                                                            |
| <b>Output Fields</b>            | Table 338 on page 4543 lists the output fields for the <b>show igmp-snooping vlans</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                            |

**Table 338: show igmp-snooping vlans Output Fields**

| Field Name             | Field Description                                                       | Level of Output |
|------------------------|-------------------------------------------------------------------------|-----------------|
| <b>VLAN</b>            | Name of the VLAN.                                                       | All levels      |
| <b>IGMP-L2-Querier</b> | Source address for IGMP snooping queries (if switch is an IGMP querier) | All levels      |
| <b>Interfaces</b>      | Number of interfaces in the VLAN.                                       | All levels      |
| <b>Groups</b>          | Number of groups in the VLAN.                                           | All levels      |
| <b>MRouters</b>        | Number of multicast routers associated with the VLAN.                   | All levels      |
| <b>Receivers</b>       | Number of host receivers in the VLAN.                                   | All levels      |

Table 338: show igmp-snooping vlans Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                          | Level of Output |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Tag                | Numerical identifier of the VLAN.                                                                                                                                                                          | detail          |
| tagged   untagged  | Interface participates in a tagged (802.1Q) or untagged (native) VLAN.                                                                                                                                     | detail          |
| vlan-interface     | Internal VLAN interface identifier.                                                                                                                                                                        | detail          |
| Membership timeout | Membership timeout value.                                                                                                                                                                                  | detail          |
| Querier timeout    | Timeout value for interfaces dynamically marked as router or switch interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface. | detail          |
| Interface          | Name of the interface.                                                                                                                                                                                     | detail          |
| Reporters          | Number of dynamic groups on an interface.                                                                                                                                                                  | detail          |

## Sample Output

### show igmp-snooping vlans

```

user@switch> show igmp-snooping vlans
VLAN Interfaces Groups MRouters Receivers
default 0 0 0 0
v1 11 50 0 0
v10 1 0 0 0
v11 1 0 0 0
v180 3 0 1 0
v181 3 0 0 0
v182 3 0 0 0

```

### show igmp-snooping vlans vlan

```

user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN Interfaces Groups MRouters Receivers
v10 1 0 0 0

```

### show igmp-snooping vlans vlan detail

```

user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
 Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2

```

## show msdp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show msdp<br><brief   detail><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display Multicast Source Discovery Protocol (MSDP) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display standard MSDP information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display information about the specified peer only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp source on page 4547</a></li> <li>• <a href="#">show msdp source-active on page 4549</a></li> <li>• <a href="#">show msdp statistics on page 4552</a></li> </ul>                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show msdp on page 4546</a><br><a href="#">show msdp brief on page 4546</a><br><a href="#">show msdp detail on page 4546</a>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 339 on page 4545</a> describes the output fields for the <b>show msdp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                |

**Table 339: show msdp Output Fields**

| Field Name    | Field Description                                                                        | Level of Output |
|---------------|------------------------------------------------------------------------------------------|-----------------|
| Peer address  | IP address of the peer.                                                                  | All levels      |
| Local address | Local address of the peer.                                                               | All levels      |
| State         | Status of the MSDP connection: <b>Listen</b> , <b>Established</b> , or <b>Inactive</b> . | All levels      |
| Last up/down  | Time at which the most recent peer-state change occurred.                                | All levels      |

Table 339: show msdp Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                   | Level of Output |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Peer-Group           | Peer group name.                                                                                                                                                                    | All levels      |
| SA Count             | Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> . | All levels      |
| Peer Connect Retries | Number of peer connection retries.                                                                                                                                                  | detail          |
| State timer expires  | Number of seconds before another message is sent to a peer.                                                                                                                         | detail          |
| Peer Times out       | Number of seconds to wait for a response from the peer before the peer is declared unavailable.                                                                                     | detail          |
| SA accepted          | Number of entries in the source-active cache accepted from the peer.                                                                                                                | detail          |
| SA received          | Number of entries in the source-active cache received by the peer.                                                                                                                  | detail          |

## Sample Output

### show msdp

```

user@host> show msdp
Peer address Local address State Last up/down Peer-Group SA Count
198.32.8.193 198.32.8.195 Established 5d 19:25:44 North23 120/150
198.32.8.194 198.32.8.195 Established 3d 19:27:27 North23 300/345
198.32.8.196 198.32.8.195 Established 5d 19:39:36 North23 10/13
198.32.8.197 198.32.8.195 Established 5d 19:32:27 North23 5/6
198.32.8.198 198.32.8.195 Established 3d 19:33:04 North23 2305/3000

```

### show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 4546](#).

### show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

## show msdp source

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show msdp source &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-address&gt;</pre>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Display standard MSDP source information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-address</b>—(Optional) IP address and optional prefix length. Display information for the specified source address only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp on page 4545</a></li> <li>• <a href="#">show msdp source-active on page 4549</a></li> <li>• <a href="#">show msdp statistics on page 4552</a></li> </ul>                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show msdp source on page 4548</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Output Fields** Table 340 on page 4548 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

**Table 340: show msdp source Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source address | IP address of the source.                                                                                                                                                                                                                                               |
| /Len           | Length of the prefix for this IP address.                                                                                                                                                                                                                               |
| Type           | Discovery method for this multicast source: <ul style="list-style-type: none"> <li>• <b>Configured</b>—Source-active limit explicitly configured for this source.</li> <li>• <b>Dynamic</b>—Source-active limit established when this source was discovered.</li> </ul> |
| Maximum        | Source-active limit applied to this source.                                                                                                                                                                                                                             |
| Threshold      | Source-active threshold applied to this source.                                                                                                                                                                                                                         |
| Exceeded       | Number of source-active messages received from this source exceeding the established maximum.                                                                                                                                                                           |

## Sample Output

**show msdp source**

```

user@host> show msdp source
Source address /Len Type Maximum Threshold Exceeded
0.0.0.0 /0 Configured 5 none 0
10.1.0.0 /16 Configured 500 none 0
10.1.1.1 /32 Configured 10000 none 0
10.1.1.2 /32 Dynamic 6936 none 0
10.1.5.5 /32 Dynamic 500 none 123
10.2.1.1 /32 Dynamic 2 none 0

```



## show msdp source-active

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show msdp source-active &lt;brief   detail&gt; &lt;group <i>group</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;local&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;originator <i>originator</i>&gt; &lt;peer <i>peer-address</i>&gt; &lt;source <i>source-address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Display the Multicast Source Discovery Protocol (MSDP) source-active cache.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display standard MSDP source-active cache information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group <i>group</i></b>—(Optional) Display source-active cache information for the specified group.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance.</p> <p><b>local</b>—(Optional) Display all source-active caches originated by this router.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>originator <i>originator</i></b>—(Optional) Display information about the peer that originated the source-active cache entries.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display the source-active cache of the specified peer.</p> <p><b>source <i>source-address</i></b>—(Optional) Display the source-active cache of the specified source.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp on page 4545</a></li> <li>• <a href="#">show msdp source on page 4547</a></li> <li>• <a href="#">show msdp statistics on page 4552</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show msdp source-active on page 4550</a></p> <p><a href="#">show msdp source-active brief on page 4550</a></p> <p><a href="#">show msdp source-active detail on page 4551</a></p> <p><a href="#">show msdp source-active source on page 4551</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | Table 341 on page 4550 describes the output fields for the <b>show msdp source-active</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 341: show msdp source-active Output Fields

| Field Name                              | Field Description                                                                                                                                          |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                                                   |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                                                        |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.                                         |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).                                |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                                                        |
| Group address                           | Multicast address of the group.                                                                                                                            |
| Source address                          | IP address of the source.                                                                                                                                  |
| Peer address                            | IP address of the peer.                                                                                                                                    |
| Originator                              | Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured. |
| Flags                                   | Flags: Accept, Reject, or Filtered.                                                                                                                        |

## Sample Output

### show msdp source-active

```

user@host> show msdp source-active
Group address Source address Peer address Originator Flags
230.0.0.0 192.168.195.46 local 10.255.14.30 Accept
230.0.0.1 192.168.195.46 local 10.255.14.30 Accept
230.0.0.2 192.168.195.46 local 10.255.14.30 Accept
230.0.0.3 192.168.195.46 local 10.255.14.30 Accept
230.0.0.4 192.168.195.46 local 10.255.14.30 Accept

```

### show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 4550](#).

### show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 4550](#).

### show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

| Group address | Source address  | Peer address   | Originator     | Flags  |
|---------------|-----------------|----------------|----------------|--------|
| 226.2.2.1     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.3     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.4     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.5     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.7     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.10    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.11    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.13    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.14    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.15    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |

## show msdp statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show msdp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Display statistics about Multicast Source Discovery Protocol (MSDP) peers.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display statistics about all MSDP peers for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics about a specific MSDP instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display statistics about a particular MSDP peer.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear msdp statistics on page 4499</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show msdp statistics on page 4554</a><br><a href="#">show msdp statistics peer on page 4554</a>                                                                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 342 on page 4552 describes the output fields for the <b>show msdp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                        |

**Table 342: show msdp statistics Output Fields**

| Field Name                              | Field Description                                                                                                           |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                    |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                         |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.          |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                         |
| Peer                                    | Address of peer.                                                                                                            |

Table 342: show msdp statistics Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                   |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Last State Change                   | How long ago the peer state changed.                                                                                                                |
| Last message received from the peer | How long ago the last message was received from the peer.                                                                                           |
| RPF Failures                        | Number of reverse path forwarding (RPF) failures.                                                                                                   |
| Remote Closes                       | Number of times the remote peer closed.                                                                                                             |
| Peer Timeouts                       | Number of peer timeouts.                                                                                                                            |
| SA messages sent                    | Number of source-active messages sent.                                                                                                              |
| SA messages received                | Number of source-active messages received.                                                                                                          |
| SA request messages sent            | Number of source-active request messages sent.                                                                                                      |
| SA request messages received        | Number of source-active request messages received.                                                                                                  |
| SA response messages sent           | Number of source-active response messages sent.                                                                                                     |
| SA response messages received       | Number of source-active response messages received.                                                                                                 |
| Active source exceeded              | Number of times this peer has exceeded configured source-active limits.                                                                             |
| Active source Maximum               | Configured number of active source messages accepted by this peer.                                                                                  |
| Active source threshold             | Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.                     |
| Active source log-warning           | Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Active source log-interval          | Time (in seconds) between consecutive log messages on this peer.                                                                                    |
| Keepalive messages sent             | Number of keepalive messages sent.                                                                                                                  |
| Keepalive messages received         | Number of keepalive messages received.                                                                                                              |
| Unknown messages received           | Number of unknown messages received.                                                                                                                |

Table 342: show msdp statistics Output Fields (*continued*)

| Field Name              | Field Description                  |
|-------------------------|------------------------------------|
| Error messages received | Number of error messages received. |

## Sample Output

### show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

### show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
 Last State Change: 8:19:23 (00:01:08)
 Last message received from peer: 8:20:05 (00:00:26)
 RPF Failures: 0
 Remote Closes: 0
 Peer Timeouts: 0
 SA messages sent: 17
 SA messages received: 16
 SA request messages sent: 0
 SA request messages received: 0
 SA response messages sent: 0
 SA response messages received: 0
 Active source exceeded: 20
 Active source Maximum: 10
 Active source threshold: 8
 Active source log-warning: 60
 Active source log-interval: 120
 Keepalive messages sent: 0

```

Keepalive messages received: 0  
Unknown messages received: 0  
Error messages received: 0

## show multicast flow-map

|                                                     |                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4556</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4556</a>                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>show multicast flow-map &lt;brief   detail&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast flow-map &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                          |
| <b>Description</b>                                  | Display configuration information about IP multicast flow maps.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about IP multicast flow maps on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show multicast flow-map on page 4557</a><br><a href="#">show multicast flow-map detail on page 4557</a>                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | <p><a href="#">Table 343 on page 4556</a> describes the output fields for the <b>show multicast flow-map</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                            |

**Table 343: show multicast flow-map Output Fields**

| Field Name           | Field Description                                         | Levels of Output |
|----------------------|-----------------------------------------------------------|------------------|
| <b>Name</b>          | Name of the flow map.                                     | All levels       |
| <b>Policy</b>        | Name of the policy associated with the flow map.          | All levels       |
| <b>Cache-timeout</b> | Cache timeout value assigned to the flow map.             | All levels       |
| <b>Bandwidth</b>     | Bandwidth setting associated with the flow map.           | All levels       |
| <b>Adaptive</b>      | Whether or not adaptive mode is enabled for the flow map. | none             |
| <b>Flow-map</b>      | Name of the flow map.                                     | <b>detail</b>    |



Table 343: show multicast flow-map Output Fields (*continued*)

| Field Name                | Field Description                                         | Levels of Output |
|---------------------------|-----------------------------------------------------------|------------------|
| <b>Adaptive Bandwidth</b> | Whether or not adaptive mode is enabled for the flow map. | <b>detail</b>    |
| <b>Redundant Sources</b>  | Redundant sources defined for the same destination group. | <b>detail</b>    |

## Sample Output

### show multicast flow-map

```

user@host> show multicast flow-map
Instance: master
Name Policy Cache timeout Bandwidth Adaptive
map2 policy2 never 2000000 no
map1 policy1 60 seconds 2000000 no

```

## Sample Output

### show multicast flow-map detail

```

user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
 Policy: policy1
 Cache Timeout: 600 seconds
 Bandwidth: 2000000
 Adaptive Bandwidth: yes
 Redundant Sources: 11.11.11.11
 Redundant Sources: 11.11.11.12
 Redundant Sources: 11.11.11.13

```

## show multicast interface

|                                                     |                                                                                                                                                                                                                                  |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4558</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4558</a>                                                                                                                 |
| <b>Syntax</b>                                       | show multicast interface<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                 |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast interface                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 8.3.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                             |
| <b>Description</b>                                  | Display bandwidth information about IP multicast interfaces.                                                                                                                                                                     |
| <b>Options</b>                                      | <b>none</b> —Display all interfaces that have multicast configured.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                        | <a href="#">show multicast interface on page 4559</a>                                                                                                                                                                            |
| <b>Output Fields</b>                                | <a href="#">Table 344 on page 4558</a> describes the output fields for the <b>show multicast interface</b> command. Output fields are listed in the approximate order in which they appear.                                      |

**Table 344: show multicast interface Output Fields**

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>                        | Name of the multicast interface.                                                                                                                                                                                                                                                                                                                            |
| <b>Maximum bandwidth (bps)</b>          | Maximum bandwidth setting, in bits per second, for this interface.                                                                                                                                                                                                                                                                                          |
| <b>Remaining bandwidth (bps)</b>        | Amount of bandwidth, in bits per second, remaining on the interface.                                                                                                                                                                                                                                                                                        |
| <b>Mapped bandwidth deduction (bps)</b> | Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.<br><br><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.<br><br>This field does not appear in the output when the no QoS adjustment feature is disabled. |

Table 344: show multicast interface Output Fields (*continued*)

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local bandwidth deduction (bps)</b>       | <p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| <b>Reverse OIF mapping</b>                   | <p>State of the reverse OIF mapping feature (<b>on</b> or <b>off</b>).</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                                    |
| <b>Reverse OIF mapping no QoS adjustment</b> | <p>State of the no QoS adjustment feature (<b>on</b> or <b>off</b>) for interfaces that are using reverse OIF mapping.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                    |
| <b>Leave timer</b>                           | <p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                          |
| <b>No QoS adjustment</b>                     | <p>State (<b>on</b>) of the no QoS adjustment feature when this feature is enabled.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                       |

## Sample Output

### show multicast interface

```

user@host> show multicast interface
Interface Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3 10000000 0
fe-0/0/3.210 10000000 -2000000
fe-0/0/3.220 100000000 100000000
fe-0/0/3.230 20000000 18000000
fe-0/0/2.200 100000000 100000000

```

## show multicast minfo

|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show multicast minfo</code><br><code>&lt;host&gt;</code>                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                  |
| <b>Description</b>              | Display configuration information about IP multicast networks, including neighboring multicast router addresses.                                                                                                          |
| <b>Options</b>                  | <b>none</b> —Display configuration information about all multicast networks.<br><br><b>host</b> —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show multicast minfo on page 4561</a>                                                                                                                                                                         |
| <b>Output Fields</b>            | <a href="#">Table 345 on page 4560</a> describes the output fields for the <b>show multicast minfo</b> command. Output fields are listed in the approximate order in which they appear.                                   |

**Table 345: show multicast minfo Output Fields**

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>source-address</i>                | Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <i>ip-address-1—&gt;ip-address-2</i> | Queried router interface address and directly attached neighbor interface address, respectively.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <i>(name or ip-address)</i>          | Name or IP address of neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>[metric/threshold/type/flags]</i> | Neighbor's multicast profile: <ul style="list-style-type: none"> <li><b>metric</b>—Always has a value of 1, because <b>minfo</b> queries the directly connected interfaces of a device.</li> <li><b>threshold</b>—Multicast threshold time-to-live (TTL). The range of values is 0 through 255.</li> <li><b>type</b>—Multicast connection type: <b>pim</b> or <b>tunnel</b>.</li> <li><b>flags</b>—Flags for this route: <ul style="list-style-type: none"> <li><b>querier</b>—Queried router is the designated router for the neighboring session.</li> <li><b>leaf</b>—Link is a leaf in the multicast network.</li> <li><b>down</b>—Link status indicator.</li> </ul> </li> </ul> |

## Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
 192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
 0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

## show multicast next-hops

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4562</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4562</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p><b>detail</b> option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                  | Display the entries in the IP multicast next-hop table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p>When you include the <b>detail</b> option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form <b>fe-0/1/2.0-(1048574)</b> where <b>1048574</b> is the next-hop ID number.</p> <p><b>identifier-number</b>—(Optional) Show a particular next hop by ID number. The range of values is 1 through <b>65,535</b>.</p> <p><b>inet   inet6</b>—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show multicast next-hops on page 4563</a><br><a href="#">show multicast next-hops (Bidirectional PIM on page 4563</a><br><a href="#">show multicast next-hops brief on page 4564</a><br><a href="#">show multicast next-hops detail on page 4564</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Output Fields** Table 346 on page 4563 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

**Table 346: show multicast next-hops Output Fields**

| Field Name                     | Field Description                                                                                                                                      |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Family</b>                  | Protocol family (such as <b>INET</b> ).                                                                                                                |
| <b>ID</b>                      | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.                                        |
| <b>Refcount</b>                | Number of cache entries that are using this next hop.                                                                                                  |
| <b>KRefcount</b>               | Kernel reference count for the next hop.                                                                                                               |
| <b>Downstream interface</b>    | Interface names associated with each multicast next-hop ID.                                                                                            |
| <b>Incoming interface list</b> | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. |

## Sample Output

### show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID Refcount KRefcount Downstream interface
262142 4 2 so-1/0/0.0
262143 2 1 mt-1/1/0.49152
262148 2 1 mt-1/1/0.32769
```

### show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID Refcount KRefcount Downstream interface
2097151 8 4 ge-0/0/1.0

Family: INET6
ID Refcount KRefcount Downstream interface
2097157 2 1 ge-0/0/1.0

Family: Incoming interface list
ID Refcount KRefcount Downstream interface
513 5 2 lo0.0
 ge-0/0/1.0
514 5 2 lo0.0
 ge-0/0/1.0
 xe-4/1/0.0
515 3 1 lo0.0
 ge-0/0/1.0
 xe-4/1/0.0
544 1 0 lo0.0
 xe-4/1/0.0
```

### show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 4563](#).

### show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID Refcount KRefCount Downstream interface
1048577 2 1 fe-0/1/2.0-(1048574)
 ge-0/2/3.0-(1048576)
```



## show multicast pim-to-igmp-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4565</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4565</a>                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                       | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                    |
| <b>Description</b>                                  | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li><a href="#">Configuring PIM-to-IGMP and PIM-to-MLD Message Translation</a></li> </ul>                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-igmp-proxy on page 4566</a><br><a href="#">show multicast pim-to-igmp-proxy instance on page 4566</a>                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | <p><a href="#">Table 347 on page 4565</a> describes the output fields for the <b>show multicast pim-to-igmp-proxy</b> command. Output fields are listed in the order in which they appear.</p>                                                                                                                                                                                                                                                      |

**Table 347: show multicast pim-to-igmp-proxy Output Fields**

| Field Name         | Field Description                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>    | Routing instance. Default instance is <b>master</b> (inet.0 routing table).                                                                           |
| <b>Proxy state</b> | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |

Table 347: show multicast pim-to-igmp-proxy Output Fields (*continued*)

| Field Name            | Field Description                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------|
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured. |

## Sample Output

### show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

### show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

## show multicast pim-to-mld-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4567</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4567</a>                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                                       | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 9.6 for EX Series switches.<br><b>instance</b> option introduced in Junos OS Release 10.3.<br><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>                                  | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                      | <b>none</b> —Display configuration information about PIM-to-MLD message translation for all routing instances.<br><br><b>instance</b> <i>instance-name</i> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.<br><br><b>logical-system</b> (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-mld-proxy on page 4568</a><br><a href="#">show multicast pim-to-mld-proxy instance on page 4568</a>                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <a href="#">Table 348 on page 4567</a> describes the output fields for the <b>show multicast pim-to-mld-proxy</b> command. Output fields are listed in the order in which they appear.                                                                                                                                                                                                                                                        |

**Table 348: show multicast pim-to-mld-proxy Output Fields**

| Field Name            | Field Description                                                                                                                                   |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Proxy state</b>    | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.                                        |

## Sample Output

### show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

### show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

## show multicast route

**List of Syntax** [Syntax on page 4569](#)  
[Syntax \(EX Series Switch and the QFX Series\) on page 4569](#)

**Syntax** show multicast route  
 <brief | detail | extensive | summary>  
 <active | all | inactive>  
 <group *group*>  
 <inet | inet6>  
 <instance *instance name*>  
 <logical-system (all | *logical-system-name*)>  
 <*regular-expression*>  
 <source-prefix *source-prefix*>

**Syntax (EX Series Switch and the QFX Series)** show multicast route  
 <brief | detail | extensive | summary>  
 <active | all | inactive>  
 <group *group*>  
 <inet | inet6>  
 <instance *instance name*>  
 <*regular-expression*>  
 <source-prefix *source-prefix*>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
**inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.  
 Support for bidirectional PIM added in Junos OS Release 12.1.

**Description** Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

**Options** **none**—Display standard information about all entries in the multicast forwarding table for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**active | all | inactive**—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

**group *group***—(Optional) Display the cache entries for a particular group.

**inet | inet6**—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**regular-expression**—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

**source-prefix source-prefix**—(Optional) Display the cache entries for a particular source prefix.

**Required Privilege Level** view

**List of Sample Output**

- [show multicast route on page 4571](#)
- [show multicast route \(Bidirectional PIM\) on page 4572](#)
- [show multicast route brief on page 4572](#)
- [show multicast route detail on page 4572](#)
- [show multicast route extensive \(Bidirectional PIM\) on page 4573](#)
- [show multicast route instance <instance-name> on page 4574](#)
- [show multicast route summary on page 4574](#)

**Output Fields** [Table 349 on page 4570](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

**Table 349: show multicast route Output Fields**

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output         |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>family</b>                        | IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ).                                                                                                                                                                                                                                                                                                                                                          | All levels              |
| <b>Group</b>                         | Group address.<br><br>For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.                                                                                                                                                                                                                                                                                               | All levels              |
| <b>Source</b>                        | Prefix and length of the source as it is in the multicast forwarding table.                                                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Incoming interface list</b>       | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.                                                                                                                                                                                                                                                                                | All levels              |
| <b>Upstream interface</b>            | Name of the interface on which the packet with this source prefix is expected to arrive.                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Downstream interface list</b>     | List of interface names to which the packet with this source prefix is forwarded.                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>Number of outgoing interfaces</b> | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>Session description</b>           | Name of the multicast session.                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Statistics</b>                    | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays <b>Forwarding statistics are not available</b> .<br><br><b>NOTE:</b> On QFX Series switches, this field does not report valid statistics. | <b>detail extensive</b> |

Table 349: show multicast route Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                | Level of Output          |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Next-hop ID</b>                            | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the <b>show multicast nexthops</b> command.                                | <b>detail extensive</b>  |
| <b>Incoming interface list ID</b>             | For bidirectional PIM, incoming interface list identifier.<br><br>Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. | <b>detail extensive</b>  |
| <b>Upstream protocol</b>                      | Protocol running on the interface on which the packet with this source prefix is expected to arrive.                                                                                                                             | <b>detail extensive</b>  |
| <b>Route type</b>                             | Type of multicast route. Values can be (S,G) or (*G).                                                                                                                                                                            | <b>summary</b>           |
| <b>Route state</b>                            | Whether the group is <b>Active</b> or <b>Inactive</b> .                                                                                                                                                                          | <b>summary extensive</b> |
| <b>Route count</b>                            | Number of multicast routes.                                                                                                                                                                                                      | <b>summary</b>           |
| <b>Forwarding state</b>                       | Whether the prefix is pruned or forwarding.                                                                                                                                                                                      | <b>extensive</b>         |
| <b>Cache lifetime/timeout</b>                 | Number of seconds until the prefix is removed from the multicast forwarding table. A value of <b>never</b> indicates a permanent forwarding entry. A value of <b>forever</b> indicates routes that do not have keepalive times.  | <b>extensive</b>         |
| <b>Wrong incoming interface notifications</b> | Number of times that the upstream interface was not available.                                                                                                                                                                   | <b>extensive</b>         |
| <b>Uptime</b>                                 | Time since the creation of a multicast route.                                                                                                                                                                                    | <b>extensive</b>         |

## Sample Output

### show multicast route

```

user@host> show multicast route
Family: INET

Group: 228.0.0.0
 Source: 10.255.14.144/32
 Upstream interface: local
 Downstream interface list:
 so-1/0/0.0

Group: 239.1.1.1
 Source: 10.255.14.144/32
 Upstream interface: local
 Downstream interface list:
 so-1/0/0.0

Group: 239.1.1.1
 Source: 10.255.70.15/32
 Upstream interface: so-1/0/0.0

```

```
Downstream interface list:
mt-1/1/0.49152
```

```
Family: INET6
```

### show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET
```

```
Group: 224.1.1.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0

Group: 225.1.1.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0

Group: 225.1.3.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0
Family: INET6
```

### show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 4571](#) or [show multicast route \(Bidirectional PIM\) on page 4572](#).

### show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
 so-1/0/0.0
Session description: Unknown
Statistics: 8 kbps, 100 pps, 45272 packets
Next-hop ID: 262142
Upstream protocol: PIM

Group: 239.1.1.1
```



```

Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
 so-1/0/0.0
Session description: Administratively Scoped
Statistics: 0 kbps, 0 pps, 13404 packets
Next-hop ID: 262142
Upstream protocol: PIM

Group: 239.1.1.1
Source: 10.255.70.15/32
Upstream interface: so-1/0/0.0
Downstream interface list:
 mt-1/1/0.49152
Session description: Administratively Scoped
Statistics: 46 kbps, 1000 pps, 921077 packets

Next-hop ID: 262143
Upstream protocol: PIM

```

Family: INET6

#### show multicast route extensive (Bidirectional PIM)

```

user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 585
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Family: INET6

#### show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
Instance: v1 Family: INET
```

```
Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
 lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
 lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.3
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
 lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Instance: v1 Family: INET6
```

#### show multicast route summary

```
user@host> show multicast route summary
Instance: master Family: INET
```

| Route type | Route state | Route count |
|------------|-------------|-------------|
| (S,G)      | Active      | 2           |
| (S,G)      | Inactive    | 3           |

```
Instance: master Family: INET6
```

## show multicast rpf

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4575</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4575</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                       | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                  | Display information about multicast reverse-path-forwarding (RPF) calculations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display RPF calculation information for all supported address families.</p> <p><b>inet   inet6</b>—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix</b>—(Optional) Display the RPF calculation information for the specified prefix.</p> <p><b>summary</b>—(Optional) Display a summary of all multicast RPF information.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast rpf on page 4576</a><br><a href="#">show multicast rpf inet6 on page 4577</a><br><a href="#">show multicast rpf prefix on page 4578</a><br><a href="#">show multicast rpf summary on page 4578</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** Table 350 on page 4576 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

**Table 350: show multicast rpf Output Fields**

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance      | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                                                                                                                                                                                                                                             |
| Source prefix | Prefix and length of the source as it exists in the multicast forwarding table.                                                                                                                                                                                                                                                                                                                               |
| Protocol      | How the route was learned.                                                                                                                                                                                                                                                                                                                                                                                    |
| Interface     | Upstream RPF interface.<br><br><b>NOTE:</b> The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured. |
| Neighbor      | Upstream RPF neighbor.<br><br><b>NOTE:</b> The displayed neighbor information does not apply to bidirectional PIM. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured.                 |

## Sample Output

### show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
 Protocol: Static

10.255.14.132/32
 Protocol: Direct
 Interface: lo0.0

10.255.245.91/32
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

```

Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21

```

### show multicast rpf inet6

```

user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
 Protocol: Direct
 Interface: lo0.0

::10.255.245.91/128
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
 Protocol: Direct
 Interface: so-1/1/1.0

::192.168.195.22/128
 Protocol: Local

::192.168.195.36/126
 Protocol: IS-IS
 Interface: so-1/1/1.0
 Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
 Protocol: Direct
 Interface: fe-2/2/0.0

::192.168.195.77/128
 Protocol: Local

```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

#### show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
 Protocol: PIM

ff02::d/128
 Protocol: PIM

...
```

#### show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

## show multicast scope

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4579</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4579</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                  | Display administratively scoped IP multicast information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p><b>inet   inet6</b>—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show multicast scope on page 4580</a><br><a href="#">show multicast scope inet on page 4580</a><br><a href="#">show multicast scope inet6 on page 4580</a>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <p><a href="#">Table 351 on page 4579</a> describes the output fields for the <b>show multicast scope</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 351: show multicast scope Output Fields**

| Field Name          | Field Description                                           |
|---------------------|-------------------------------------------------------------|
| <b>Scope name</b>   | Name of the multicast scope.                                |
| <b>Group Prefix</b> | Range of multicast groups that are scoped.                  |
| <b>Interface</b>    | Interface that is the boundary of the administrative scope. |

Table 351: show multicast scope Output Fields (*continued*)

| Field Name      | Field Description                 |
|-----------------|-----------------------------------|
| Resolve Rejects | Number of kernel resolve rejects. |

## Sample Output

### show multicast scope

```
user@host> show multicast scope
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |

### show multicast scope inet

```
user@host> show multicast scope inet
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |

### show multicast scope inet6

```
user@host> show multicast scope inet6
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |



## show multicast sessions

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4581</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4581</a>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                                       | show multicast sessions<br><brief   detail   extensive><br><logical-system (all   <i>logical-system-name</i> )><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast sessions<br><brief   detail   extensive><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Display information about announced IP multicast sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about all multicast sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>regular-expression</i></b>—(Optional) Display information about announced sessions that match a UNIX-style regular expression.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">show multicast sessions on page 4582</a><br><a href="#">show multicast sessions regular-expression detail on page 4582</a>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>                                | <a href="#">Table 352 on page 4581</a> describes the output fields for the <b>show multicast sessions</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                 |

**Table 352: show multicast sessions Output Fields**

| Field Name          | Field Description                               |
|---------------------|-------------------------------------------------|
| <i>session-name</i> | Name of the known announced multicast sessions. |

## Sample Output

### show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.
```

### show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0 Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

## show multicast usage

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4584</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4584</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                                       | <code>show multicast usage</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <code>show multicast usage</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                  | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <b>none</b> —Display multicast usage information for all supported address families for all routing instances.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>inet   inet6</b> —(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>                        | <a href="#">show multicast usage on page 4585</a><br><a href="#">show multicast usage brief on page 4585</a><br><a href="#">show multicast usage instance on page 4585</a><br><a href="#">show multicast usage detail on page 4586</a>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                                | <a href="#">Table 353 on page 4585</a> describes the output fields for the <b>show multicast usage</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 353: show multicast usage Output Fields

| Field Name      | Field Description                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                        |
| <b>Group</b>    | Group address.                                                                                                                                                                           |
| <b>Sources</b>  | Number of sources.                                                                                                                                                                       |
| <b>Packets</b>  | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays <b>unavailable</b> . |
| <b>Bytes</b>    | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays <b>unavailable</b> .     |
| <b>Prefix</b>   | IP address.                                                                                                                                                                              |
| <b>/len</b>     | Prefix length.                                                                                                                                                                           |
| <b>Groups</b>   | Number of multicast groups.                                                                                                                                                              |

## Sample Output

### show multicast usage

```

user@host> show multicast usage
Group Sources Packets Bytes
228.0.0.0 1 52847 4439148
239.1.1.1 2 13450 1125530

Prefix /len Groups Packets Bytes
10.255.14.144 /32 2 66254 5561304
10.255.70.15 /32 1 43 3374...
```

### show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 4585](#).

### show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group Sources Packets Bytes
224.2.127.254 1 5538 509496
224.0.1.39 1 13 624
224.0.1.40 1 13 624

Prefix /len Groups Packets Bytes
192.168.195.34 /32 1 5538 509496
10.255.14.30 /32 1 13 624
```

```
10.255.245.91 /32 1 13 624
...
```

#### show multicast usage detail

```
user@host> show multicast usage detail
Group Sources Packets Bytes
228.0.0.0 1 53159 4465356
 Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356
239.1.1.1 2 13450 1125530
 Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156
 Source: 10.255.70.15 /32 Packets: 43 Bytes: 3374
```

```
Prefix /len Groups Packets Bytes
10.255.14.144 /32 2 66566 5587512
 Group: 228.0.0.0 Packets: 53159 Bytes: 4465356
 Group: 239.1.1.1 Packets: 13407 Bytes: 1122156
10.255.70.15 /32 1 43 3374
 Group: 239.1.1.1 Packets: 43 Bytes: 3374
```

## show pim bootstrap

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4587</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4587</a>                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                       | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                  |
| <b>Description</b>                                  | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM bootstrap router information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                        | <a href="#">show pim bootstrap on page 4588</a><br><a href="#">show pim bootstrap instance on page 4588</a>                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | <p><a href="#">Table 354 on page 4587</a> describes the output fields for the <b>show pim bootstrap</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                       |

**Table 354: show pim bootstrap Output Fields**

| Field Name           | Field Description                                                            |
|----------------------|------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance.                                                |
| <b>BSR</b>           | Bootstrap router.                                                            |
| <b>Pri</b>           | Priority of the routing device as elected to be the bootstrap router.        |
| <b>Local address</b> | Local routing device address.                                                |
| <b>Pri</b>           | Local routing device address priority to be elected as the bootstrap router. |

Table 354: show pim bootstrap Output Fields (*continued*)

| Field Name     | Field Description                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------|
| <b>State</b>   | Local routing device election state: <b>Candidate</b> , <b>Elected</b> , or <b>Ineligible</b> .      |
| <b>Timeout</b> | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

## Sample Output

### show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

| BSR                     | Pri | Local address           | Pri | State      | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None                    | 0   | 10.255.71.46            | 0   | InEligible | 0       |
| feco:1:1:1:1:0:aff:785c | 34  | feco:1:1:1:1:0:aff:7c12 | 0   | InEligible | 0       |

### show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR  | Pri | Local address   | Pri | State      | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0   | 192.168.196.105 | 0   | InEligible | 0       |



## show pim interfaces

**List of Syntax** [Syntax on page 4589](#)  
[Syntax \(EX Series Switch and the QFX Series\) on page 4589](#)

**Syntax** show pim interfaces  
 <inet | inet6>  
 <instance (*instance-name* | all)>  
 <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switch and the QFX Series)** show pim interfaces  
 <inet | inet6>  
 <instance (*instance-name* | all)>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
**inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.  
 Support for bidirectional PIM added in Junos OS Release 12.1.  
 Support for the **instance all** option added in Junos OS Release 12.1.

**Description** Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

**Options** **none**—Display interface information for all family addresses for the main instance.

**inet | inet6**—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

**instance (*instance-name* | all)**—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show pim interfaces on page 4590](#)

**Output Fields** [Table 355 on page 4589](#) describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 355: show pim interfaces Output Fields**

| Field Name      | Field Description                                                                          |
|-----------------|--------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance.                                                              |
| <b>Name</b>     | Interface name.                                                                            |
| <b>State</b>    | State of the interface. The state also is displayed in the <b>show interfaces</b> command. |

Table 355: show pim interfaces Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mode</b>         | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> <li><b>B</b>—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers.</li> <li><b>S</b>—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic.</li> <li><b>Dense</b>—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.)</li> <li><b>Sparse-Dense</b>—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as <b>dense</b> is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as <b>sparse</b> is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.)</li> </ul> <p>When sparse-dense mode is configured, the output includes both <b>S</b> and <b>D</b>. When bidirectional-sparse mode is configured, the output includes <b>S</b> and <b>B</b>. When bidirectional-sparse-dense mode is configured, the output includes <b>B</b>, <b>S</b>, and <b>D</b>.</p> |
| <b>IP</b>           | Version number of the address family on the interface: <b>4</b> (IPv4) or <b>6</b> (IPv6).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>V</b>            | PIM version running on the interface: 1 or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>State</b>        | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> <li><b>Active</b>—Bidirectional mode is enabled on the interface and on all PIM neighbors.</li> <li><b>DR</b>—Designated router.</li> <li><b>NotCap</b>—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol.</li> <li><b>NotDR</b>—Not the designated router.</li> <li><b>P2P</b>—Point to point.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>NbrCnt</b>       | Number of neighbors that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>JoinCnt(sg)</b>  | Number of (s,g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>JointCnt(*g)</b> | Number of (*g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>DR address</b>   | Address of the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name           | Stat | Mode | IP | V | State        | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0     | Up   | S    | 4  | 2 | NotDR,NotCap | 1      | 0/0            | 40.0.0.3   |
| ge-0/3/3.50    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 9901/100       | 50.0.0.2   |
| ge-0/3/3.51    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 0/0            | 51.0.0.2   |
| pe-1/2/0.32769 | Up   | S    | 4  | 2 | P2P,NotCap   | 0      | 0/0            |            |

## show pim join

---

**List of Syntax**   [Syntax on page 4592](#)  
                              [Syntax \(EX Series Switch and the QFX Series\) on page 4592](#)

**Syntax**   show pim join  
              <brief | detail | extensive | summary>  
              <inet | inet6>  
              <instance *instance-name*>  
              <logical-system (all | *logical-system-name*)>  
              <range>

**Syntax (EX Series Switch and the QFX Series)**   show pim join  
                                                          <brief | detail | extensive | summary>  
                                                          <inet | inet6>  
                                                          <instance *instance-name*>  
                                                          <range>

**Release Information**   Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              **summary** option introduced in Junos OS Release 9.6.  
                              **inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                              Support for bidirectional PIM added in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**   Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (\*G-range) for each active bidirectional RP group range, in addition to each of the joined (\*G) routes.

**Options**   **none**—Display the standard information about PIM groups for all supported family addresses for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**inet | inet6**—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**range**—(Optional) Address range of the group, specified as *prefix/prefix-length*.

**Required Privilege Level**   view

**Related Documentation**   • [clear pim join on page 4505](#)

**List of Sample Output**

- [show pim join summary on page 4596](#)
- [show pim join \(PIM Sparse Mode\) on page 4596](#)
- [show pim join \(Bidirectional PIM\) on page 4596](#)
- [show pim join instance <instance-name> on page 4597](#)
- [show pim join detail on page 4597](#)
- [show pim join extensive \(PIM Sparse Mode\) on page 4598](#)
- [show pim join extensive \(Bidirectional PIM\) on page 4599](#)
- [show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 4600](#)
- [show pim join instance <instance-name> extensive on page 4600](#)

**Output Fields** [Table 356 on page 4593](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

**Table 356: show pim join Output Fields**

| Field Name                        | Field Description                                                                                                                                      | Level of Output                     |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Instance                          | Name of the routing instance.                                                                                                                          | brief detail extensive summary none |
| Family                            | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                 | brief detail extensive summary none |
| Route type                        | Type of multicast route: (S,G) or (*,G).                                                                                                               | summary                             |
| Route count                       | Number of (S,G) routes and number of (*,G) routes.                                                                                                     | summary                             |
| R                                 | Rendezvous Point Tree.                                                                                                                                 | brief detail extensive none         |
| S                                 | Sparse.                                                                                                                                                | brief detail extensive none         |
| W                                 | Wildcard.                                                                                                                                              | brief detail extensive none         |
| Group                             | Group address.                                                                                                                                         | brief detail extensive none         |
| Bidirectional group prefix length | For bidirectional PIM, length of the IP prefix for RP group ranges.                                                                                    | All levels                          |
| Source                            | Multicast source: <ul style="list-style-type: none"> <li>• * (wildcard value)</li> <li>• <i>ipv4-address</i></li> <li>• <i>ipv6-address</i></li> </ul> | brief detail extensive none         |
| RP                                | Rendezvous point for the PIM group.                                                                                                                    | brief detail extensive none         |

Table 356: show pim join Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                    |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Flags</b>              | PIM flags: <ul style="list-style-type: none"> <li>• <b>bidirectional</b>—Bidirectional mode entry.</li> <li>• <b>dense</b>—Dense mode entry.</li> <li>• <b>rptree</b>—Entry is on the rendezvous point tree.</li> <li>• <b>sparse</b>—Sparse mode entry.</li> <li>• <b>spt</b>—Entry is on the shortest-path tree for the source.</li> <li>• <b>wildcard</b>—Entry is on the shared tree.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive none</b> |
| <b>Upstream interface</b> | RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G).<br><br>For bidirectional PIM, <b>RP Link</b> means that the interface is directly connected to a subnet that contains a phantom RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail extensive none</b> |
| <b>Upstream neighbor</b>  | Information about the upstream neighbor: <b>Direct</b> , <b>Local</b> , <b>Unknown</b> , or a specific IP address.<br><br>For bidirectional PIM, <b>Direct</b> means that the interface is directly connected to a subnet that contains a phantom RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                   |
| <b>Upstream state</b>     | Information about the upstream interface: <ul style="list-style-type: none"> <li>• <b>Join to RP</b>—Sending a join to the rendezvous point.</li> <li>• <b>Join to Source</b>—Sending a join to the source.</li> <li>• <b>Local RP</b>—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point.</li> <li>• <b>Local Source</b>—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device.</li> <li>• <b>Prune to RP</b>—Sending a prune to the rendezvous point.</li> <li>• <b>Prune to Source</b>—Sending a prune to the source.</li> </ul> <p><b>NOTE:</b> RP group range entries have <b>None</b> in the <b>Upstream state</b> field because RP group ranges do not trigger actual PIM join messages between routers.</p> | <b>extensive</b>                   |

Table 356: show pim join Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Downstream neighbors</b>               | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> <li>• <b>Interface</b>—Interface name for the downstream neighbor.</li> </ul> <p><b>NOTE:</b> A pseudo PIM-SM interface appears for all IGMP-only interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Interface address</b>—Address of the downstream neighbor.</li> <li>• <b>State</b>—Information about the downstream neighbor: <b>join</b> or <b>prune</b>.</li> <li>• <b>Flags</b>—PIM join flags: <b>R (RPtree)</b>, <b>S (Sparse)</b>, <b>W (Wildcard)</b>, or <b>zero</b>.</li> <li>• <b>Uptime</b>—Time since the downstream interface joined the group.</li> <li>• <b>Time since last Join</b>—Time since the last join message was received from the downstream interface.</li> <li>• <b>Time since last Prune</b>—Time since the last prune message was received from the downstream interface.</li> </ul> | <b>extensive</b> |
| <b>Number of downstream interfaces</b>    | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b> |
| <b>Assert Timeout</b>                     | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| <b>Keepalive timeout</b>                  | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, <b>Keepalive timeout</b> is <b>Infinity</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |
| <b>Uptime</b>                             | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| <b>Bidirectional accepting interfaces</b> | <p>Interfaces on the router that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (<b>DF Winner</b>), or the interface is the reverse path forwarding (RPF) interface toward the RP (<b>RPF</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b> |

## Sample Output

### show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type Route count
(s,g) 2
(*,g) 1

Instance: PIM.master Family: INET6
```

### show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
```



```

RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

```

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

### show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: SRW Timeout: 174
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: SRW Timeout: Infinity
 Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local Source, Local RP
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: S Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
 Interface: Pseudo-GMP
 fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
 Interface: so-1/0/0.0 (pruned)
 10.111.10.2 State: Prune Flags: SR Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3
```

Instance: PIM.master Family: INET6  
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

### show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Number of downstream interfaces: 0

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Downstream neighbors:
 Interface: lt-1/0/10.24
 10.0.24.4 State: Join RW Timeout: 185
 Interface: lt-1/0/10.23
 10.0.23.3 State: Join RW Timeout: 184
 Number of downstream interfaces: 2

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

**show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)**

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0
```

**show pim join instance <instance-name> extensive**

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
 Source: *
 RP: 10.10.47.100
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: mt-1/1/0.32768
 10.10.47.101 State: Join Flags: SRW Timeout: 156
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Number of downstream interfaces: 1

Group: 235.1.1.2
 Source: 192.168.195.74
 Flags: sparse,spt
 Upstream interface: at-0/3/1.0
 Upstream neighbor: 10.111.30.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52

Group: 235.1.1.2
 Source: 192.168.195.169
 Flags: sparse
 Upstream interface: so-1/0/1.0
 Upstream neighbor: 10.111.20.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52
```

## show pim neighbors

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4601</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4601</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                     |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (<i>instance-name</i>   all)</b>—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim neighbors on page 4603</a><br><a href="#">show pim neighbors brief on page 4603</a><br><a href="#">show pim neighbors instance on page 4603</a><br><a href="#">show pim neighbors detail on page 4603</a><br><a href="#">show pim neighbors detail (With BFD) on page 4604</a>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                                | <p><a href="#">Table 357 on page 4602</a> describes the output fields for the <b>show pim neighbors</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 357: show pim neighbors Output Fields

| Field Name                                       | Field Description                                                                                                                                                                                                                                                                                              | Level of Output   |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Instance</b>                                  | Name of the routing instance.                                                                                                                                                                                                                                                                                  | All levels        |
| <b>Interface</b>                                 | Interface through which the neighbor is reachable.                                                                                                                                                                                                                                                             | All levels        |
| <b>Neighbor addr</b>                             | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                 | All levels        |
| <b>IP</b>                                        | IP version: 4 or 6.                                                                                                                                                                                                                                                                                            | All levels        |
| <b>V</b>                                         | PIM version running on the neighbor: 1 or 2.                                                                                                                                                                                                                                                                   | All levels        |
| <b>Mode</b>                                      | PIM mode of the neighbor: <b>Sparse</b> , <b>Dense</b> , <b>SparseDense</b> , or <b>Unknown</b> . When the neighbor is running PIM version 2, this mode is always <b>Unknown</b> .                                                                                                                             | All levels        |
| <b>Option</b>                                    | Can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>B</b>—Bidirectional Capable.</li> <li>• <b>H</b>—Hello Option Holdtime.</li> <li>• <b>G</b>—Generation Identifier.</li> <li>• <b>P</b>—Hello Option DR Priority.</li> <li>• <b>L</b>—Hello Option LAN Prune Delay.</li> </ul> | <b>brief</b> none |
| <b>Uptime</b>                                    | Time the neighbor has been operational since the PIM process was last initialized, in the format <b>dd:hh:mm:ss ago</b> for less than a week and <b>nwnd:hh:mm:ss ago</b> for more than a week.                                                                                                                | All levels        |
| <b>Address</b>                                   | Address of the neighboring PIM router.                                                                                                                                                                                                                                                                         | <b>detail</b>     |
| <b>BFD</b>                                       | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: <b>Enabled</b> , <b>Operational state is up</b> , or <b>Disabled</b> .                                                                                                                                 | <b>detail</b>     |
| <b>Hello Option Holdtime</b>                     | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>Hello Default Holdtime</b>                    | Default holdtime and the time remaining if the <b>holdtime</b> option is not in the received hello message.                                                                                                                                                                                                    | <b>detail</b>     |
| <b>Hello Option DR Priority</b>                  | Designated router election priority. The range of values is 0 through 255.                                                                                                                                                                                                                                     | <b>detail</b>     |
| <b>Hello Option Generation ID</b>                | 9-digit or 10-digit number used to tag hello messages.                                                                                                                                                                                                                                                         | <b>detail</b>     |
| <b>Hello Option Bi-Directional PIM supported</b> | Neighbor can process bidirectional PIM messages.                                                                                                                                                                                                                                                               | <b>detail</b>     |
| <b>Hello Option LAN Prune Delay</b>              | Time to wait before the neighbor receives prune messages, in the format <b>delay nnn ms override nnnn ms</b> .                                                                                                                                                                                                 | <b>detail</b>     |

Table 357: show pim neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                         | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression.                                                                                                                                                                                                                                                  | detail          |
| Rx Join                    | Information about joins received from the neighbor. <ul style="list-style-type: none"> <li>• <b>Group</b>—Group addresses in the join message.</li> <li>• <b>Source</b>—Address of the source in the join message.</li> <li>• <b>Timeout</b>—Time for which the join is valid.</li> </ul> | detail          |

## Sample Output

### show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface IP V Mode Option Uptime Neighbor addr
so-1/0/0.0 4 2 HPLG 00:07:10 10.111.10.2

```

### show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 4603](#).

### show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface IP V Mode Option Uptime Neighbor addr
at-0/3/1.0 4 2 HPLG 00:07:54 10.111.30.2
mt-1/1/0.32768 4 2 HPLG 00:07:22 10.10.47.101
so-1/0/1.0 4 2 HPLG 00:07:50 10.111.20.2

```

### show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsf
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
 BFD: Disabled
 Hello Option Holdtime: 105 seconds 93 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1734018161
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1997462267
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

#### show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 836607909
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
 BFD: Enabled, Operational state is up
 Hello Default Holdtime: 105 seconds 104 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1907549685
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
 BFD: Disabled
 Hello Default Holdtime: 105 seconds 80 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1971554705
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```



## show pim rps

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4605</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4605</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>group-address</b>—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim rps on page 4608</a><br><a href="#">show pim rps brief on page 4608</a><br><a href="#">show pim rps &lt;group-address&gt; (Bidirectional PIM) on page 4608</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

[show pim rps <group-address> \(PIM Dense Mode\) on page 4608](#)  
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 4608](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 4609](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 4609](#)  
[show pim rps instance on page 4609](#)  
[show pim rps extensive \(PIM Sparse Mode\) on page 4609](#)  
[show pim rps extensive \(Bidirectional PIM\) on page 4610](#)  
[show pim rps extensive \(PIM Anycast RP in Use\) on page 4610](#)

**Output Fields** [Table 358 on page 4606](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

**Table 358: show pim rps Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Instance</b>                 | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Family or Address family</b> | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>RP address</b>               | Address of the rendezvous point.                                                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Type</b>                     | Type of RP: <ul style="list-style-type: none"> <li><b>auto-rp</b>—Address of the RP known through the Auto-RP protocol.</li> <li><b>bootstrap</b>—Address of the RP known through the bootstrap router protocol (BSR).</li> <li><b>embedded</b>—Address of the RP known through an embedded RP (IPv6).</li> <li><b>static</b>—Address of RP known through static configuration.</li> </ul> | <b>brief none</b>       |
| <b>Holdtime</b>                 | How long to keep the RP active, with time remaining, in seconds.                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Timeout</b>                  | How long until the local routing device determines the RP to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                   | All levels              |
| <b>Groups</b>                   | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Group prefixes</b>           | Addresses of groups that this RP can span.                                                                                                                                                                                                                                                                                                                                                 | <b>brief none</b>       |
| <b>Learned via</b>              | Address and method by which the RP was learned.                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Mode</b>                     | The PIM mode of the RP: bidirectional or sparse.<br><br>If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.                                                                                                                                                                                                        | All levels              |
| <b>Time Active</b>              | How long the RP has been active, in the format <b>hh:mm:ss</b> .                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |

Table 358: show pim rps Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                                     |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>Device Index</b>                   | Index value of the order in which Junos OS finds and initializes the interface.<br><br>For bidirectional RPs, the <b>Device Index</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>                             |
| <b>Subunit</b>                        | Logical unit number of the interface.<br><br>For bidirectional RPs, the <b>Subunit</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>                             |
| <b>Interface</b>                      | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.<br><br>For bidirectional RPs, the <b>Interface</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>                             |
| <b>Group Ranges</b>                   | Addresses of groups that this RP spans.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br><br><i>group-address</i> |
| <b>Active groups using RP</b>         | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>                             |
| <b>total</b>                          | Total number of active groups for this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>                             |
| <b>Register State for RP</b>          | Current register state for each group: <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:</li> <li>• <b>First Hop</b>—PIM-designated routing device that sent the Register message (the source address in the IP header).</li> <li>• <b>RP Address</b>—RP to which the Register message was sent (the destination address in the IP header).</li> <li>• <b>State</b>: <ul style="list-style-type: none"> <li>On the designated router: <ul style="list-style-type: none"> <li>• <b>Send</b>—Sending Register messages.</li> <li>• <b>Probe</b>—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.</li> <li>• <b>Suppress</b>—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to <b>Probe</b> state.</li> </ul> </li> <li>On the RP: <ul style="list-style-type: none"> <li>• <b>Receive</b>—Receiving Register messages.</li> </ul> </li> </ul> </li> </ul> | <b>extensive</b>                                    |
| <b>Anycast-PIM rpset</b>              | If anycast RP is configured, the addresses of the RPs in the set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>                                    |
| <b>Anycast-PIM local address used</b> | If anycast RP is configured, the local address used by the RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>                                    |

Table 358: show pim rps Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output      |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| <b>Anycast-PIM Register State</b> | <p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.</li> <li>• <b>Origin</b>—How the information was obtained: <ul style="list-style-type: none"> <li>• <b>DIRECT</b>—From a local attachment</li> <li>• <b>MSDP</b>—From the Multicast Source Discovery Protocol (MSDP)</li> <li>• <b>DR</b>—From the designated router</li> </ul> </li> </ul> | <b>extensive</b>     |
| <b>RP selected</b>                | For sparse mode and bidirectional mode, the identity of the RP for the specified group address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <i>group-address</i> |

## Sample Output

### show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address Type Mode Holdtime Timeout Groups Group prefixes
10.10.1.3 static bidir 150 None 2 224.1.3.0/24
 225.1.3.0/24
10.10.13.2 static bidir 150 None 2 224.1.1.0/24
 225.1.1.0/24

```

### show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 4608](#).

### show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
 11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

### show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

### show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```

user@host> show pim rps 224.1.1.1

```

Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75

RP selected: 11.4.12.75

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75 (Bidirectional)

RP selected: (null)

#### show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

| RP address   | Type   | Holdtime | Timeout | Groups | Group prefixes |
|--------------|--------|----------|---------|--------|----------------|
| 10.10.47.100 | static | 0        | None    | 1      | 224.0.0.0/4    |

Address family INET6

#### show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

| Group     | Source         | FirstHop      | RP Address    | State   | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0       |

**show pim rps extensive (Bidirectional PIM)**

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
 224.1.3.0/24
 225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
 224.1.1.0/24
 225.1.1.0/24
```

**show pim rps extensive (PIM Anycast RP in Use)**

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
 224.0.0.0/4
Active groups using RP:
 224.10.10.10

 total 1 groups active

Anycast-PIM rpset:
 10.100.111.34
 10.100.111.17
 10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:

```

| Group        | Source     | Origin |
|--------------|------------|--------|
| 224.1.1.1    | 10.10.95.2 | DIRECT |
| 224.1.1.2    | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP   |
| 224.10.10.11 | 10.10.70.1 | MSDP   |
| 224.20.20.1  | 10.10.71.1 | DR     |

```
Address family INET6

Anycast-PIM rpset:
```

```
ab::1
ab::2
Anycast-PIM local address used: cd::1
```

Anycast-PIM Register State:

| Group         | Source       | Origin |
|---------------|--------------|--------|
| ::224.1.1.1   | ::10.10.95.2 | DIRECT |
| ::224.1.1.2   | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR     |

## show pim source

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4612</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4612</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                                       | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-prefix</b>—(Optional) Display the state for source RPF states in the given range.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">show pim source on page 4613</a><br><a href="#">show pim source brief on page 4613</a><br><a href="#">show pim source detail on page 4613</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                                | <a href="#">Table 359 on page 4613</a> describes the output fields for the <b>show pim source</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



Table 359: show pim source Output Fields

| Field Name         | Field Description                                                     |
|--------------------|-----------------------------------------------------------------------|
| Instance           | Name of the routing instance.                                         |
| Source             | Address of the source or reverse path.                                |
| Prefix/length      | Prefix and prefix length for the route used to reach the RPF address. |
| Upstream interface | RPF interface toward the source address.                              |
| Upstream Neighbor  | Address of the RPF neighbor used to reach the source address.         |

## Sample Output

### show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

### show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 4613](#).

### show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local
 Active groups:228.0.0.0
 239.1.1.1
 239.1.1.1

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2
 Active groups:239.1.1.1

```

Instance: PIM.master Family: INET6

## show pim statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 4615</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 4615</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                                       | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                              |
| <b>Description</b>                                  | Display Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM statistics.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear pim statistics on page 4509</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show pim statistics on page 4622</a><br><a href="#">show pim statistics inet interface &lt;interface-name&gt; on page 4624</a><br><a href="#">show pim statistics inet6 interface &lt;interface-name&gt; on page 4624</a><br><a href="#">show pim statistics instance &lt;instance-name&gt; on page 4625</a><br><a href="#">show pim statistics interface &lt;interface-name&gt; on page 4626</a>                                                                                                                                                         |
| <b>Output Fields</b>                                | <p><a href="#">Table 360 on page 4616</a> describes the output fields for the <b>show pim statistics</b> command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                  |

Table 360: show pim statistics Output Fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>         | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet</b> interface <i>interface-name</i></li> <li>• <b>inet6</b> interface <i>interface-name</i></li> <li>• <b>interface</b> <i>interface-name</i></li> </ul>                                                                                                 |
| <b>Family</b>           | <p>Output is for IPv4 or IPv6 PIM statistics. <b>INET</b> indicates IPv4 statistics, and <b>INET6</b> indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet</b> interface <i>interface-name</i></li> <li>• <b>inet6</b> interface <i>interface-name</i></li> <li>• <b>interface</b> <i>interface-name</i></li> </ul> |
| <b>PIM statistics</b>   | PIM statistics for all interfaces or for the specified interface.                                                                                                                                                                                                                                                                                                                                                       |
| <b>PIM message type</b> | Message type for which statistics are displayed.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Received</b>         | Number of received statistics.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Sent</b>             | Number of messages sent of a certain type.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx errors</b>        | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>V2 Hello</b>         | PIM version 2 hello packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Register</b>      | PIM version 2 register packets.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>V2 Register Stop</b> | PIM version 2 register stop packets.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>V2 Join Prune</b>    | PIM version 2 join and prune packets.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>V2 Bootstrap</b>     | PIM version 2 bootstrap packets.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>V2 Assert</b>        | PIM version 2 assert packets.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>V2 Graft</b>         | PIM version 2 graft packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Graft Ack</b>     | PIM version 2 graft acknowledgment packets.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>V2 Candidate RP</b>  | PIM version 2 candidate RP packets.                                                                                                                                                                                                                                                                                                                                                                                     |

Table 360: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>V2 State Refresh</b>                 | PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.<br><br>State refresh is an extension to PIM-DM. It not supported in Junos OS. |
| <b>V2 DF Election</b>                   | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.                                                       |
| <b>V1 Query</b>                         | PIM version 1 query packets.                                                                                                                                   |
| <b>V1 Register</b>                      | PIM version 1 register packets.                                                                                                                                |
| <b>V1 Register Stop</b>                 | PIM version 1 register stop packets.                                                                                                                           |
| <b>V1 Join Prune</b>                    | PIM version 1 join and prune packets.                                                                                                                          |
| <b>V1 RP Reachability</b>               | PIM version 1 RP reachability packets.                                                                                                                         |
| <b>V1 Assert</b>                        | PIM version 1 assert packets.                                                                                                                                  |
| <b>V1 Graft</b>                         | PIM version 1 graft packets.                                                                                                                                   |
| <b>V1 Graft Ack</b>                     | PIM version 1 graft acknowledgment packets.                                                                                                                    |
| <b>AutoRP Announce</b>                  | Auto-RP announce packets.                                                                                                                                      |
| <b>AutoRP Mapping</b>                   | Auto-RP mapping packets.                                                                                                                                       |
| <b>AutoRP Unknown type</b>              | Auto-RP packets with an unknown type.                                                                                                                          |
| <b>Anycast Register</b>                 | Auto-RP announce packets.                                                                                                                                      |
| <b>Anycast Register Stop</b>            | Auto-RP announce packets.                                                                                                                                      |
| <b>Global Statistics</b>                | Summary of PIM statistics for all interfaces.                                                                                                                  |
| <b>Hello dropped on neighbor policy</b> | Number of hello packets dropped because of a configured neighbor policy.                                                                                       |
| <b>Unknown type</b>                     | Number of PIM control packets received with an unknown type.                                                                                                   |
| <b>V1 Unknown type</b>                  | Number of PIM version 1 control packets received with an unknown type.                                                                                         |
| <b>Unknown Version</b>                  | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.                                                     |

Table 360: show pim statistics Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor unknown</b>                | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.                    |
| <b>Bad Length</b>                      | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.       |
| <b>Bad Checksum</b>                    | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| <b>Bad Receive If</b>                  | Number of PIM control packets received on an interface that does not have PIM configured.                                 |
| <b>Rx Bad Data</b>                     | Number of PIM control packets received that contain data for TCP Bad register packets.                                    |
| <b>Rx Intf disabled</b>                | Number of PIM control packets received on an interface that has PIM disabled.                                             |
| <b>Rx V1 Require V2</b>                | Number of PIM version 1 control packets received on an interface configured for PIM version 2.                            |
| <b>Rx V2 Require V1</b>                | Number of PIM version 2 control packets received on an interface configured for PIM version 1.                            |
| <b>Rx Register not RP</b>              | Number of PIM register packets received when the router is not the RP for the group.                                      |
| <b>Rx Register no route</b>            | Number of PIM register packets received when the RP does not have a unicast route back to the source.                     |
| <b>Rx Register no decap if</b>         | Number of PIM register packets received when the RP does not have a de-encapsulation interface.                           |
| <b>Null Register Timeout</b>           | Number of NULL register timeout packets.                                                                                  |
| <b>RP Filtered Source</b>              | Number of PIM packets received when the router has a source address filter configured for the RP.                         |
| <b>Rx Unknown Reg Stop</b>             | Number of register stop messages received with an unknown type.                                                           |
| <b>Rx Join/Prune no state</b>          | Number of join and prune messages received for which the router has no state.                                             |
| <b>Rx Join/Prune on upstream if</b>    | Number of join and prune messages received on the interface used to reach the upstream router, toward the RP.             |
| <b>Rx Join/Prune for invalid group</b> | Number of join or prune messages received for invalid multicast group addresses.                                          |

Table 360: show pim statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                           |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rx Join/Prune messages dropped</b> | Number of join and prune messages received and dropped.                                                                                                     |
| <b>Rx sparse join for dense group</b> | Number of PIM sparse mode join messages received for a group that is configured for dense mode.                                                             |
| <b>Rx Graft/Graft Ack no state</b>    | Number of graft and graft acknowledgment messages received for which the router or switch has no state.                                                     |
| <b>Rx Graft on upstream if</b>        | Number of graft messages received on the interface used to reach the upstream router, toward the RP.                                                        |
| <b>Rx CRP not BSR</b>                 | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.                                                 |
| <b>Rx BSR when BSR</b>                | Number of BSR messages received in which the PIM message type is Bootstrap.                                                                                 |
| <b>Rx BSR not RPF if</b>              | Number of BSR messages received on an interface that is not the RPF interface.                                                                              |
| <b>Rx unknown hello opt</b>           | Number of PIM hello packets received with options that Junos OS does not support.                                                                           |
| <b>Rx data no state</b>               | Number of PIM control packets received for which the router has no state for the data type.                                                                 |
| <b>Rx RP no state</b>                 | Number of PIM control packets received for which the router has no state for the RP.                                                                        |
| <b>Rx aggregate</b>                   | Number of PIM aggregate MDT packets received.                                                                                                               |
| <b>Rx malformed packet</b>            | Number of PIM control packets received with a malformed IP unicast or multicast address family.                                                             |
| <b>No RP</b>                          | Number of PIM control packets received with no RP address.                                                                                                  |
| <b>No register encaps if</b>          | Number of PIM register packets received when the first-hop router does not have an encapsulation interface.                                                 |
| <b>No route upstream</b>              | Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, toward the RP. |
| <b>Nexthop Unusable</b>               | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.                                  |
| <b>RP mismatch</b>                    | Number of PIM control packets received for which the router has an RP mismatch.                                                                             |

Table 360: show pim statistics Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RP mode mismatch</b>                    | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>RPF neighbor unknown</b>                | Number of PIM control packets received for which the router has an unknown RPF neighbor for the source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Tx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Embedded-RP invalid addr</b>            | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Embedded-RP limit exceed</b>            | Number of times the limit configured with the <b>maximum-rps</b> statement is exceeded. The <b>maximum-rps</b> statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Embedded-RP added</b>                   | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the router:</p> <ul style="list-style-type: none"> <li>• Multicast Listener Discovery (MLD) report for an embedded RP multicast group address</li> <li>• PIM join message with an embedded RP multicast group address</li> <li>• Static embedded RP multicast group address associated with an interface</li> <li>• Packets sent to an embedded RP multicast group address received on the DR</li> </ul> <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| <b>Embedded-RP removed</b>                 | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Rx Register msgs filtering drop</b>     | Number of received register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Tx Register msgs filtering drop</b>     | Number of register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Rx Bidir Join/Prune on non-Bidir if</b> | Error counter for join and prune messages received on non-bidirectional PIM interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



Table 360: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                   |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rx Bidir Join/Prune on non-DF if</b> | Error counter for join and prune messages received on non-designated forwarder interfaces.                                                                                          |
| <b>V4 (S,G) Maximum</b>                 | Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| <b>V4 (S,G) Accepted</b>                | Number of accepted (S,G) IPv4 multicast routes.                                                                                                                                     |
| <b>V4 (S,G) Threshold</b>               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).                                            |
| <b>V4 (S,G) Log Interval</b>            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V6 (S,G) Maximum</b>                 | Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| <b>V6 (S,G) Accepted</b>                | Number of accepted (S,G) IPv6 multicast routes.                                                                                                                                     |
| <b>V6 (S,G) Threshold</b>               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).                                            |
| <b>V6 (S,G) Log Interval</b>            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V4 (grp-prefix, RP) Maximum</b>      | Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.        |
| <b>V4 (grp-prefix, RP) Accepted</b>     | Number of accepted group-to-RP IPv4 multicast mappings.                                                                                                                             |
| <b>V4 (grp-prefix, RP) Threshold</b>    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).                                    |
| <b>V4 (grp-prefix, RP) Log Interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V6 (grp-prefix, RP) Maximum</b>      | Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.                           |
| <b>V6 (grp-prefix, RP) Accepted</b>     | Number of accepted group-to-RP IPv6 multicast mappings.                                                                                                                             |

Table 360: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                 |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>V6 (grp-prefix, RP) Threshold</b>    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).                                                  |
| <b>V6 (grp-prefix, RP) Log Interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                               |
| <b>V4 Register Maximum</b>              | Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP. |
| <b>V4 Register Accepted</b>             | Number of accepted IPv4 PIM registers.                                                                                                                                                            |
| <b>V4 Register Threshold</b>            | Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).                                                                   |
| <b>V4 Register Log Interval</b>         | Time (in seconds) between consecutive log messages.                                                                                                                                               |
| <b>V6 Register Maximum</b>              | Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP. |
| <b>V6 Register Accepted</b>             | Number of accepted IPv6 PIM registers.                                                                                                                                                            |
| <b>V6 Register Threshold</b>            | Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).                                                                   |
| <b>V6 Register Log Interval</b>         | Time (in seconds) between consecutive log messages.                                                                                                                                               |

## Sample Output

### show pim statistics

```

user@host> show pim statistics
PIM Message type Received Sent Rx errors
V2 Hello 15 32 0
V2 Register 0 362 0
V2 Register Stop 483 0 0
V2 Join Prune 18 518 0
V2 Bootstrap 0 0 0
V2 Assert 0 0 0
V2 Graft 0 0 0
V2 Graft Ack 0 0 0
V2 Candidate RP 0 0 0
V2 State Refresh 0 0 0
V2 DF Election 0 0 0
V1 Query 0 0 0
V1 Register 0 0 0

```

|                       |   |   |   |
|-----------------------|---|---|---|
| V1 Register Stop      | 0 | 0 | 0 |
| V1 Join Prune         | 0 | 0 | 0 |
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

## Global Statistics

|                                  |   |
|----------------------------------|---|
| Hello dropped on neighbor policy | 0 |
| Unknown type                     | 0 |
| V1 Unknown type                  | 0 |
| Unknown Version                  | 0 |
| Neighbor unknown                 | 0 |
| Bad Length                       | 0 |
| Bad Checksum                     | 0 |
| Bad Receive If                   | 0 |
| Rx Bad Data                      | 0 |
| Rx Intf disabled                 | 0 |
| Rx V1 Require V2                 | 0 |
| Rx V2 Require V1                 | 0 |
| Rx Register not RP               | 0 |
| Rx Register no route             | 0 |
| Rx Register no decap if          | 0 |
| Null Register Timeout            | 0 |
| RP Filtered Source               | 0 |
| Rx Unknown Reg Stop              | 0 |
| Rx Join/Prune no state           | 0 |
| Rx Join/Prune on upstream if     | 0 |
| Rx Join/Prune for invalid group  | 5 |
| Rx Join/Prune messages dropped   | 0 |
| Rx sparse join for dense group   | 0 |
| Rx Graft/Graft Ack no state      | 0 |
| Rx Graft on upstream if          | 0 |
| Rx CRP not BSR                   | 0 |
| Rx BSR when BSR                  | 0 |
| Rx BSR not RPF if                | 0 |
| Rx unknown hello opt             | 0 |
| Rx data no state                 | 0 |
| Rx RP no state                   | 0 |
| Rx aggregate                     | 0 |
| Rx malformed packet              | 0 |
| Rx illegal TTL                   | 0 |
| Rx illegal destination address   | 0 |
| No RP                            | 0 |
| No register encap if             | 0 |
| No route upstream                | 0 |
| Nexthop Unusable                 | 0 |
| RP mismatch                      | 0 |
| RP mode mismatch                 | 0 |
| RPF neighbor unknown             | 0 |
| Rx Joins/Prunes filtered         | 0 |
| Tx Joins/Prunes filtered         | 0 |
| Embedded-RP invalid addr         | 0 |
| Embedded-RP limit exceed         | 0 |
| Embedded-RP added                | 0 |

```
Embedded-RP removed 0
Rx Register msgs filtering drop 0
Tx Register msgs filtering drop 0
Rx Bidir Join/Prune on non-Bidir if 0
Rx Bidir Join/Prune on non-DF if 0
```

## Sample Output

**show pim statistics inet interface <interface-name>**

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| V1 Query              | 0        | 0    | 0         |
| V1 Register           | 0        | 0    | 0         |
| V1 Register Stop      | 0        | 0    | 0         |
| V1 Join Prune         | 0        | 0    | 0         |
| V1 RP Reachability    | 0        | 0    | 0         |
| V1 Assert             | 0        | 0    | 0         |
| V1 Graft              | 0        | 0    | 0         |
| V1 Graft Ack          | 0        | 0    | 0         |
| AutoRP Announce       | 0        | 0    | 0         |
| AutoRP Mapping        | 0        | 0    | 0         |
| AutoRP Unknown type   | 0        |      |           |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## Sample Output

**show pim statistics inet6 interface <interface-name>**

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## show pim statistics instance &lt;instance-name&gt;

```

user@host> show pim statistics instance VPN-A
PIM Message type Received Sent Rx errors
V2 Hello 31 37 0
V2 Register 0 0 0
V2 Register Stop 0 0 0
V2 Join Prune 0 16 0
V2 Bootstrap 0 0 0
V2 Assert 0 0 0
V2 Graft 0 0 0
V2 Graft Ack 0 0 0
V2 Candidate RP 0 0 0
V2 State Refresh 0 0 0
V2 DF Election 0 0 0
V1 Query 0 0 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 0 0 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
AutoRP Announce 0 0 0
AutoRP Mapping 0 0 0
AutoRP Unknown type 0 0 0
Anycast Register 0 0 0
Anycast Register Stop 0 0 0

```

## Global Statistics

```

Hello dropped on neighbor policy 0
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Bad Data 0
Rx Intf disabled 0
Rx V1 Require V2 0
Rx V2 Require V1 0
Rx Register not RP 0
Rx Register no route 0
Rx Register no decap if 0
Null Register Timeout 0
RP Filtered Source 0
Rx Unknown Reg Stop 0
Rx Join/Prune no state 0
Rx Join/Prune on upstream if 0
Rx Join/Prune for invalid group 0
Rx Join/Prune messages dropped 0
Rx sparse join for dense group 0
Rx Graft/Graft Ack no state 0
Rx Graft on upstream if 0
Rx CRP not BSR 0
Rx BSR when BSR 0
Rx BSR not RPF if 0
Rx unknown hello opt 0
Rx data no state 0

```

|                                     |     |
|-------------------------------------|-----|
| Rx RP no state                      | 0   |
| Rx aggregate                        | 0   |
| Rx malformed packet                 | 0   |
| Rx illegal TTL                      | 0   |
| Rx illegal destination address      | 0   |
| No RP                               | 0   |
| No register encap if                | 0   |
| No route upstream                   | 28  |
| Nexthop Unusable                    | 0   |
| RP mismatch                         | 0   |
| RP mode mismatch                    | 0   |
| RPF neighbor unknown                | 0   |
| Rx Joins/Prunes filtered            | 0   |
| Tx Joins/Prunes filtered            | 0   |
| Embedded-RP invalid addr            | 0   |
| Embedded-RP limit exceed            | 0   |
| Embedded-RP added                   | 0   |
| Embedded-RP removed                 | 0   |
| Rx Register msgs filtering drop     | 0   |
| Tx Register msgs filtering drop     | 0   |
| Rx Bidir Join/Prune on non-Bidir if | 0   |
| Rx Bidir Join/Prune on non-DF if    | 0   |
| V4 (S,G) Maximum                    | 10  |
| V4 (S,G) Accepted                   | 9   |
| V4 (S,G) Threshold                  | 80  |
| V4 (S,G) Log Interval               | 80  |
| V6 (S,G) Maximum                    | 8   |
| V6 (S,G) Accepted                   | 8   |
| V6 (S,G) Threshold                  | 50  |
| V6 (S,G) Log Interval               | 100 |
| V4 (grp-prefix, RP) Maximum         | 100 |
| V4 (grp-prefix, RP) Accepted        | 5   |
| V4 (grp-prefix, RP) Threshold       | 80  |
| V4 (grp-prefix, RP) Log Interval    | 10  |
| V6 (grp-prefix, RP) Maximum         | 20  |
| V6 (grp-prefix, RP) Accepted        | 0   |
| V6 (grp-prefix, RP) Threshold       | 90  |
| V6 (grp-prefix, RP) Log Interval    | 20  |
| V4 Register Maximum                 | 100 |
| V4 Register Accepted                | 10  |
| V4 Register Threshold               | 80  |
| V4 Register Log Interval            | 10  |
| V6 Register Maximum                 | 20  |
| V6 Register Accepted                | 0   |
| V6 Register Threshold               | 90  |
| V6 Register Log Interval            | 20  |

## Sample Output

show pim statistics interface <interface-name>

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|------------------|----------|------|-----------|
| V2 Hello         | 0        | 3    | 0         |
| V2 Register      | 0        | 0    | 0         |
| V2 Register Stop | 0        | 0    | 0         |
| V2 Join Prune    | 0        | 0    | 0         |

|                       |   |   |   |
|-----------------------|---|---|---|
| V2 Bootstrap          | 0 | 0 | 0 |
| V2 Assert             | 0 | 0 | 0 |
| V2 Graft              | 0 | 0 | 0 |
| V2 Graft Ack          | 0 | 0 | 0 |
| V2 Candidate RP       | 0 | 0 | 0 |
| V1 Query              | 0 | 0 | 0 |
| V1 Register           | 0 | 0 | 0 |
| V1 Register Stop      | 0 | 0 | 0 |
| V1 Join Prune         | 0 | 0 | 0 |
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 3    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## show system statistics igmp

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 4628</a><br><a href="#">Syntax (EX Series Switches) on page 4628</a><br><a href="#">Syntax (TX Matrix Router) on page 4628</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 4628</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                         | show system statistics igmp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switches)</b>    | show system statistics igmp<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (TX Matrix Router)</b>      | show system statistics igmp<br><all-chassis   all-lcc   lcc <i>number</i>   scc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (TX Matrix Plus Router)</b> | show system statistics igmp<br><all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>            | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                    | Display system-wide Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                        | <b>none</b> —Display system statistics for IGMP.<br><br><b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only ) (Optional) Display system statistics for IGMP for all the routers in the chassis.<br><br><b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.<br><br><b>all-members</b> —(EX4200 switches only) (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.<br><br><b>lcc <i>number</i></b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li></ul> |



- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Display IGMP statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches only) (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**scc**—(TX Matrix routers only) (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system statistics igmp** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system statistics igmp on page 4629](#)  
[show system statistics igmp \(EX Series Switches\) on page 4629](#)  
[show system statistics igmp \(TX Matrix Plus Router\) on page 4630](#)

## Sample Output

### show system statistics igmp

```
user@host> show system statistics igmp
igmp:
 17178 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

### show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
```

```
igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid fields
 0 membership reports received
 0 membership reports received with invalid fields
 0 membership reports received for groups to which we belong
 0 Membership reports sent
```

### show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:
```

```

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

```
lcc0-re0:
```

```

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

```
lcc1-re0:
```

```

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
 0 membership reports received
 0 membership reports received with invalid field(s)
 0 membership reports received for groups to which we belong
 0 membership reports sent
```

```
lcc2-re0:
```

```

igmp:
 0 messages received
 0 messages received with too few bytes
 0 messages received with bad checksum
 0 membership queries received
 0 membership queries received with invalid field(s)
```

```
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

lcc3-re0:

-----  
igmp:

```
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

## test msdp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>test msdp (dependent-peers <i>prefix</i>   rpf-peer <i>originator</i>)</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Find Multicast Source Discovery Protocol (MSDP) peers.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>dependent-peers <i>prefix</i></b> —Find downstream dependent MSDP peers.<br><br><b>rpf-peer <i>originator</i></b> —Find the MSDP reverse-path-forwarding (RPF) peer for the originator.<br><br><b>instance <i>instance-name</i></b> —(Optional) Find MDSP peers for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">test msdp dependent-peers on page 4632</a>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```

## PART 17

# Security

- [Overview on page 4635](#)
- [Configuration on page 4697](#)
- [Administration on page 4845](#)
- [Troubleshooting on page 4873](#)



## CHAPTER 51

# Overview

- [Firewall Filters on page 4635](#)
- [Policers on page 4663](#)
- [Port Security on page 4673](#)
- [Device Security on page 4694](#)

## Firewall Filters

---

- [Overview of Firewall Filters on page 4635](#)
- [Understanding Filter-Based Forwarding on page 4637](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Understanding How Firewall Filters Control Packet Flows on page 4640](#)
- [Understanding Firewall Filter Match Conditions on page 4641](#)
- [Firewall Filter Match Conditions and Actions on page 4644](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 4656](#)
- [Understanding Firewall Filter Planning on page 4657](#)
- [Planning the Number of Firewall Filters to Create on page 4658](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 4662](#)
- [Applying Firewall Filters to Interfaces on page 4663](#)

## Overview of Firewall Filters

Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface on a QFX Series product. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). You configure firewall filters to determine whether to accept or discard a packet before it enters or exits any of these:

- Port
- VLAN

- Layer 3 (routed) interface
- Routed VLAN interface (RVI)

An *ingress* firewall filter is applied to packets that are entering an interface or VLAN, and an *egress* firewall filter is applied to packets that are exiting an interface or VLAN.



**NOTE:** Firewall filters are sometimes called *access control lists (ACLs)*.

- [Firewall Filter Types on page 4636](#)
- [Firewall Filter Components on page 4636](#)
- [Firewall Filter Processing on page 4637](#)

---

### Firewall Filter Types

The following firewall filter types are supported on the QFX Series:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 traffic transiting system ports.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter on IPv4 or IPv6 Layer 3 (routed) interfaces and routed VLAN interfaces (RVI). You can also apply a router firewall filter in the ingress direction (only) on a loopback interface, which filters traffic sent to the switch itself.



**NOTE:** You can apply a firewall filter to a management interface (for example, `me0`) on a QFX3500 device. You cannot apply a firewall filter to a management interface on a QFX3000-G or QFX3000-M system.

To apply a firewall filter:

1. Configure the firewall filter.
2. Apply the firewall filter to a port, VLAN, or router interface.



**NOTE:** You can apply only one firewall filter to a port, VLAN, or interface for a given direction. For example, for interface `ge-0/0/6.0`, you can apply one filter for the ingress direction and one for the egress direction.

---

### Firewall Filter Components

In a firewall filter, you first define the family address type (ethernet-switching or inet) and then define one or more terms that specify the filtering criteria and the action to take if a match occurs.



Each term consists of the following components:

- **Match conditions**—Specify values that a packet must contain to be considered a match. You can specify values for most fields in the IP, TCP, UDP, or ICMP headers. You can also match on interface names.
- **Action**—Specifies what to do if a packet matches the match conditions. A filter can accept, discard, or reject a matching packet and then perform additional actions, such as counting, classifying, and policing. If no action is specified for a term, the default is to accept the matching packet.

### Firewall Filter Processing

If there are multiple terms in a filter, the order of the terms is important. If a packet matches the first term, the switch executes the action defined by that term, and no other terms are evaluated. If the switch does not find a match between the packet and the first term, it compares the packet to the next term. If no match occurs between the packet and the second term, the system continues to compare the packet to each successive term in the filter until a match is found. If the packet does not match any terms in the filter, the switch discards the packet by default.

#### Related Documentation

- [Understanding Firewall Filter Planning on page 4657](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 4662](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Understanding Firewall Filter Match Conditions on page 4641](#)
- [Overview of Policers on page 4664](#)
- [Configuring Firewall Filters on page 4747](#)

### Understanding Filter-Based Forwarding

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or other security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment. For example, you might want to ensure that the highest-priority traffic is forwarded over a 40-Gigabit Ethernet link. You might also use filter-based forwarding to obtain more control over load balancing than dynamic routing protocols provide.



**NOTE:** You can create as many as 128 filters or terms that direct packets to a given virtual routing instance.

Filters used for filter-based forwarding consume memory in two ternary content addressable memories (TCAMs), and this affects the number of supported filters. See [“Planning the Number of Firewall Filters to Create” on page 4658](#) and [“Understanding FIP Snooping, FBF, and MVR Filter Scalability” on page 5038](#) for more information. The section *FBF Filter VFP TCAM Consumption* in the latter topic specifically addresses the number of supported filters when using filter-based forwarding.

**Related  
Documentation**

- [Understanding Virtual Router Routing Instances on page 2822](#)
- [Overview of Firewall Filters on page 4635](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 4697](#)

## Understanding How Firewall Filters Are Evaluated

A firewall filter consists of one or more terms, and the order of the terms within a filter is important. Before you configure firewall filters, you should understand how the QFX Series products evaluate the terms within a filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.
- If the packet does not match all the conditions, the switch discards it.

When a firewall filter consists of more than one term, the filter is evaluated sequentially:

1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

This process continues until the packet matches all the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.

4. If a packet passes through all the terms in the filter without a match, the switch discards it.




**NOTE:** The order of conditions in a **from** statement is not important because a packet must match all the conditions to be considered a match.

---

[Figure 169 on page 4639](#) shows how QFX Series products evaluate the terms within a firewall filter.

Figure 169: Evaluation of Terms Within a Firewall Filter



If you do not include a **from** statement in a term, all packets will match the term and be processed by the **then** statement. If a term does not contain a **then** statement or if an action has not been configured in the **then** statement, the term accepts any matching packets.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
 then discard;
}
```

Consequently, a packet that does not match any of the terms in a firewall filter is discarded. If you configure a filter that has no terms, all packets that pass through the filter are discarded.



**NOTE:** Firewall filtering is supported on packets that are at least 64 bytes long.

#### Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Understanding Firewall Filter Match Conditions on page 4641](#)
- [Overview of Policers on page 4664](#)
- [Configuring Firewall Filters on page 4747](#)

## Understanding How Firewall Filters Control Packet Flows

A switch supports firewall filters that allow you to control flows of data packets and local packets. *Data packets* transit a switch as they are forwarded from a source to a destination. *Local packets* are destined for or sent by a Routing Engine (they do not transit a switch). Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, or data for administrative protocols such as the Internet Control Message Protocol (ICMP).

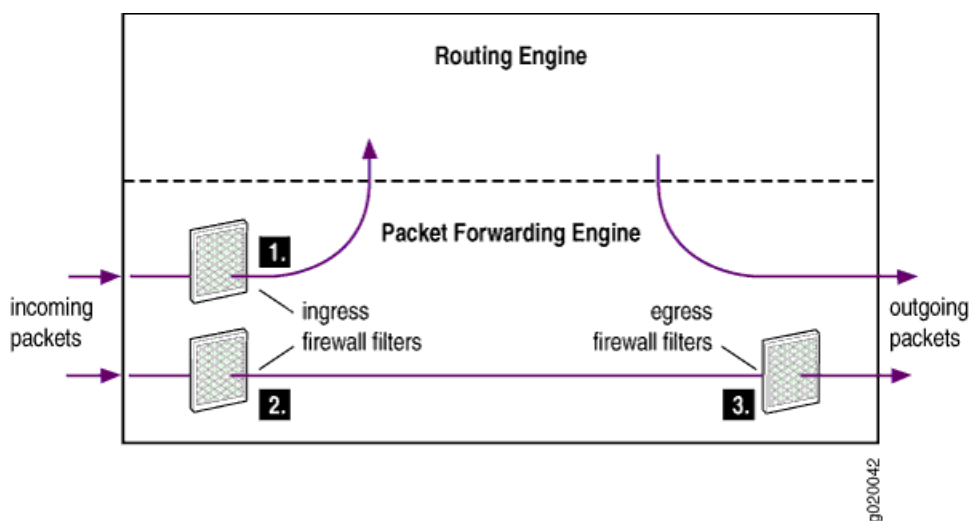
Firewall filters affect packet flows entering into or exiting from a switch as follows:

- Ingress firewall filters affect the flow of data packets that are received on switch interfaces. When a switch receives a data packet, the Packet Forwarding Engine in the system that contains the ingress interface determines where to forward the packet by looking in its Layer 2 or Layer 3 forwarding table for the best route to the destination. Data packets are forwarded to an egress interface. Locally destined packets are forwarded to the Routing Engine.
- Egress firewall filters affect data packets that are transiting a switch but do not affect packets sent by the Routing Engine. These filters are applied by the Packet Forwarding Engine in the system that contains the egress interface.

Figure 170 on page 4640 illustrates the application of ingress and egress firewall filters to control the flow of packets through a switch:

1. Ingress firewall filter applied to locally destined packets that are received on switch interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to data packets that are received on switch interfaces and will transit the switch.
3. Egress firewall filter applied to data packets that are transiting the switch.

Figure 170: Application of Firewall Filters to Control Packet Flow



## Related Documentation

- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 4662](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Configuring Firewall Filters on page 4747](#)

## Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the conditions in a term are handled and how to specify interface, numeric, address, and bit-field filter match conditions to achieve the desired filter results.

- [Filter Match Conditions on page 4641](#)
- [Numeric Filter Match Conditions on page 4641](#)
- [Interface Filter Match Conditions on page 4642](#)
- [IP Address Filter Match Conditions on page 4642](#)
- [MAC Address Filter Match Conditions on page 4643](#)
- [Bit-Field Filter Match Conditions on page 4643](#)

### Filter Match Conditions

In the **from** statement of a firewall filter term, you specify the conditions that the packet must match for the action in the **then** statement to be taken. All conditions must match for the action to be implemented. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify multiple values for the same condition, a match on any one of those values matches that condition. For example, if you specify multiple IP source addresses using the **source-address** statement, a packet that contains any one of those IP source addresses matches the condition. In some cases you can specify multiple values for the same condition by enclosing the possible values in square brackets, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set protocol (icmp | udp)
```

In other cases you must enter multiple statements, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.1.1
user@switch# set source-address 10.1.1.2
```

If you specify no match conditions in a term, that term matches all packets.



**NOTE:** Unlike traditional Junos OS firewall filters, you cannot use **except** in a condition statement to negate the condition.

### Numeric Filter Match Conditions

You can specify numeric filter match conditions that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify the

condition and a single value that a field in a packet must contain to be considered a match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 23
```

- Text synonym for a single number—A match occurs if the value of the field matches the number that corresponds to the synonym. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port telnet
```

- To specify multiple values for the same match condition in a filter term, enter each value in its own match statement. For example, a match occurs in the following term if the value of the source port in the packet is 22 or 23.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 22
user@switch# set source-port 23
```

---

### Interface Filter Match Conditions

You can specify an interface filter match condition to match an interface on which a packet is received or transmitted. For example, if you apply a filter to a VLAN you might want the filter to match on some interfaces that participate in the VLAN and not match on other interfaces in the VLAN. When you specify the name of the interface, you must include a logical unit.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/0/6.0
```

In this example, the final character (0) specifies the logical unit. You can include the wildcard (\*) as part of the interface name. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/*/6.0
user@switch# set interface ge-0/1/*.0
user@switch# set interface ge-0/0/6.*
```

Note that you must specify a value or a wildcard for the logical unit.

---

### IP Address Filter Match Conditions

You can specify an address filter match condition to match an IP source or destination address or prefix in a packet. Specify the address or prefix type and the address or prefix itself. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10.2.1.0/24;
```

If you omit the prefix length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@switch# show
destination-address {
```

```
10.0.0.0/32;
}
```

To specify more than one IP address or prefix in a filter term, enter each address or prefix in its own match statement. For example, a match occurs in the following term if the source address of a packet matches either of the following prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.0.0/16
user@switch# set source-address 10.2.0.0/16
```

### MAC Address Filter Match Conditions

You can specify a MAC address filter match condition to match a source or destination MAC address. You specify the address type and value that a packet must contain to be considered a match.

You can specify the MAC address as six hexadecimal bytes in any of the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 001122334455
```

Regardless of the formats you use, the system resolves the address to the standard format, in this case 00:11:22:33:44:55.

To specify more than one MAC address in a filter term, enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the MAC source address matches either of the following addresses:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-mac-address 00:11:22:33:44:55
user@switch# set source-mac-address 00:11:22:33:20:15
```

### Bit-Field Filter Match Conditions

You can specify bit-field filter match conditions to match particular bits within certain fields in Ethernet frames and IP, TCP, UDP, and ICMP headers. You usually specify the field and the bit within the field that must be set in a packet to be considered a match.

In most cases you can use a keyword to specify the bit you want to match on. For example, to match on a TCP SYN packet you can enter **syn**, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags syn
```

You can also enter **0x02** because the SYN bit is the third least-significant bit of the 8-bit tcp-flags field:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags 0x02
```

To match multiple bit-field values, use the logical operators, which are described in [Table 361 on page 4644](#). The operators are listed in order from highest precedence to lowest precedence. Operations are evaluated from left to right.

Table 361: Actions for Firewall Filters

| Logical Operators | Description |
|-------------------|-------------|
| !                 | Negation    |
| &                 | Logical AND |
|                   | Logical OR  |

If you use a logical operator, enclose the values in quotation marks and do not include any spaces. For example, the following statement matches the second packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&ack"
```

To negate a match, precede the value with an exclamation point. For example, the following statement matches only the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&!ack"
```

You can use text synonyms to specify some common bit-field matches. For example, the following statement also matches the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-initial
```

#### Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 4656](#)
- [Firewall Filter Match Conditions and Actions on page 4644](#)
- [Configuring Firewall Filters on page 4747](#)

## Firewall Filter Match Conditions and Actions

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the fields and values that a packet must contain to be considered a match. You can define single or multiple match conditions in *match statements*. You can also include no match statement, in which case the term matches all packets.

When a packet matches a filter, a switch takes the action specified in the term. In addition, you can specify action modifiers to count, mirror, rate-limit, and classify packets. If no match conditions are specified for the term, the switch accepts the packet by default.

This topic describes the various match conditions, actions, and action modifiers that you can define in a firewall filter.

- [Table 362 on page 4645](#) describes the match conditions you can specify when configuring a firewall filter. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. To see a list of all the synonyms for a match condition, type ? at the appropriate place in a statement.



- [Table 363 on page 4653](#) shows the actions that you can specify in a term.
- [Table 364 on page 4654](#) shows the action modifiers you can use to count, mirror, rate-limit, and classify packets.

Table 362: Supported Match Conditions for Firewall Filters

| Match Condition                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Direction and Interface                                                                                                                       |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>destination-address</b><br><i>ip-address</i>   | IP destination address field, which is the address of the final destination node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. |
| <b>destination-mac-address</b> <i>mac-address</i> | Destination media access control (MAC) address of the packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Ingress ports, VLANs and IPv4 (inet) interfaces.<br><br>Egress ports and VLANs.                                                               |
| <b>destination-port</b> <i>value</i>              | TCP or UDP destination port field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. For the following well-known ports you can specify text synonyms (the port numbers are also listed):<br><br><b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67),<br><br><b>cmd</b> (514), <b>cvspserver</b> (2401),<br><br><b>dhcp</b> (67), <b>domain</b> (53),<br><br><b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512),<br><br><b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20), | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces.                              |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Direction and Interface                                                                                                                             |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                   | <p>http (80), https (443),</p> <p>ident (113), imap (143),</p> <p>kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544),</p> <p>ldap (389), login (513),</p> <p>mobileip-agent (434), mobilip-mn (435), msdp (639),</p> <p>netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123),</p> <p>pop3 (110), pptp (1723), printer (515),</p> <p>radacct (1813), radius (1812), rip (520), rkinit (2108),</p> <p>smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514),</p> <p>tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525),</p> <p>who (513),</p> <p>xdmcp (177),</p> <p>zephyr-clt (2103), zephyr-hm (2104)</p> |                                                                                                                                                     |
| <b>destination-prefix-list</b> <i>prefix-list</i> | IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the [edit policy-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p> |
| <b>dot1q-tag</b> <i>number</i>                    | 802.1Q VLAN ID field in the Ethernet frame. The tag values can be 1–4094.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs (<i>Number</i> must be the VLAN ID of the VLAN you want to match).</p>                    |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Direction and Interface                                                                        |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| <b>dot1q-user-priority</b> <i>number</i> | <p>802.1Q priority field in the Ethernet frame (used for class-of-service priorities). Values can be 0–7.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>best-effort (0)</b>—Best effort</li> <li>• <b>background (1)</b>—Background</li> <li>• <b>standard (2)</b>—Standard or spare</li> <li>• <b>excellent-load (3)</b>—Excellent load</li> <li>• <b>controlled-load (4)</b>—Controlled load</li> <li>• <b>video (5)</b>—Video</li> <li>• <b>voice (6)</b>—Voice</li> <li>• <b>network-control (7)</b>—Network control reserved traffic</li> </ul>                                                                                                                                                                                      | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                 |
| <b>dscp</b> <i>value</i>                 | <p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>ef (46)</b>—as defined in <a href="#">RFC 3246</a>, <i>An Expedited Forwarding PHB</i>.</li> <li>• <b>af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38)</b></li> </ul> <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB</i>.</p> | <p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Direction and Interface                                        |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------|
| <b>ether-type value</b>     | <p>Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>aarp (0x80F3)</b>—EtherType value AARP</li> <li>• <b>appletalk (0x809B)</b>—EtherType value AppleTalk</li> <li>• <b>arp (0x0806)</b>—EtherType value ARP</li> <li>• <b>fcoe (0x8906)</b>—EtherType value FCoE</li> <li>• <b>fip (0x8914)</b>—EtherType value FIP</li> <li>• <b>ipv4 (0x0800)</b>—EtherType value IPv4</li> <li>• <b>ipv6 (0x08DD)</b>—EtherType value IPv6</li> <li>• <b>mpls-multicast (0x8848)</b>—EtherType value MPLS multicast</li> <li>• <b>mpls-unicast (0x8847)</b>—EtherType value MPLS unicast</li> <li>• <b>oam (0x88A8)</b>—EtherType value OAM</li> <li>• <b>ppp (0x880B)</b>—EtherType value PPP</li> <li>• <b>pppoe-discovery (0x8863)</b>—EtherType value PPPoE Discovery Stage</li> <li>• <b>pppoe-session (0x8864)</b>—EtherType value PPPoE Session Stage</li> <li>• <b>sna (0x80D5)</b>—EtherType value SNA</li> </ul> | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p> |
| <b>fragment-flags value</b> | <p>IP fragmentation flags. In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>is-fragment</b></li> <li>• <b>dont-fragment (0x4000)</b></li> <li>• <b>more-fragments (0x2000)</b></li> <li>• <b>reserved (0x8000)</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Ingress ports and VLANs.                                       |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Direction and Interface                                                                                                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <code>icmp-code value</code> | <p>ICMP code field. Because the meaning of the value depends upon the associated <code>icmp-type</code>, you must specify a value for <code>icmp-type</code> along with a value for <code>icmp-code</code>. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> <li><i>IPv4:</i><br/>parameter-problem—ip-header-bad (0), required-option-missing (1)</li> <li><i>IPv6:</i><br/>parameter-problem—ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2)</li> <li>redirect—redirect-for-network (0), redirect-for-host (1), redirect-for-tos-and-net (2), redirect-for-tos-and-host (3)</li> <li>time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)</li> <li><i>IPv4:</i><br/>unreachable—network-unreachable (0), host-unreachable (1), protocol-unreachable (2), port-unreachable (3), fragmentation-needed (4), source-route-failed (5), destination-network-unknown (6), destination-host-unknown (7), source-host-isolated (8), destination-network-prohibited (9), destination-host-prohibited (10), network-unreachable-for-TOS (11), host-unreachable-for-TOS (12), communication-prohibited-by-filtering (13), host-precedence-violation (14), precedence-cutoff-in-effect (15)</li> <li><i>IPv6:</i><br/>unreachable—address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)</li> </ul> | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Direction and Interface                                                                                                                             |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>icmp-type</b> <i>value</i>          | <p>ICMP message type field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p><i>IPv4:</i> echo-reply (0), destination unreachable (3), source-quench (4), redirect (5), echo-request (8), IPv4 (inet)-advertisement (9), IPv4 (inet)-solicit (10), time-exceeded (11), parameter-problem (12), timestamp (13), timestamp-reply (14), info-request (15), info-reply (16), mask-request (17), mask-reply (18)</p> <p><i>IPv6:</i> destination-unreachable (1), packet-too-big (2), time-exceeded (3), parameter-problem (4), echo-request (128), echo-reply (129), membership-query (130), membership-report (131), membership-termination (132), router-solicit (133), router-advertisement (134), neighbor-solicit (135), neighbor-advertisement (136), redirect (137), router-renumbering (138), node-information-request (139), node-information-reply (140)</p> <p>See also <b>icmp-code</b> <i>variable</i>.</p> | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                             |
| <b>interface</b> <i>interface-name</i> | <p>Interface on which the packet is received, including the logical unit. You can include the wildcard character (*) as part of an interface name or logical unit.</p> <p><b>NOTE:</b> An interface from which a packet is sent cannot be used as a match condition.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p> |
| <b>ip-options</b>                      | Specify <b>any</b> to create a match if anything is specified in the options field in the IP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                                      |
| <b>is-fragment</b>                     | Using this condition causes a match if the More Fragments flag is enabled in the IP header or if the fragment offset is not zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                                      |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Direction and Interface                                                                                      |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>l2-encap-type llc-non-snap</b> | Match on logical link control (LLC) layer packets for non-Subnet Access Protocol (SNAP) Ethernet Encapsulation type.                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Ingress ports and VLANs.<br>Egress ports and VLANs.                                                          |
| <b>next-header</b>                | <p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>                                                                                                                                  | <p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>             |
| <b>packet-length</b>              | Packet length in bytes. You must enter a value between 0 and 65535.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>Ingress ports, VLANs, IPv4 (inet), and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |
| <b>payload-protocol</b>           | <p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>                                                                                                                                  | <p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>             |
| <b>precedence value</b>           | <p>IP precedence bits in the type-of-service (ToS) byte in the IP header. (This byte can also be used for the DiffServ DSCP.) In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>routine (0)</b></li> <li>• <b>priority (1)</b></li> <li>• <b>immediate (2)</b></li> <li>• <b>flash (3)</b></li> <li>• <b>flash-override (4)</b></li> <li>• <b>critical-ecp (5)</b></li> <li>• <b>internet-control (6)</b></li> <li>• <b>net-control (7)</b></li> </ul> | <p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>               |

Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                     | Direction and Interface                                                                                           |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>protocol type</b>                         | IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):<br><br><b>hop-by-hop (0), icmp (1), icmp6, igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b> | Ingress ports, VLANs and IPv4 (inet) interfaces.<br><br>Egress IPv4 (inet) interfaces.                            |
| <b>source-address</b><br><b>ip-address</b>   | IP source address field, which is the address of the node that sent the packet.                                                                                                                                                                                                                                                                                                                                                 | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces..<br><br>Egress IPv4 (inet) interfaces. |
| <b>source-mac-address</b> <i>mac-address</i> | Source media access control (MAC) address of the packet.                                                                                                                                                                                                                                                                                                                                                                        | Ingress ports and VLANs.<br><br>Egress ports and VLANs.                                                           |
| <b>source-port value</b>                     | TCP or UDP source port. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. In place of the numeric field, you can specify one of the text synonyms listed under <b>destination-port</b> .                                                                                                                                                                                               | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces.  |
| <b>source-prefix-list</b> <i>prefix-list</i> | IP source prefix list. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the <b>[edit policy-options]</b> hierarchy level.                                                                                                                                                                                                                                           | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces.  |
| <b>tcp-established</b>                       | Match packets of an established TCP connection. This condition matches packets other than those used to set up a TCP connection—that is, three-way handshake packets are not matched.<br><br>When you specify <b>tcp-established</b> , a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.                                                             | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces.  |
| <b>tcp-flags value</b>                       | One or more TCP flags:<br><ul style="list-style-type: none"><li>• <b>ack (0x10)</b></li><li>• <b>fin (0x01)</b></li><li>• <b>push (0x08)</b></li><li>• <b>rst (0x04)</b></li><li>• <b>syn (0x02)</b></li><li>• <b>urgent (0x20)</b></li></ul>                                                                                                                                                                                   | Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.<br><br>Egress IPv4 (inet) interfaces.  |



Table 362: Supported Match Conditions for Firewall Filters (*continued*)

| Match Condition                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Direction and Interface                                                                                                 |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>tcp-initial</b>                 | <p>Match the first TCP packet of a connection. A match occurs when the TCP flag <b>SYN</b> is set and the TCP flag <b>ACK</b> is not set.</p> <p>When you specify <b>tcp-initial</b>, a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.</p>                                                                                                                                                                                                                                                                                                                                                    | <p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p> |
| <b>traffic-class</b>               | <p>8-bit field that specifies the class-of-service (CoS) priority of the packet. The traffic-class field is used to specify a DiffServ code point (DSCP) value. This field was previously used as the type-of-service (ToS) field in IPv4, and, the semantics of this field (for example, DSCP) are identical to those of IPv4.</p> <p>You can specify one of the following text synonyms (the field values are also listed):</p> <p><b>af11 (10), af12 (12), af13 (14), af21 (18), af22 (20), af23 (22), af31 (26), af32 (28), af33 (30), af41 (34), af42 (36), af43 (38), cs0 (0), cs1 (8), cs2 (16), cs3 (24), cs4 (32), cs5 (40), cs6 (48), cs7 (56), ef (46)</b></p> | <p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>                        |
| <b>ttl value</b>                   | IP Time-to-live (TTL) field in decimal. The value can be 1-255.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>                                            |
| <b>vlan (vlan-name   vlan-id )</b> | VLAN names or ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>                                                          |

Use **then** statements to define actions that should occur if a packet matches all conditions in a **from** statement. [Table 363 on page 4653](#) shows the actions that you can specify in a term. (If you do not include a **then** statement, the system accepts packets that match the filter.)

Table 363: Actions for Firewall Filters

| Action         | Description                                                                                    |
|----------------|------------------------------------------------------------------------------------------------|
| <b>accept</b>  | Accept a packet. This is the default action for packets that match a term.                     |
| <b>discard</b> | Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message. |

Table 363: Actions for Firewall Filters (*continued*)

| Action                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>reject</b> <i>message-type</i>            | <p>Discard a packet and send a “destination unreachable” ICMPv4 message (type 3). To log rejected packets, configure the <b>syslog</b> action modifier.</p> <p>You can specify one of the following message types: <b>administratively-prohibited</b> (default), <b>bad-host-tos</b>, <b>bad-network-tos</b>, <b>host-prohibited</b>, <b>host-unknown</b>, <b>host-unreachable</b>, <b>network-prohibited</b>, <b>network-unknown</b>, <b>network-unreachable</b>, <b>port-unreachable</b>, <b>precedence-cutoff</b>, <b>precedence-violation</b>, <b>protocol-unreachable</b>, <b>source-host-isolated</b>, <b>source-route-failed</b>, or <b>tcp-reset</b>.</p> <p>If you specify <b>tcp-reset</b>, the system sends a TCP reset if the packet is a TCP packet; otherwise nothing is sent.</p> <p>If you do not specify a message type, the ICMP notification “destination unreachable” is sent with the default message “communication administratively filtered.”</p> <p><b>NOTE:</b> The <b>reject</b> action is supported on ingress interfaces only.</p> |
| <b>routing-instance</b> <i>instance-name</i> | Forward matched packets to a virtual routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>routing-instance</b> <i>instance-name</i> | Forward matched packets to a virtual routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>vlan</b> <i>VLAN-name</i>                 | Forward matched packets to a specific VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                              | <b>NOTE:</b> The <b>vlan</b> action is supported on ingress interfaces only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

You can also specify the action modifiers listed in [Table 364 on page 4654](#) to count, mirror, rate-limit, and classify packets.

Table 364: Action Modifiers for Firewall Filters

| Action Modifier                      | Description                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>analyzer</b> <i>analyzer-name</i> | <p>Mirror traffic (copy packets) to an analyzer configured at the <b>[edit ethernet-switching-options analyzer]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>                                                                                                                        |
| <b>count</b> <i>counter-name</i>     | Count the number of packets that match the term.                                                                                                                                                                                                                                                                                                                     |
| <b>forwarding-class</b> <i>class</i> | <p>Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>assured-forwarding</b></li> <li>• <b>best-effort</b></li> <li>• <b>expedited-forwarding</b></li> <li>• <b>network-control</b></li> </ul> <p><b>NOTE:</b> The <b>forwarding-class</b> action modifier is supported on ingress interfaces only.</p> |

Table 364: Action Modifiers for Firewall Filters (*continued*)

| Action Modifier                                              | Description                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>log</b>                                                   | <p>Log the packet's header information in the Routing Engine. To view this information, enter the <b>show firewall log</b> operational mode command.</p> <p><b>NOTE:</b> The <b>log</b> action modifier is supported on ingress interfaces only.</p>                                                                                     |
| <b>loss-priority (low   medium-low   medium-high   high)</b> | <p>Set the packet loss priority (PLP).</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is supported on ingress interfaces only.</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is not supported in combination with the <b>policer</b> action.</p>                                                            |
| <b>policer <i>policer-name</i></b>                           | <p>Send packets to a policer (for the purpose of applying rate limiting).</p> <p>You can specify a policer for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p>                               |
| <b>syslog</b>                                                | <p>Log an alert for this packet.</p> <p><b>NOTE:</b> The <b>syslog</b> action modifier is supported on ingress interfaces only.</p>                                                                                                                                                                                                      |
| <b>three-color-policer <i>three-color-policer-name</i></b>   | <p>Send packets to a three-color policer (for the purpose of applying rate limiting).</p> <p>You can specify a three-color policer for ingress and egress port, VLAN, and IPv4 (inet) firewall filters.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p> |

**Related  
Documentation**

- [Understanding Firewall Filter Match Conditions on page 4641](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 4656](#)
- [Overview of Policers on page 4664](#)
- [Understanding Port Mirroring on page 4887](#)
- [Configuring Firewall Filters on page 4747](#)

## Understanding How a Firewall Filter Tests a Protocol

When examining match conditions in a firewall filter, a switch tests only the fields that you specify. It does not implicitly test any fields that you do not explicitly configure. For example, if you specify a match condition of **source-port ssh**, there is no implied test to determine if the protocol is TCP. In this case, the switch considers any packet that has a value of **22** (decimal) in the 2-byte field that follows a *presumed* IP header to be a match. To ensure that the term matches on TCP packets, you also specify a **protocol tcp** match condition.

For the following match conditions, you should explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify **protocol tcp** or **protocol udp**.
- **icmp-code**—Specify **protocol icmp** and **icmp-type**.
- **icmp-type**—Specify **protocol tcp** or **protocol udp**.
- **source-port**—Specify **protocol tcp** or **protocol udp**.
- **tcp-flags**—Specify **protocol tcp**.

### Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Understanding Firewall Filter Match Conditions on page 4641](#)
- [Configuring Firewall Filters on page 4747](#)

## Understanding Firewall Filter Planning

Before you create a firewall filter and apply it, determine what you want the filter to accomplish and how to use its match conditions and actions to achieve your goals. It is important that you understand how packets are matched, the default and configured actions of the firewall filter, and where to apply the firewall filter.

You can apply no more than one firewall filter per port, VLAN, or router interface per direction (input and output). For example, for a given port you can apply at most one filter in the input direction and one filter in the output direction. You should try to be conservative in the number of terms (rules) that you include in each firewall filter, because a large number of terms requires longer processing time during a commit operation and can make testing and troubleshooting more difficult.

Before you configure and apply firewall filters, answer the following questions for each of them:

1. What is the purpose of the filter?

For example, the system can drop packets based on header information, rate-limit traffic, classify packets into forwarding classes, log and count packets, or prevent denial-of-service attacks.

2. What are the appropriate match conditions? Determine the packet header fields that the packet must contain for a match. Possible fields include:

- Layer 2 header fields—Source and destination MAC addresses, 802.1Q tag, Ethernet type, or VLAN.
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, or TTL type).
- TCP header fields—Source and destination ports and flags.
- ICMP header fields—Packet type and code.

3. What are the appropriate actions to take if a match occurs?

The system can accept, discard, or reject packets.

4. What additional action modifiers might be required?

For example, you can configure the system to mirror (copy) packets to a specified port, count matching packets, apply traffic management, or police packets.

5. On what port, router interface, or VLAN should the firewall filter be applied?

Start with the following basic guidelines:

- If packets entering or leaving a Layer 2 interface (port) need to be filtered, apply the filter at the **[edit family ethernet switching filter]** hierarchy level. This is a port filter.
- If packets entering or leaving any port in a specific VLAN need to be filtered, use a VLAN filter.

- If packets entering or leaving a Layer 3 (routed) interface or routed VLAN interface (RVI) need to be filtered, use a router firewall filter. Apply the filter to the interface at the **[edit family inet]** hierarchy level. You can also apply a router firewall filter on a loopback interface in the ingress direction.

Before you choose the interface or VLAN on which to apply a firewall filter, understand how that placement can affect traffic flow to other interfaces. In general, apply a filter close to the source device if the filter matches on source or destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP or UDP port numbers. However, you should apply a filter close to the destination device if the filter matches *only* on a source IP address. When you apply a filter too close to the source device, the filter could prevent that source device from accessing other services that are available on the network.



**NOTE:** Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

---

6. In which direction should the firewall filter be applied?

You typically configure different actions for traffic entering an interface than you configure for traffic exiting an interface.

7. How many filters should I create?

See “[Planning the Number of Firewall Filters to Create](#)” on page 4658 for information about how many firewall filters you can apply.

**Related  
Documentation**

- [Overview of Firewall Filters on page 4635](#)
- [Overview of Policers on page 4664](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Planning the Number of Firewall Filters to Create on page 4658](#)
- [Configuring Firewall Filters on page 4747](#)

## Planning the Number of Firewall Filters to Create

- [Understanding How Many Firewall Filters Are Supported on page 4658](#)
- [Egress Filters on page 4660](#)
- [Avoid Configuring too Many Filters on page 4660](#)
- [Policers can Limit Egress Filters on page 4660](#)
- [Planning for Filter-Specific Policers on page 4661](#)
- [Planning for Filter-Based Forwarding on page 4662](#)

---

### Understanding How Many Firewall Filters Are Supported

QFX3500 and QFX3600 standalone switches and QFabric Node devices support the following maximum number of firewall filter terms per type of attachment point:

- 768 terms for ingress filters
- 1024 terms for egress filters

These totals are applied in aggregate. That is, you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction.



**NOTE:** If you want to create more than 512 egress VLAN filters, your first VLAN ID should be 6 and the subsequent VLAN IDs should increase by 1. For example, to create 1024 egress VLAN filters, the first VLAN ID would be 6, the second ID would be 7, and the sequence would continue through VLAN ID 1029. Similarly, if you want to create fewer than 512 egress VLAN filters but want the total number of terms in those filters to exceed 512, you should number your VLAN IDs in the same manner. If you do not use this approach to create your VLAN IDs, the total number of allowed terms or filters will be less than 1024 and might be 512.

The ternary content addressable memory (TCAM) for firewall filters is divided into slices that accommodate 256 terms, and all the terms in a memory slice must be in filters of the same type and applied in the same direction. A memory slice is reserved as soon as you commit a filter. For example, if you create a port filter and apply it in the input direction, a memory slice is reserved that will only store ingress port filters. If you create and apply only one ingress port filter and that filter has only one term, the rest of this slice is unused and is unavailable for other filter types.

Continuing with the above example, assume that you create and apply 256 ingress port filters with one term each so that one memory slice is filled. This leaves two more memory slices available for ingress filters. (Remember that the maximum number of ingress terms is 768.) If you then create and apply an ingress Layer 3 filter with one term, another memory slice is reserved for ingress Layer 3 filters. As before, the rest of the slice is unused and is unavailable for different filter types. At this point there is one memory slice available for any ingress filter type.

Now assume that you create and apply a VLAN ingress filter. The final memory slice is reserved for VLAN ingress filters. Memory allocation for ingress filters (once again assuming one term per filter) is as follows:

- Slice 1: Filled with 256 ingress port filters. You cannot commit any more ingress port filters.
- Slice 2: Contains one ingress Layer 3 filter with one term. You can commit 255 more terms in ingress Layer 3 filters.
- Slice 3: Contains one ingress VLAN filter with one term. You can commit 255 more terms in ingress VLAN filters.

Here is another example. Assume that you create 257 ingress port filters with one term per filter—that is, you create one more term than a single memory slice can accommodate.

When you apply the filters and commit the configuration, the filter memory allocation is:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress port filter. You can apply 255 more terms in ingress port filters.
- Slice 3: This slice is unassigned. You can create and apply 256 terms in ingress filters of any type (port, Layer 3, or VLAN), but all the filters must be of the same type.

### Egress Filters

---

All of the preceding principles also apply to egress filters, but four memory slices are used because IPv4 Layer 3 filters and IPv6 Layer 3 filters are stored in separate slices. The memory slices for egress filters are the same size as those for ingress filters, so the maximum number of egress filter terms is therefore 1024.

### Avoid Configuring too Many Filters

---

If you violate any of these restrictions and commit a configuration that is not in compliance, Junos OS rejects the excessive filters. For example, if you configure 300 ingress port filters and 300 ingress Layer 3 filters and try to commit the configuration, Junos OS does the following (again assuming one term per filter):

- Accepts the 300 ingress port filters (storing them in two memory slices).
- Accepts the first 256 ingress Layer 3 filters it processes (storing them in the third memory slice).
- Rejects the remaining 44 ingress Layer 3 filters.



**NOTE:** In this situation, be sure to delete excessive filters (for example, the remaining 44 ingress Layer 3 filters) from the configuration before you reboot the device. If you reboot a device that has a noncompliant configuration, you cannot predict which filters are installed after the reboot. Using the example above, the 44 ingress Layer 3 filters that were originally rejected might be installed, and 44 of the port filters that were originally accepted might be rejected.

---

### Policers can Limit Egress Filters

---

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress



firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem from occurring by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

### Planning for Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented above to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Planning for Filter-Based Forwarding

---

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to. Filters used in this way also consume memory in an additional TCAM. See [“Understanding FIP Snooping, FBF, and MVR Filter Scalability” on page 5038](#) for more information. The section *FBF Filter VFP TCAM Consumption* in this topic specifically addresses the number of supported filters when using filter-based forwarding.

### Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Understanding How Firewall Filters Are Evaluated on page 4638](#)
- [Understanding Firewall Filter Planning on page 4657](#)
- [Configuring Firewall Filters on page 4747](#)
- [Understanding Filter-Based Forwarding on page 4637](#)

## Understanding Firewall Filter Processing Points for Bridged and Routed Packets

You apply firewall filters at multiple processing points in the forwarding path. At each processing point, the action to be taken on a packet is determined by the configuration of the filter and the results of the lookup in the forwarding or routing table.

For both bridged (Layer 2) unicast packets and routed (Layer 3) unicast packets, firewall filters are applied in the prescribed order shown below (assuming that each filter is present and a packet is accepted by each one).

Bridged packets:

1. Ingress port filter
2. Ingress VLAN filter
3. Egress VLAN filter
4. Egress port filter

Routed packets:

1. Ingress port firewall filter
2. Ingress VLAN firewall filter (Layer 2 CoS)
3. Ingress router firewall filter (Layer 3 CoS)
4. Egress router firewall filter
5. Egress VLAN firewall filter
6. Egress port filter



**NOTE:** MAC learning occurs before filters are applied, so QFX Series products learn the MAC addresses of packets that are dropped by ingress filters.

#### Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Understanding How Firewall Filters Control Packet Flows on page 4640](#)
- [Configuring Firewall Filters on page 4747](#)

## Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family family-name filter (input |
output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface. (You cannot apply an output filter to a loopback interface.)



**NOTE:** When you create a loopback interface, it is important to apply a filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface lo0, you include the **apply-groups** statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including lo0 and other loopback interfaces.

#### Related Documentation

- [Configuring Firewall Filters on page 4747](#)

## Policers

- [Overview of Policers on page 4664](#)
- [Understanding Policers with Link Aggregation Groups on page 4669](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 4669](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 4670](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 4671](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 4672](#)

## Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 4664](#)
- [Policer Types on page 4664](#)
- [Policer Actions on page 4665](#)
- [Policer Colors on page 4666](#)
- [Filter-Specific Policers on page 4666](#)
- [Suggested Naming Convention for Policers on page 4667](#)
- [Policer Counters on page 4667](#)
- [Policer Algorithms on page 4667](#)
- [How Many Policers are Supported? on page 4668](#)
- [Policers can Limit Egress Firewall Filters on page 4668](#)

### Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 171 on page 4664](#) illustrates this process.

**Figure 171: Flow of Tricolor Marking Policer Operation**



g017049

After you name and configure a policer, you use it by specifying it as an action in one or more firewall filters.

### Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit.

You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- Single-rate three-color marker—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate three-color marker—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 365 on page 4666](#) for information about how metering results are applied for each of these policer types.

### **Policer Actions**

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 365 on page 4666](#) describes the policer actions.

Table 365: Policer Actions

| Policer                 | Marking                        | Implicit Action                  | Configurable Action |
|-------------------------|--------------------------------|----------------------------------|---------------------|
| Single-rate two-color   | Green (conforming)             | Assign low loss priority         | None                |
|                         | Red (nonconforming)            | None                             | Discard             |
| Single-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the EBS)            | Assign high loss priority        | Discard             |
| Two-rate three-color    | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the PIR and PBS)    | Assign high loss priority        | Discard             |



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is discard.

### Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

### Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth

allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 4658](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

### Suggested Naming Convention for Policers

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

### Policer Counters

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

### Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

### How Many Policers are Supported?

---

You can configure and commit the following numbers of policers on QFX3500 and QFX3600 standalone switches and QFabric Node devices:

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

### Policers can Limit Egress Firewall Filters

---

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.



You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related Documentation**

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 4669](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 4671](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 4670](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 4672](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

## Understanding Policers with Link Aggregation Groups

If you apply a policer to a link aggregation group (LAG) on a QFX3500 switch or node, the policer applies to all the interfaces in the LAG in aggregate. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer (by using a firewall filter) to a LAG that has two member interfaces on a single switch or node, the total allowed throughput for both members is 1 Gbps.

If you apply a policer to a LAG that has members on different nodes in a QFabric network Node group or redundant server Node group, the configured rate applies to the interface on each node. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer to a LAG that has one member on server node A and one member on server node B, the allowed throughput for each member is 1 Gbps, for a total allowed throughput of 2 Gbps.

**Related Documentation**

- [Overview of Policers on page 4664](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

## Understanding Color-Blind Mode for Single-Rate Tricolor Marking

With the color-blind mode of single-rate tricolor marking, all packets are evaluated against the CBS. If a packet exceeds the CBS, it is evaluated against the EBS. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

Packets that exceed the CBS but are below the EBS are marked yellow (medium-high). Packets that exceed the EBS are marked red (high), as shown in [Table 366 on page 4669](#).

**Table 366: Color-Blind Mode TCM Color-to-PLP Mapping**

| Color  | PLP         | Meaning                                                     |
|--------|-------------|-------------------------------------------------------------|
| Green  | low         | Conforming.                                                 |
| Yellow | medium-high | Packet exceeds the CIR and CBS but does not exceed the EBS. |
| Red    | high        | Packet exceeds the EBS.                                     |

- Related Documentation**
- [Overview of Policers on page 4664](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)

## Understanding Color-Aware Mode for Single-Rate Tricolor Marking

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase a preassigned PLP but cannot decrease it.

### Summary of PLP Changes

Table 367 on page 4670 shows how a packet's incoming priority can be modified with single-rate marking.

**Table 367: Color-Aware Mode Single-Rate PLP Mapping**

| Incoming PLP | Packet Metered Against      | Possible Cases                                              | Outgoing PLP |
|--------------|-----------------------------|-------------------------------------------------------------|--------------|
| low          | CIR, CBS, and EBS           | Conforming                                                  | low          |
|              |                             | Packet exceeds the CIR and CBS but does not exceed the EBS. | medium-high  |
|              |                             | Packet exceeds the EBS.                                     | high         |
| medium-low   | EBS only                    | Packet does not exceed the EBS.                             | medium-low   |
|              |                             | Packet exceeds the EBS.                                     | high         |
| medium-high  | EBS only                    | Packet does not exceed the EBS.                             | medium-high  |
|              |                             | Packet exceeds the EBS.                                     | high         |
| high         | Not metered by the policer. | All cases.                                                  | high         |

The following sections describe single-rate color-aware PLP mapping in more detail.

### ***Effect on Green Packets (Low PLP)***

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the PLP unchanged or increase it to medium-high or high, so these packets are therefore metered against both the CBS and the EBS. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as low PLP.
- If bursts exceed the CBS but not the EBS, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.
- If bursts exceed the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

***Effect on Yellow Packets (Medium PLP)***

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high, so these packets are therefore metered against the EBS only. For example, if a behavior aggregate or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-high PLP.

***Effect on Red Packets (High PLP)***

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CBS or the EBS.

**Related Documentation**

- [Overview of Policers on page 4664](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)

**Understanding Color-Blind Mode for Two-Rate Tricolor Marking**

With the color-blind mode of two-rate tricolor marking, all packets are evaluated against the committed information rate (CIR). If a packet exceeds the CIR, it is evaluated against the peak information rate (PIR). Packets that exceed the CIR but are below the PIR are marked yellow (medium-high). Packets that exceed the PIR are marked red (high).

**Table 368: Color-Blind Mode TCM Color-to-PLP Mapping**

| Color  | PLP         | Meaning                                             |
|--------|-------------|-----------------------------------------------------|
| Green  | low         | Packet does not exceed the CIR.                     |
| Yellow | medium-high | Packet exceeds the CIR but does not exceed the PIR. |

Table 368: Color-Blind Mode TCM Color-to-PLP Mapping (*continued*)

| Color | PLP  | Meaning                 |
|-------|------|-------------------------|
| Red   | high | Packet exceeds the PIR. |

**Related Documentation**

- [Overview of Policers on page 4664](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)

## Understanding Color-Aware Mode for Two-Rate Tricolor Marking

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase the preassigned PLP but cannot decrease it.

### Summary of PLP Changes

Table 369 on page 4672 shows how a packet's incoming priority can be modified with two-rate marking.

Table 369: Color-Aware Mode Two-Rate PLP Mapping

| Incoming PLP | Packet Metered Against      | Possible Cases                          | Outgoing PLP |
|--------------|-----------------------------|-----------------------------------------|--------------|
| low          | CIR and PIR                 | Packet does not exceed the CIR.         | low          |
|              |                             | Packet exceeds the CIR but not the PIR. | medium-high  |
|              |                             | Packet exceeds the PIR.                 | high         |
| medium-low   | PIR only                    | Packet does not exceed the PIR.         | medium-low   |
|              |                             | Packet exceeds the PIR.                 | high         |
| medium-high  | PIR only                    | Packet does not exceed the PIR.         | medium-high  |
|              |                             | Packet exceeds the PIR.                 | high         |
| high         | Not metered by the policer. | All cases.                              | high         |

The following sections describe color-aware two-rate PLP mapping in more detail.

### Effect on Green Packets (Low PLP)

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. These packets are therefore metered against both the CIR and the PIR. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, some of the packets are marked as medium-high PLP and some of the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as low PLP.

### Effect on Yellow Packets (Medium PLP)

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high. These packets are therefore metered against the PIR only. For example, if a behavior aggregate (BA) or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-high PLP.

### Effect on Red Packets (High PLP)

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CIR or the PIR.

#### Related Documentation

- [Overview of Policers on page 4664](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)

## Port Security

- [Overview of Access Port Protection on page 4674](#)
- [Port Security Overview on page 4676](#)

- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [Understanding DAI for Port Security on page 4686](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- [Understanding Trusted and Untrusted Ports on page 4690](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4690](#)
- [Understanding DHCP Option 82 for Port Security on page 4691](#)
- [Understanding Static ARP Entries on page 4693](#)

## Overview of Access Port Protection

Port security features can protect a switch against various types of attacks. Protection methods against some common attacks are:

- [Mitigation of Ethernet Switching Table Overflow Attacks on page 4674](#)
- [Mitigation of Rogue DHCP Server Attacks on page 4674](#)
- [Protection Against ARP Spoofing Attacks on page 4675](#)
- [Protection Against DHCP Snooping Database Alteration Attacks on page 4675](#)
- [Protection Against DHCP Starvation Attacks on page 4676](#)

### Mitigation of Ethernet Switching Table Overflow Attacks

---

In an overflow attack on an Ethernet switching table, an intruder sends so many requests from new MAC addresses that the table cannot learn all the addresses. The attack forces the switch to send broadcast messages when it needs to send traffic to addresses for which it lacks MAC addresses. In addition to generating unnecessary traffic, the attacker might be able to sniff the broadcast packets.

To mitigate such attacks, you can configure a limit for learned MAC addresses or allow only specific MAC addresses. Use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface or interfaces. By setting the MAC addresses that are explicitly allowed, you ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

### Mitigation of Rogue DHCP Server Attacks

---

By default, all access ports are untrusted, and all trunk ports are trusted with regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. If someone connects an unauthorized DHCP server to a trusted port, the unauthorized server can start issuing IP addresses and configuration information to the network's DHCP clients. The information provided to the clients by this server can disrupt their network access. The unauthorized server might also assign itself as the default gateway device for the network. An attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice.

To mitigate this problem, set the interface to which the unauthorized server is connected as untrusted. That action blocks all ingress DHCP server messages from that interface.



**NOTE:** The switch logs all DHCP server packets that are received on untrusted ports. For example:

```
5 untrusted DHCPOFFER received, interface xe-0/0/2.0[65], vlan v1[10] server
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac
12.12.12.253/00:AA:BB:CC:DD:01
```

You can use these messages to detect unauthorized DHCP servers on the network.



**NOTE:** If you attach a DHCP server to an access port, you must configure the port as trusted.

### Protection Against ARP Spoofing Attacks

In ARP spoofing, an attacker sends faked ARP messages on the network. The attacker associates its own MAC address with the IP address of a network device connected to the switch. Any traffic sent to that IP address is instead sent to the attacker. Now the attacker can create various types of problems, including sniffing the packets that were meant for another host and perpetrating man-in-the middle attacks. (In a man-in-the-middle attack, the attacker intercepts messages between two hosts, reads them, and perhaps alters them, all without the original hosts knowing that their communications have been compromised.)

To protect against ARP spoofing on your switch, enable both DHCP snooping and dynamic ARP inspection (DAI). DHCP snooping builds and maintains the DHCP snooping table. That table contains the MAC addresses, IP addresses, lease times, binding types, VLAN information, and interface information for the untrusted interfaces on the switch. DAI uses the information in the DHCP snooping table to validate ARP packets. Invalid ARP packets are blocked, and when they are blocked, a system log message is recorded that includes the type of ARP packet and the sender's IP address and MAC address.

See [“Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks”](#) on page 4732.

### Protection Against DHCP Snooping Database Alteration Attacks

In an attack designed to alter the DHCP snooping database, an intruder introduces a DHCP client on one of the switch's untrusted access interfaces that has a MAC address identical to that of a client on another untrusted port. The intruder acquires the DHCP lease, which results in changes to the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

To protect against this type of alteration of the DHCP snooping database, configure MAC addresses that are explicitly allowed on the interface. See [“Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks”](#) on page 4737.

## Protection Against DHCP Starvation Attacks

---

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses so that trusted DHCP servers cannot keep up with requests from legitimate DHCP clients. The address space of those servers is completely used up, so they can no longer assign IP addresses and lease times to clients. DHCP requests from those clients are either dropped—that is, the result is a denial of service (DoS)—or directed to a rogue DHCP server set up by the attacker to imitate a legitimate DHCP server.

To protect the switch from DHCP starvation attacks, use the MAC limiting feature. Specify the maximum number of MAC addresses that the switch can learn on the access interfaces to which DHCP clients connect. The DHCP server or servers can then supply only the specified number of IP addresses over each of those interfaces. If a DHCP starvation attack occurs after the maximum number of IP addresses has been assigned, the attack fails.

### Related Documentation

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- [Configuring MAC Limiting on page 4758](#)
- [Verifying That MAC Limiting Is Working Correctly on page 4850](#)
- [Understanding DHCP Option 82 for Port Security on page 4691](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 4715](#)
- [Understanding DAI for Port Security on page 4686](#)

## Port Security Overview

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your switch against the loss of information and productivity that can result from such attacks.

Juniper Networks Junos operating system (Junos OS) provides features to help secure ports on the switch. Ports can be categorized as either trusted or untrusted. You apply policies appropriate to each category to protect ports against various types of attacks.

Basic port security features are enabled in the switch's default configuration. You can configure additional features with minimal configuration steps.

Depending on the particular feature, you can configure the feature either on:

- VLANs—A specific VLAN or all VLANs
- Interfaces—A specific interface or all interfaces





**NOTE:** If you configure one of the port security features on all VLANs or all interfaces, those port security features are thereby enabled on all VLANs or all interfaces that are not explicitly configured with other port security features.

However, if you explicitly configure one of the port security features on a specific VLAN or on a specific interface, you must explicitly configure any additional port security features that you want to apply to that VLAN or interface. Otherwise, the switch software automatically applies the default values for the other port security features.

For example, if you enable DHCP snooping on all VLANs and enable IP source guard (not supported on the QFX Series switch) only on a specific VLAN, you must also explicitly enable DHCP snooping on that specific VLAN. If you do not explicitly enable DHCP snooping on that VLAN, the default value of no DHCP snooping applies to it.

Port security features on switches are:

- DHCP option 82—Also known as the DHCP relay agent information option. This feature helps protect the switch against attacks such as spoofing of IP addresses and media access control (MAC) addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- DHCP snooping—Filters and blocks ingress DHCP server messages on untrusted ports; builds and maintains an IP address/MAC address binding database (called the DHCP snooping database). You enable this feature on VLANs.
- Dynamic ARP inspection (DAI)—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. You enable this feature on VLANs.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is forwarded if the source IP address to source MAC address binding is valid; if the binding is not valid, the packet is discarded. You enable this feature on VLANs.



**NOTE:** IP source guard is not supported on the QFX Series.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You can enable this feature on:
  - Access interfaces (ports)
  - Access interface based on its membership within a specific VLAN



**NOTE:** You can configure a MAC limit on a VLAN in the [edit vlans] hierarchy. However, configuring the MAC limit on a VLAN does not provide protection against flooding of the Ethernet switching table. When incoming packets exceed the MAC limit on a VLAN, the event is logged, but the packets are not dropped. No other action can be configured. Therefore, setting the MAC limit on a VLAN is not considered a port security feature.

- MAC move limiting—Tracks MAC movement and detects MAC spoofing on access ports. You enable this feature on VLANs.
- Persistent MAC learning—Also known as sticky MAC. Persistent MAC learning allows dynamically learned MAC addresses to persist through switch reboots. You enable this feature on interfaces.
- Trusted DHCP server—Configuring the DHCP server on a trusted port protects against rogue DHCP servers sending leases. You enable this feature on interfaces (ports). By default, access ports are untrusted, and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect to other Ethernet switches or to routers.)

#### Related Documentation

- *Security Features for EX Series Switches Overview*
- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [Understanding DAI for Port Security on page 4686](#)
- *Understanding IP Source Guard for Port Security on EX Series Switches*
- *Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches*
- *Understanding DHCP Option 82 for Port Security on EX Series Switches*
- *Understanding Persistent MAC Learning (Sticky MAC)*
- *Understanding How to Protect Access Ports on EX Series Switches from Common Attacks*
- [Example: Configuring Basic Port Security Features on page 4706](#)

## Understanding DHCP Snooping for Port Security

DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information and build and maintain a database of valid IP address to MAC address (IP-MAC) bindings called the DHCP snooping database. Only clients with valid bindings are allowed access to the network.

- [DHCP Snooping Basics on page 4679](#)
- [DHCP Snooping Process on page 4680](#)
- [DHCP Server Access on page 4681](#)

- [DHCP Snooping Table on page 4684](#)
- [Static IP Address Additions to the DHCP Snooping Database on page 4684](#)
- [Snooping DHCP Packets That Have Invalid IP Addresses on page 4684](#)
- [Prioritizing Snooped Packets on page 4685](#)

### DHCP Snooping Basics

---

Dynamic Host Configuration Protocol (DHCP) allocates IP addresses dynamically, “leasing” addresses to devices so that the addresses can be reused when no longer needed. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port). By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping. You can modify these defaults on each of the switch's interfaces.

When DHCP snooping is enabled, the lease information from the switch (which is a DHCP client) is used to create the DHCP snooping database, a mapping of IP address to VLAN–MAC-address pairs. For each VLAN–MAC-address pair, the database stores the corresponding IP address.

Entries in the DHCP database are updated in these events:

- When a DHCP client releases an IP address (sends a DHCPRELEASE message), the associated mapping entry is deleted from the database.
- If you move a network device from one VLAN to another, typically the device has to acquire a new IP address, so its entry in the database, including the VLAN ID, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires, the associated entry is deleted from the database.



**TIP:** By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

You can configure the switch to snoop DHCP server responses only from particular VLANs. Doing this prevents spoofing of DHCP server messages.

You configure DHCP snooping per VLAN, not per interface (port). By default, DHCP snooping is disabled for all VLANs. You can enable DHCP snooping on all VLANs or on specific VLANs.



**NOTE:** If you configure DHCP for all VLANs and you enable a different port security feature on a specific VLAN, you must also explicitly enable DHCP on that VLAN. Otherwise, the default value of no DHCP snooping applies to that VLAN.



**TIP:** For private VLANs (PVLANS), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

---

### DHCP Snooping Process

---

The basic process of DHCP snooping entails the following steps:

1. Device sends DHCPDISCOVER to request IP address.
2. Switch forwards the packet to the DHCP server.
3. Server sends DHCPOFFER to offer an address. If the DHCPOFFER is from a trusted interface, switch forwards the packet to the DHCP client.
4. Device sends DHCPREQUEST to accept the IP address. Switch snoops this packet and adds IP-MAC placeholder binding to the database. The entry is considered a placeholder until a DHCPACK is received from the server. Until then, the IP address could still be assigned to some other host.
5. Server sends DHCPACK to assign the IP address or DHCPNAK to deny the address request.
6. Switch updates the DHCP database in accordance with the type of packet received:
  - Upon receipt of DHCPACK, switch updates lease information for the IP-MAC binding in its database.
  - Upon receipt of DHCPNACK, switch deletes the placeholder.



**NOTE:** DHCPDISCOVER and DHCPOFFER packets are not snooped. The DHCP database is updated only after the DHCPREQUEST packet has been sent.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the [Junos OS System Basics Configuration Guide](#).

## DHCP Server Access

Switch access to the DHCP server can be configured in three ways:


- [Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN on page 4681](#)
- [Switch Acts as DHCP Server on page 4682](#)
- [Switch Acts as Relay Agent on page 4683](#)

### ***Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN***


When the switch, DHCP clients, and DHCP server are all members of the same VLAN, the DHCP server can be connected to the switch in one of two ways:

- The server is directly connected to the same switch as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). You must configure the port that connects the server to the switch as a trusted port. See [Figure 172 on page 4681](#).
- The server is directly connected to a switch that is itself directly connected through a trunk port to the switch that the DHCP clients are connected to. The trunk port is configured by default as a trusted port. The switch that the DHCP server is connected to is not configured for DHCP snooping. See [Figure 173 on page 4682](#)—in the figure, `ge-0/0/11` is a trusted trunk port.

**Figure 172: DHCP Server Connected Directly to Switch**



**Figure 173: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port**




#### **Switch Acts as DHCP Server**



**NOTE:** This is not supported on the QFX Series switch.

The switch itself is configured as a DHCP server; this is known as a “local” configuration. See [Figure 174 on page 4683](#).

Figure 174: Switch Is the DHCP Server


**Switch Acts as Relay Agent**

The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface (on the switch, these interfaces are configured as routed VLAN interfaces, or RVIs). These trunk interfaces are trusted by default.

These two scenarios illustrate the switch acting as a relay agent:

- The DHCP server and clients are in different VLANs.
- The switch is connected to a router that is in turn connected to the DHCP server. See [Figure 175 on page 4684](#).

Figure 175: Switch Acting as Relay Agent Through Router to DHCP Server



### DHCP Snooping Table

The software creates a DHCP snooping information table that displays the content of the DHCP snooping database. The table shows current IP-MAC bindings, as well as lease time, type of binding, names of associated VLANs, and associated interface. To view the table, type **show dhcp snooping binding** at the operational mode prompt:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

| MAC address       | IP address | Lease (seconds) | Type    | VLAN     | Interface  |
|-------------------|------------|-----------------|---------|----------|------------|
| 00:05:85:3A:82:77 | 192.0.2.17 | 600             | dynamic | employee | ge-0/0/1.0 |
| 00:05:85:3A:82:79 | 192.0.2.18 | 653             | dynamic | employee | ge-0/0/1.0 |
| 00:05:85:3A:82:80 | 192.0.2.19 | 720             | dynamic | employee | ge-0/0/2.0 |

### Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add static IP addresses, you supply the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. No lease time is assigned to the entry. The statically configured entry never expires.

### Snooping DHCP Packets That Have Invalid IP Addresses

If you enable DHCP snooping on a VLAN and then devices on that VLAN send DHCP packets that request invalid IP addresses, these invalid IP addresses will be stored in the DHCP snooping database until they are deleted when their default timeout is reached. To eliminate this unnecessary consumption of space in the DHCP snooping database,



the switch drops the DHCP packets that request invalid IP addresses, preventing the snooping of these packets. The invalid IP addresses are:

- 0.0.0.0
- 128.0.x.x
- 191.255.x.x
- 192.0.0.x
- 223.255.255.x
- 224.x.x.x
- 240.x.x.x to 255.255.255.255

### Prioritizing Snooped Packets



**NOTE:** This is not supported on the QFX Series.

You can use CoS forwarding classes and queues to prioritize DHCP snooped packets for a specified VLAN. This type of configuration places the DHCP snooped packets for that VLAN in the desired egress queue, so that the security procedure does not interfere with the transmittal of high-priority traffic.

#### Related Documentation

- [Port Security Overview on page 4676](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4690](#)
- [Understanding DHCP Option 82 for Port Security on EX Series Switches](#)
- [Understanding DHCP Option 82 for Port Security on page 4691](#)
- [Understanding DHCP Services for Switches on page 63](#)
- [DHCP/BOOTP Relay for Switches Overview](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Enabling DHCP Snooping \(CLI Procedure\) on page 4764 and Enabling DHCP Snooping \(J-Web Procedure\)](#)
- [Troubleshooting Port Security](#)

## Understanding DAI for Port Security

Dynamic ARP inspection (DAI) protects switches against ARP spoofing.

DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address to entries in the database. If the MAC address or IP address in an ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are sent to the Routing Engine and are rate-limited to protect the switch from CPU overload.

- [Address Resolution Protocol on page 4686](#)
- [ARP Spoofing on page 4686](#)
- [Dynamic ARP Inspection \(DAI\) on page 4687](#)
- [Prioritizing Inspected Packets on page 4687](#)

---

### Address Resolution Protocol

Sending IP packets on a multiaccess network requires mapping an IP address to an Ethernet media access control (MAC) address.

Ethernet LANs use Address Resolution Protocol (ARP) to map MAC addresses to IP addresses.

The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

---

### ARP Spoofing

ARP spoofing (also known as ARP poisoning or ARP cache poisoning) is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switch sending traffic to the proper network device, it sends it to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that should have gone to another device. The result is that traffic from the switch is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an attacker can poison the ARP cache of a network device by sending an ARP response to

the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, the switches examine ARP responses through DAI.

### Dynamic ARP Inspection (DAI)

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid.

Junos OS for EX Series switches and the QFX Series uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, so ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs. You can set an interface to be trusted for ARP packets by setting **dhcp-trusted** on that port.

For packets directed to the switch to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Routing Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

### Prioritizing Inspected Packets



**NOTE:** This is not supported on the QFX Series.

You can use CoS forwarding classes and queues to prioritize DAI packets for a specified VLAN. This type of configuration places inspected packets for that VLAN in the desired egress queue, ensuring that the security procedure does not interfere with the transmittal of high-priority traffic.

#### Related Documentation

- [Port Security Overview on page 4676](#)
- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)

- *Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic*
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4766](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)

## Understanding MAC Limiting and MAC Move Limiting for Port Security

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

- [MAC Limiting on page 4688](#)
- [MAC Move Limiting on page 4689](#)
- [Actions for MAC Limiting on page 4689](#)
- [MAC Addresses That Exceed the MAC Limit or MAC Move Limit on page 4689](#)

### MAC Limiting

---

MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.



**NOTE:** If you do not want the system to log messages about invalid MAC addresses received by an interface that has been configured for allowed MAC addresses, disable the logging by configuring the [no-allowed-mac-log](#) statement.

You configure MAC limiting per interface, not per VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface (including tagged-access interfaces) or on all Layer 2 access interfaces.

## MAC Move Limiting

MAC move limiting causes the switch to track the number of times a MAC address can move to a new interface (port). It can help to prevent MAC spoofing, and it can also detect and prevent loops.

If a MAC address moves more than the configured number of times within 1 second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.

## Actions for MAC Limiting

You can choose to have one of the following actions performed when the limit of MAC addresses or the limit of MAC moves is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you configure the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this is not configured, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

See descriptions of results of these various action settings in “[Verifying That MAC Limiting Is Working Correctly](#)” on page 4850.

If you set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying action **none**. See “[Configuring the none Action to Override a MAC Limit Applied to All Interfaces \(CLI Procedure\)](#)” on page 4762

## MAC Addresses That Exceed the MAC Limit or MAC Move Limit

If you have configured the **port-error-disable** statement, you can view which interfaces are temporarily disabled because the MAC limit or MAC move limit was exceeded. Use the **show ethernet-switching interfaces** command.

The log messages that indicate the MAC limit or MAC move limit has been exceeded include the offending MAC addresses.

### Related Documentation

- [Port Security Overview](#) on page 4676
- [Configuring MAC Limiting](#) on page 4758
- [Configuring MAC Move Limiting \(CLI Procedure\)](#) on page 4760
- [Verifying That MAC Limiting Is Working Correctly](#) on page 4850
- [Verifying That MAC Move Limiting Is Working Correctly](#) on page 4853
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#) on page 4715
- [Example: Configuring Basic Port Security Features](#) on page 4706

- [no-allowed-mac-log on page 4821](#)

## Understanding Trusted and Untrusted Ports

By default, all access ports are untrusted and all trunk ports are trusted in regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. Untrusted ports drop traffic from DHCP servers to prevent unauthorized servers from providing any configuration information to clients.

If you attach a DHCP server to an access port, you must configure the port as trusted. Before you do so, ensure that the server is physically secure—that is, that access to the server is monitored and controlled.

### Related Documentation

- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Enabling a Trusted Port for DHCP on page 4769](#)

## Understanding Trusted DHCP Servers for Port Security

Any interface on the switch that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

### Related Documentation

- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4768](#)
- [Enabling a Trusted DHCP Server \(J-Web Procedure\)](#)

## Understanding DHCP Option 82 for Port Security

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on Ethernet LAN switches send requests for IP addresses in order to access the Internet. The switch forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to perpetrate address spoofing and starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Juniper Networks Junos operating system (Junos OS) implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Processing on page 4691](#)
- [Suboption Components of Option 82 on page 4692](#)
- [Configurations That Support Option 82 on page 4692](#)

### DHCP Option 82 Processing

If DHCP option 82 is enabled on the switch, then when a DHCP client that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “[Suboption Components of Option 82](#)” on page 4692 for details about option 82 information.

You can enable DHCP option 82 on a single VLAN or on all VLANs on the switch. You can also configure it on Layer 3 interfaces (in routed VLAN interfaces, or RVIs) when the switch is functioning as a relay agent.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards or relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.



**NOTE:** To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If it is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information in setting parameters and it does not echo the information in its response message.

---

### Suboption Components of Option 82

---

When configuring DHCP option 82, you can use the following suboptions:

- **circuit ID**—Identifies the circuit (interface and/or VLAN) on the switch on which the request was received. The circuit ID contains the interface name and/or VLAN name, with the two elements separated by a colon—for example, **xe-0/0/10:vlan1**. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, **xe-0/0/10**.

Use the **prefix** option to add an optional prefix to the circuit ID. If you enable the **prefix** option, the hostname for the switch is used as the prefix; for example, **switch1:xe-0/0/10:vlan1**.

You can also specify that the interface description be used rather than the interface name and that the VLAN ID be used rather than the VLAN name.

- **remote ID**—Identifies the host. By default, the remote ID is the MAC address of the switch. You can specify that the remote ID be the hostname of the switch, the interface description, or a character string of your choice. You can also add an optional prefix to the remote ID.
- **vendor ID**—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value **Juniper** is used. To specify a value, you type a character string.

---

### Configurations That Support Option 82


---

You can use option 82 with the following configurations:

- The DHCP client and the DHCP server are on the same VLAN. In this case the switch forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. For this configuration, you set DHCP option 82 at the **[edit ethernet-switching-options secure-access-port vlan]** hierarchy level.
- The DHCP client or the DHCP server is connected to the switch through a Layer 3 interface and the switch is configured to relay DHCP requests. [Figure 176 on page 4693](#) illustrates a scenario for the switch-as-relay-agent; in this instance, the switch relays requests through a router to the server.



Figure 176: Switch Relays DHCP Requests to Server



For the configuration shown in [Figure 176 on page 4693](#), you set DHCP option 82 at the `[edit forwarding-options helpers bootp]` hierarchy level.

#### Related Documentation

- [Overview of Access Port Protection on page 4674](#)
- [DHCP and BOOTP Relay Overview on page 4892](#)
- [dhcp-option82 on page 4807](#)
- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740](#)
- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4770](#)
- [Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4773](#)

## Understanding Static ARP Entries

You can create explicit mappings between IP addresses and MAC addresses, which are called static ARP table entries. Unlike dynamically learned ARP entries, static entries do not age out. You might want to create static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically for any reason.

#### Related Documentation

- [Configuring Static ARP Entries on page 2045](#)
- [arp on page 2072](#)

## Device Security

---

- [Understanding Storm Control on page 4694](#)

### Understanding Storm Control

A traffic storm occurs when broadcast packets prompt receiving devices to broadcast packets in response. This prompts further responses, creating a snowball effect. The switch is flooded with packets, which creates unnecessary traffic that leads to poor performance or even a complete loss of service by some clients. Storm control causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

By default, storm control is enabled for ingress traffic on all interfaces at a level of 80 percent of the available bandwidth. You can change the storm control level by configuring it as a specific bandwidth value. (The **level** configuration statement, which allows you to configure the storm control level as a percentage of the combined broadcast and unknown unicast streams, is deprecated and might be removed from future releases. We recommend that you phase out its use and replace it with the **bandwidth** statement.)



**NOTE:** When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually enforced. For example, if you configure a bandwidth limit of 150 Kbps, storm control enforces a bandwidth limit of 128 Kbps.



**NOTE:** Storm control must be disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped. Configuring storm control on an Ethernet interface and including that interface in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. After disabling storm control on all interfaces, enable storm control on any interfaces that are not part of an FCoE-FC gateway on which you want to use storm control.



**CAUTION:** The Junos OS allows you to configure a storm control value that exceeds the bandwidth of the interface. If you configure an interface this way, storm control does not drop broadcast or unknown unicast packets even if they consume all the available bandwidth.

Broadcast, multicast, and unicast packets are sent as part of normal operations, so to recognize a storm, you must be able to identify when traffic has reached an abnormal

level. Suspect a storm when operations begin timing out and network response times slow down. Users might be unable to access expected services. Monitor the percentage of broadcast and unknown unicast traffic in the network when it is operating normally. This data can then be used as a benchmark to determine when traffic levels are too high. You can then configure storm control to set the level at which you want to drop broadcast and unknown unicast traffic.

**Related  
Documentation**

- [Example: Configuring Storm Control on page 4713](#)
- [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\) on page 4762](#)
- [Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway on page 5189](#)
- [action-shutdown on page 4836](#)
- [interface \(Storm Control\) on page 2095](#)
- [port-error-disable on page 4823](#)
- [storm-control on page 2132](#)



## CHAPTER 52

# Configuration

- [Firewall and Policer Configuration Examples on page 4697](#)
- [Port Security Configuration Examples on page 4706](#)
- [Firewall and Policer Configuration Tasks on page 4746](#)
- [Port Security Configuration Tasks on page 4755](#)
- [Configuration Statements for Firewall Filters on page 4775](#)
- [Configuration Statements for Policers on page 4783](#)
- [Configuration Statements for Port Security on page 4801](#)
- [Configuration Statements for Device Security on page 4835](#)

### Firewall and Policer Configuration Examples

---

- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 4697](#)
- [Example: Using Two-Color Policers and Prefix Lists on page 4701](#)
- [Example: Using Policers to Manage Oversubscription on page 4704](#)

### Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device

You can configure filter-based forwarding by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

- [Requirements on page 4697](#)
- [Overview and Topology on page 4698](#)
- [Configuration on page 4698](#)
- [Verification on page 4700](#)

#### Requirements

---

This example requires Junos OS Release 12.2X50-D20 or later.

## Overview and Topology

---

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address of the source application server. Any matching packets are routed to a virtual routing instance that sends the traffic to a security device. In this case, the security device must be able to forward the traffic to the destination application server. For this example, assume that the address of the destination application server is 192.168.0.1.

## Configuration

---

To configure filter-based forwarding:

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste them into the CLI at the **[edit]** hierarchy level.

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces xe-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter f1 term t1 from source-address 10.1.0.50/32
set firewall family inet filter f1 term t1 from protocol tcp
set interfaces xe-0/0/0 unit 0 family inet filter input f1
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface xe-0/0/3.0
set routing-instances vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
set firewall family inet filter f1 term t1 then routing-instance vrf01
```

### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To configure filter-based forwarding:

1. Configure an interface to connect to the application server:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet address 10.1.0.1/24
```
2. Configure an interface to connect to the security device:  

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 10.1.3.1/24
```
3. Create a firewall filter that matches packets based on the address of the application server that the traffic will be sent from. Also configure the filter so that it matches only TCP packets:  

```
[edit firewall]
user@switch# set family inet filter f1 term t1 from source-address 10.1.0.50/32
user@switch# set firewall family inet filter f1 term t1 from protocol tcp
```
4. Apply the filter to the interface that connects to the source application server and configure it to match incoming packets:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet filter input f1
```
5. Create a virtual router:  

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```
6. Associate the virtual router with the interface that connects to the security device:

- ```
[edit routing-instances]
user@switch# set vrf01 interface xe-0/0/3.0
```
7. Configure the routing information for the virtual routing instance:

```
[edit routing-instances]
user@switch# set vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
```
 8. Set the filter to forward packets to the virtual router:

```
[edit firewall]
user@switch# set family inet filter f1 term t1 then routing-instance vrf01
```

Results

Check the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input f1;
        }
        address 10.1.0.1/24;
      }
    }
  }
  xe-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.3.1/24;
      }
    }
  }
}
firewall {
  family inet {
    filter f1 {
      term t1 {
        from {
          source-address {
            10.1.0.50/32;
          }
          protocol tcp;
        }
        then {
          routing-instance vrf01;
        }
      }
    }
  }
}
routing-instances {
  vrf01 {
    instance-type virtual-router;
    interface xe-0/0/1.0;
    routing-options {
```

```

static {
    route 12.34.56.0/24 next-hop 10.1.3.254;
}
}
}
}

```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Filter-Based Forwarding Was Configured on page 4700](#)

Verifying That Filter-Based Forwarding Was Configured

Purpose Verify that filter-based forwarding was properly enabled on the switch.

Action 1. Use the `show interfaces filters` command:

```

user@switch> show interfaces filters xe-0/0/0.0
Interface      Admin Link Proto Input Filter      Output Filter
xe-0/0/0.0      up    down inet f1

```

2. Use the `show route forwarding-table` command:

```

user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          user  1 0:12:f2:21:cf:0 ucst  331  4 me0.0
default          perm  0                rjct  36   3
0.0.0.0/32       perm  0                dscd  34   1
10.1.0.0/24      ifdn  0                rslv  613  1 xe-0/0/0.0
10.1.0.0/32      iddn  0 10.1.0.0        recv  611  1 xe-0/0/0.0
10.1.0.1/32      user  0                rjct  36   3
10.1.0.1/32      intf  0 10.1.0.1        locl  612  2
10.1.0.1/32      iddn  0 10.1.0.1        locl  612  2
10.1.0.255/32    iddn  0 10.1.0.255      bcst  610  1 xe-0/0/0.0
10.1.1.0/26      ifdn  0                rslv  583  1 vlan.0
10.1.1.0/32      iddn  0 10.1.1.0        recv  581  1 vlan.0
10.1.1.1/32      user  0                rjct  36   3
10.1.1.1/32      intf  0 10.1.1.1        locl  582  2
10.1.1.1/32      iddn  0 10.1.1.1        locl  582  2
10.1.1.63/32     iddn  0 10.1.1.63       bcst  580  1 vlan.0
255.255.255.255/32 perm  0                bcst  32   1

```

```

Routing table: vrf01.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  559  2
0.0.0.0/32       perm  0                dscd  545  1
10.1.3.0/24      ifdn  0                rslv  617  1 xe-0/0/3.0
10.1.3.0/32      iddn  0 10.1.3.0        recv  615  1 xe-0/0/3.0
10.1.3.1/32      user  0                rjct  559  2
192.168.0.1/24   user  0 10.1.3.254      ucst  616  2 xe-0/0/3.0
192.168.0.1/24   user  0 10.1.3.254      ucst  616  2 xe-0/0/3.0
10.1.3.255/32    iddn  0 10.1.3.255      bcst  614  1 xe-0/0/3.0
224.0.0.0/4      perm  0                mdsc  546  1
224.0.0.1/32     perm  0 224.0.0.1       mcst  529  1

```



```

255.255.255.255/32 perm      0                bcst    543      1

Routing table: default.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct    60       1

Routing table: vrf01.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct    600      1

```

Meaning The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

- Related Documentation**
- [Configuring Firewall Filters on page 4747](#)
 - [Understanding Filter-Based Forwarding on page 4637](#)
 - [Understanding Virtual Router Routing Instances on page 2822](#)

Example: Using Two-Color Policers and Prefix Lists

If you provide specific amounts of bandwidth to internal or external customers, you can use policing to make sure that customers do not consume more bandwidth than they should receive. For example, you might connect many customers to one 10-Gbps interface and want to ensure that none of them congest the interface by using more bandwidth than they have been allotted.

You could accomplish this by creating a two-color policer similar to the following for each customer:

```

firewall {
  policer Limit-Customer-1 {
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 150m;
    }
    then discard;
  }
}

```

Creating a policer for each customer is clearly not a scalable solution, however. As an alternative, you can create prefix lists that group classes of customers and then create policers for each prefix list. For example, you could create prefix lists such as **Class-A-Customer-Prefixes**, **Class-B-Customer-Prefixes**, and **Class-C-Customer-Prefixes** (at the **[edit policy-options]** hierarchy level) and create the following corresponding policers:

```

firewall {
  policer Class-A {
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 150m;
    }
  }
}

```

```
    }
    then discard;
  }
  policer Class-B {
    if-exceeding {
      bandwidth-limit 75m;
      burst-size-limit 100m;
    }
    then discard;
  }
  policer Class-C {
    if-exceeding {
      bandwidth-limit 50m;
      burst-size-limit 75m;
    }
    then discard;
  }
}
```

You must create filter terms that specify the prefix lists in their **from** statements and the corresponding policers in their **then** statements similar to the following:

```
firewall
family inet {
  filter Class-A-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-A-Customer-Prefixes;
        }
      }
      then policer Class-A;
    }
  }
  filter Class-B-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-B-Customer-Prefixes;
        }
      }
      then policer Class-B;
    }
  }
  filter Class-C-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-C-Customer-Prefixes;
        }
      }
      then policer Class-C;
    }
  }
}
```

Here are the steps to create this firewall configuration:

1. Create the first policer:

```
[edit firewall]
user@switch# set policer Class-A if-exceeding bandwidth-limit 100m burst-size-limit 150m
user@switch# set policer Class-A then discard
```

2. Create the second policer:

```
[edit firewall]
user@switch# set policer Class-B if-exceeding bandwidth-limit 75m burst-size-limit 100m
user@switch# set policer Class-B then discard
```

3. Create the third policer:

```
[edit firewall]
user@switch# set policer Class-C if-exceeding bandwidth-limit 50m burst-size-limit 75m
user@switch# set policer Class-C then discard
```

4. Create a filter for class A customers:

```
[edit firewall]
user@switch# edit family inet filter Class-A-Customers
```

5. Configure the filter to send packets matching the Class-A-Customer-Prefixes prefix list to the Class-A policer:

```
[edit firewall family inet filter Class-A-Customers]
user@switch# set term term-1 from source-prefix-list Class-A-Customers
user@switch# set term term-1 then policer Class-A
```

6. Create a filter for class B customers:

```
[edit firewall]
user@switch# edit family inet filter Class-B-Customers
```

7. Configure the filter to send packets matching the Class-B-Customer-Prefixes prefix list to the Class-B policer:

```
[edit firewall family inet filter Class-B-Customers]
user@switch# set term term-1 from source-prefix-list Class-B-Customers
user@switch# set term term-1 then policer Class-B
```

8. Create a filter for class C customers:

```
[edit firewall]
user@switch# edit family inet filter Class-C-Customers
```

9. Configure the filter to send packets matching the Class-C-Customer-Prefixes prefix list to the Class-C policer:

```
[edit firewall family inet filter Class-C-Customers]
user@switch# set term term-1 from source-prefix-list Class-C-Customers
user@switch# set term term-1 then policer Class-C
```

10. Apply the filters you created to the appropriate interfaces in the output direction.



NOTE: Note that the implicit deny statement in this filter will block traffic from any source that does not match one of the prefix lists. If you want the filter to allow this traffic, you must include an explicit term for this purpose.

Related Documentation

- [Overview of Policers on page 4664](#)
- [Applying Firewall Filters to Interfaces on page 4663](#)

- *prefix-list*

Example: Using Policers to Manage Oversubscription

You might want to use a policer when an interface is oversubscribed and you want to control what will happen if congestion occurs. For example, you might have servers connected to a switch as listed in [Table 370 on page 4704](#).

Table 370: Servers Connected to Switch

Server Type	Connection	IP Address
Network application server	1-gigabit interface	10.0.0.1
Authentication server	1-gigabit interface	10.0.0.2
Database server	10-gigabit interface	10.0.0.3

In this example, users access services provided by the network application server, which requests information from the database server as appropriate. When it receives a request from a user, the network application server first contacts the authentication server to verify the user's credentials. When a user is authenticated and the network application server provides the requested service, all the packets sent from the database server to the application server must transit the 1-Gigabit Ethernet interface connected to the application server twice—once on ingress to the application server and again on egress to the user.

The sequence of events for a user session is as follows:

1. A user connects to the application server and requests a service.
2. The application server requests the user's credentials and relays them to the authentication server.
3. If the authentication server verifies the credentials, the application server initiates the requested service.
4. The application server requests the files necessary to meet the user's request from the database server.
5. The database server sends the requested files to the application server.
6. The application server includes the requested files in its response to the user.

Traffic from the database server to the application server might congest the 1-gigabit interface to which that the application server is connected. This congestion might prevent the server from responding to requests from users and creating new sessions for them. You can use policing to make sure that this does not occur.

To create this firewall configuration, perform the following steps on the database server:

1. Create a policer to drop traffic from the database server to the application server if it exceeds certain limits:

```
[edit firewall]
user@switch# set policer Database-Egress-Policer if-exceeding bandwidth-limit 400
burst-size-limit 500m
user@switch# set policer Database-Egress-Policer then discard
```

2. Create a filter to examine traffic from the database server to the application server:

```
[edit firewall]
user@switch# edit family inet filter Database-Egress-Filter
```

3. Configure the filter to apply the policer to traffic egressing the database server and destined for the application server:

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-1 from destination-address 10.0.0.1
user@switch# set term term-1 then policer Database-Egress-Policer
```

4. If required, configure a term to allow traffic from the database server to other destinations (otherwise the traffic will be dropped by the implicit deny statement):

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-2 then accept
```

Note that omitting a **from** statement causes the term to match all packets, which is the desired behavior.

5. Install the egress filter as an output filter on the database server interface that is connected the application server:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet filter output Database-Egress-Filter
```

Here is how the final configuration would appear:

```
firewall {
  policer Database-Egress-Policer {
    if-exceeding {
      bandwidth-limit 400;
      burst-size-limit 500m;
    }
    then discard;
  }
  family inet {
    filter Database-Egress-Filter {
      term term-1 {
        from {
          destination-address {
            10.0.0.1/24;
          }
        }
        then policer Database-Egress-Policer;
      }
      term term-2 { # If required, include this term so that traffic from the database server
                    # to other destinations is allowed.
        then accept;
      }
    }
  }
}
```

```
}  
]
```

**Related
Documentation**

- [Overview of Policers on page 4664](#)

Port Security Configuration Examples

- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring Storm Control on page 4713](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 4715](#)
- [Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 4718](#)
- [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
- [Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 4737](#)
- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740](#)
- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744](#)

Example: Configuring Basic Port Security Features

You can configure DHCP snooping, dynamic ARP inspection (DAI), MAC limiting, persistent MAC learning, and MAC move limiting on the access ports of switches to protect the switches and the Ethernet LAN against address spoofing and Layer 2 denial-of-service (DoS) attacks. You can also configure a trusted DHCP server and specific (allowed) MAC addresses for the switch interfaces.

This example describes how to configure basic port security features on a switch:

- [Requirements on page 4706](#)
- [Overview and Topology on page 4707](#)
- [Configuration on page 4709](#)
- [Verification on page 4710](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or one QFX3500 switch

- Junos OS Release 11.4 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure basic port security features, be sure you have:

- Connected the DHCP server to the switch.
- Configured a VLAN on the switch. See the task for your platform:
 - *Configuring VLANs for EX Series Switches (CLI Procedure)*
 - “Configuring VLANs” on page 2048 for the QFX Series



NOTE: In this example, the DHCP server and its clients are all members of a single VLAN on the switch.

Overview and Topology


Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. To protect the devices from such attacks, you can configure:

- DHCP snooping to validate DHCP server messages
- DAI to protect against MAC spoofing
- MAC limiting to constrain the number of MAC addresses the switch adds to its MAC address cache
- MAC move limiting to help prevent MAC spoofing
- Persistent MAC learning (sticky MAC) to constrain the MAC addresses that can be learned on an interface to the first ones learned, even after a reboot of the switch
- Trusted DHCP server configured on a trusted port to protect against rogue DHCP servers sending leases

This example shows how to configure these security features on a switch connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. [Figure 177 on page 4708](#) illustrates the topology for this example.

Figure 177: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 371 on page 4708](#).

Table 371: Components of the Port Security Topology

Properties	Settings
Switch hardware	One EX Series switch or one QFX3500 switch
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is subnet's broadcast address
Interfaces in employee-vlan	ge-0/0/1 , ge-0/0/2 , ge-0/0/3 , ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, the switch is initially configured with the default port security setup. In the default switch configuration:

- Secure port access is activated on the switch.
- DHCP snooping and DAI are disabled on all VLANs.
- All access ports are untrusted, and all trunk ports are trusted for DHCP snooping.

In the configuration tasks for this example, you set the DHCP server as trusted; you enable DHCP snooping, DAI, and MAC move limiting on a VLAN; you set a value for a MAC limit on some interfaces; you configure some specific (allowed) MAC addresses on an interface; and you configure persistent MAC learning on an interface.

Configuration

To configure basic port security on a switch whose DHCP server and client ports are in a single VLAN:

CLI Quick Configuration

To quickly configure basic port security on the switch, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/1 mac-limit 4
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
set interface ge-0/0/2 mac-limit 4
set interface ge-0/0/1 persistent-learning
set interface ge-0/0/8 dhcp-trusted
set vlan employee-vlan arp-inspection
set vlan employee-vlan examine-dhcp
set vlan employee-vlan mac-move-limit 5
```

Step-by-Step Procedure

Configure basic port security on the switch:

1. Enable DHCP snooping on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan examine-dhcp
```

2. Specify the interface (port) from which DHCP responses are allowed:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/8 dhcp-trusted
```

3. Enable dynamic ARP inspection (DAI) on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan arp-inspection
```

4. Configure a MAC limit of 4 and use the default action, **drop**. (Packets are dropped, and the MAC address is not added to the Ethernet switching table if the MAC limit is exceeded on the interfaces):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 mac-limit 4
user@switch# set interface ge-0/0/2 mac-limit 4
```

5. Allow learned MAC addresses for a particular interface to persist across restarts of the switch and interface-down events by enabling persistent MAC learning:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 persistent-learning
```

6. Configure a MAC move limit of 5 and use the default action, **drop**. (Packets are dropped, and the MAC address is not added to the Ethernet switching table if a MAC address has exceeded the MAC move limit):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan mac-move-limit 5
```

7. Configure allowed MAC addresses:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
```

```
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
    mac-limit 4;
    persistent-learning;
}
interface ge-0/0/2.0 {
    allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83
    00:05:85:3a:82:85 00:05:85:3a:82:88 ];
    mac-limit 4;
}
interface ge-0/0/8.0 {
    dhcp-trusted;
}
vlan employee-vlan {
    arp-inspection
    examine-dhcp;
    mac-move-limit 5;
}
```

Verification

To confirm that the configuration is working properly:

- [Verifying That DHCP Snooping Is Working Correctly on the Switch on page 4710](#)
- [Verifying That DAI Is Working Correctly on the Switch on page 4711](#)
- [Verifying That MAC Limiting, MAC Move Limiting, and Persistent MAC Learning Are Working Correctly on the Switch on page 4712](#)
- [Verifying That Allowed MAC Addresses Are Working Correctly on the Switch on page 4712](#)

Verifying That DHCP Snooping Is Working Correctly on the Switch

Purpose Verify that DHCP snooping is working on the switch.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
DHCP Snooping Information:
MAC Address      IP Address      Lease      Type      VLAN      Interface
-----
00:05:85:3A:82:77 192.0.2.17      600      dynamic  employee-vlan ge-0/0/1.0
00:05:85:3A:82:79 192.0.2.18      653      dynamic  employee-vlan ge-0/0/1.0
00:05:85:3A:82:80 192.0.2.19      720      dynamic  employee-vlan ge-0/0/2.0
00:05:85:3A:82:81 192.0.2.20      932      dynamic  employee-vlan ge-0/0/2.0
00:05:85:3A:82:83 192.0.2.21      1230     dynamic  employee-vlan ge-0/0/2.0
00:05:85:27:32:88 192.0.2.22      3200     dynamic  employee-vlan ge-0/0/2.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database, and nothing would be shown in the output of the **show dhcp snooping binding** command.

Verifying That DAI Is Working Correctly on the Switch

Purpose Verify that DAI is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface      Packets received  ARP inspection pass  ARP inspection failed
-----
ge-0/0/1.0      7                 5                    2
ge-0/0/2.0     10                10                   0
ge-0/0/3.0     12                12                   0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Verifying That MAC Limiting, MAC Move Limiting, and Persistent MAC Learning Are Working Correctly on the Switch

Purpose Verify that MAC limiting, MAC move limiting, and persistent MAC learning are working on the switch.

Action Suppose that two packets have been sent from hosts on **ge-0/0/1** and five packets from hosts on **ge-0/0/2**, with both interfaces set to a MAC limit of 4 with the default action **drop** and **ge-0/0/1** enabled for persistent MAC learning.

Display the MAC addresses learned:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 4 learned, 2 persistent entries
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	*	Flood	-	All-members
employee-vlan	00:05:85:3A:82:77	Persistent	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:79	Persistent	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	ge-0/0/2.0

Now suppose packets have been sent from two of the hosts on **ge-0/0/2** after they have been moved to other interfaces more than five times in 1 second, with **employee-vlan** set to a MAC move limit of 5 with the default action **drop**.

Display the MAC addresses in the table:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 2 learned, 2 persistent entries
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	*	Flood	-	All-members
employee-vlan	00:05:85:3A:82:77	Persistent	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:79	Persistent	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0

Meaning The first sample output shows that with a MAC limit of 4 for each interface, the fifth MAC address on **ge-0/0/2** was not learned because it exceeded the MAC limit. The second sample output shows that MAC addresses for three of the hosts on **ge-0/0/2** were not learned, because the hosts had been moved back more than five times in 1 second.

Interface **ge-0/0/1.0** was enabled for persistent MAC learning, so the MAC addresses associated with this interface are of the type **persistent**.

Verifying That Allowed MAC Addresses Are Working Correctly on the Switch

Purpose Verify that allowed MAC addresses are working on the switch.

Action Display the MAC cache information after five allowed MAC addresses have been configured on interface `ge-0/0/2`:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0

Meaning Because the MAC limit value for this interface has been set to 4, only four of the five configured allowed addresses are learned.

- Related Documentation**
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
 - [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
 - [Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 4737](#)
 - [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
 - [Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 4718](#)
 - [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)
 - [Configuring Port Security \(CLI Procedure\) on page 4756](#)
 - [Configuring Port Security \(J-Web Procedure\)](#)
 - [secure-access-port](#)
 - [secure-access-port on page 4826](#)
 - [show arp inspection statistics on page 4861](#)
 - [show dhcp snooping binding on page 4862](#)
 - [show ethernet-switching table](#)
 - [show ethernet-switching table on page 2211](#)

Example: Configuring Storm Control

Using storm control can prevent problems caused by broadcast storms. You can configure storm control to rate-limit broadcast traffic and unknown unicast traffic at a specified level and to drop packets when the specified traffic level is exceeded, which prevents packets from proliferating and degrading service or causing a security issue. You can also

configure the switch to shut down or temporarily disable an interface when the storm control limit is exceeded.

This example shows how to configure storm control:

- [Requirements on page 4714](#)
- [Overview and Topology on page 4714](#)
- [Configuration on page 4714](#)

Requirements

This example uses the following hardware and software components:

- A switch
- Junos OS Release 11.1 or later

Overview and Topology

A traffic storm occurs when broadcast packets prompt receiving devices to broadcast packets in response. This prompts further responses, creating a snowball effect. The switch is flooded with packets, and the resulting unnecessary traffic leads to poor performance or even a complete loss of service by some clients. Storm control causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service.

Storm control monitors the incoming broadcast traffic and unknown unicast traffic and compares it with the level that you specify. If broadcast traffic and unknown unicast traffic exceed the specified level, the switch drops packets for the controlled traffic types. Storm control is enabled by default on all interfaces, and the default level is 80 percent of the available bandwidth.

This example shows how to configure the storm control level on interface **xe-0/0/0** by setting the level to a traffic rate of 5000000 Kbps, based on the total of the combined broadcast and unknown unicast streams. If broadcast traffic and unknown unicast traffic exceed these levels, the switch drops packets for the controlled traffic types.

Configuration

Step-by-Step Procedure

To configure storm control for a 10-Gigabit Ethernet interface to the equivalent of 50 percent of the available bandwidth:

- Specify the level of allowed broadcast traffic and unknown unicast traffic on a specific interface:

```
[edit ethernet-switching-options]  
user@switch# set storm-control interface xe-0/0/0 bandwidth 5000000
```

Results Display the results of the configuration:

```
[edit ethernet-switching-options]  
user@switch# show storm-control  
interface xe-0/0/0 {
```

```
    bandwidth 5000000;
}
```

Related Documentation

- [Understanding Storm Control on page 4694](#)

Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses. The switch's trusted DHCP server or servers cannot keep up with the requests and can no longer assign IP addresses and lease times to legitimate DHCP clients on the switch. Requests from those clients are either dropped or directed to a rogue DHCP server set up by the attacker.

This example describes how to configure MAC limiting, a port security feature, to protect the switch against DHCP starvation attacks:

- [Requirements on page 4715](#)
- [Overview and Topology on page 4715](#)
- [Configuration on page 4716](#)
- [Verification on page 4717](#)

Requirements

This example uses the following hardware and software components:

- One QFX3500 switch
- Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure MAC limiting, a port security feature, to mitigate DHCP starvation attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN **employee-vlan** on the switch. See [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#).


Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch against one common type of attack, a DHCP starvation attack.

This example shows how to configure port security features on a switch that is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. [Figure 178 on page 4716](#) illustrates the topology for this example.

Figure 178: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 372 on page 4716](#).

Table 372: Components of the Port Security Topology

Properties	Settings
Switch hardware	One QFX3500 switch
VLAN name and ID	employee-vlan
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- No MAC limit is set on any of the interfaces.
- DHCP snooping is disabled on the VLAN **employee-vlan**.
- All access interfaces are untrusted, which is the default setting.

Configuration

To configure the MAC limiting port security feature to protect the switch against DHCP starvation attacks:

CLI Quick Configuration To quickly configure MAC limiting, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 mac-limit 3 action drop
user@switch# set interface ge-0/0/2 mac-limit 3 action drop
```

Step-by-Step Procedure Configure MAC limiting:

1. Configure a MAC limit of **3** on **ge-0/0/1** and specify that packets with new addresses be dropped if the limit has been exceeded on the interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 mac-limit (Access Port Security) 3 action drop
```

2. Configure a MAC limit of **3** on **ge-0/0/2** and specify that packets with new addresses be dropped if the limit has been exceeded on the interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 mac-limit 3 action drop
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
    mac-limit 3 action drop;
}
interface ge-0/0/2.0 {
    mac-limit 3 action drop;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That MAC Limiting Is Working Correctly on the Switch

Purpose Verify that MAC limiting is working on the switch.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the MAC addresses learned when DHCP requests are sent from hosts on **ge-0/0/1** and from hosts on **ge-0/0/2**, with both interfaces set to a MAC limit of **3** with the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
```

VLAN	MAC address	Type	Age	Interfaces
default	*	Flood	-	ge-0/0/2.0
default	00:05:85:3A:82:77	Learn	0	ge-0/0/1.0
default	00:05:85:3A:82:79	Learn	0	ge-0/0/1.0
default	00:05:85:3A:82:80	Learn	0	ge-0/0/1.0
default	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
default	00:05:85:3A:82:83	Learn	0	ge-0/0/2.0
default	00:05:85:3A:82:85	Learn	0	ge-0/0/2.0

Meaning The sample output shows that with a MAC limit of **3** for each interface, the DHCP request for a fourth MAC address on **ge-0/0/2** was dropped because it exceeded the MAC limit.

Because only 3 MAC addresses can be learned on each of the two interfaces, attempted DHCP starvation attacks fail.

Related Documentation

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- [Configuring MAC Limiting on page 4758](#)

Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks

In an Ethernet switching table overflow attack, an intruder sends so many requests from new MAC addresses that the Ethernet switching table fills up and then overflows, forcing the switch to broadcast all messages.

This example describes how to configure MAC limiting and allowed MAC addresses, two port security features, to protect the switch from Ethernet switching table attacks:

- [Requirements on page 4718](#)
- [Overview and Topology on page 4719](#)
- [Configuration on page 4720](#)
- [Verification on page 4721](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or QFX3500 switch

- Junos OS Release 9.0 or later for EX Series switches or Junos OS 12.1 or later for the QFX Series.
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure specific port security features to mitigate common access-interface attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured a VLAN on the switch. See the task for your platform:


Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from an attack on the Ethernet switching table that causes the table to overflow and thus forces the switch to broadcast all messages.

This example shows how to configure port security features on a switch connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches* and “*Example: Setting Up Bridging with Multiple VLANs*” on page 1971 for the QFX Series. That procedure is not repeated here. Figure 179 on page 4719 illustrates the topology for this example.

Figure 179: Network Topology for Basic Port Security



The components of the topology for this example are shown in Table 373 on page 4720.

Table 373: Components of the Port Security Topology

Properties	Settings
Switch hardware	One EX Series switch or one QFX3500 switch
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is subnet's broadcast address
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface. Use the allowed MAC addresses feature to ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- No MAC limit is set on any of the interfaces.
- All access interfaces are untrusted, which is the default setting.

Configuration

To configure MAC limiting and some allowed MAC addresses to protect the switch against Ethernet switching table overflow attacks:

CLI Quick Configuration

To quickly configure MAC limiting, clear the MAC forwarding table, and configure some allowed MAC addresses, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/1 mac-limit 4 action drop
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
exit
exit
clear ethernet-switching-table interface ge-0/0/1
```

Step-by-Step Procedure

Configure MAC limiting and some allowed MAC addresses:

1. Configure a MAC limit of 4 on **ge-0/0/1** and specify that incoming packets with different addresses be dropped once the limit is exceeded on the interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 mac-limit (Access Port Security) 4 action drop
```

2. Clear the current entries for interface ge-0/0/1 from the MAC address forwarding table :

```
user@switch# clear ethernet-switching-table interface ge-0/0/1
```

3. Configure the allowed MAC addresses on ge-0/0/2:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
  mac-limit 4 action drop;
}
interface ge-0/0/2.0 {
  allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83 00:05:85:3a:82:85 ];
}
```

Verification

To confirm that the configuration is working properly:

- [Verifying That MAC Limiting Is Working Correctly on the Switch on page 4721](#)

Verifying That MAC Limiting Is Working Correctly on the Switch

Purpose Verify that MAC limiting is working on the switch.

Action Display the MAC cache information after DHCP requests have been sent from hosts on ge-0/0/1, with the interface set to a MAC limit of 4 with the action **drop**, and after four allowed MAC addresses have been configured on interface ge-0/0/2:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:71	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:74	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:77	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:79	Learn	0	ge-0/0/1.0
employee-vlan	*	Flood	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0

Meaning The sample output shows that with a MAC limit of 4 for the interface, the DHCP request for a fifth MAC address on ge-0/0/1 was dropped because it exceeded the MAC limit and

that only the specified allowed MAC addresses have been learned on the **ge-0/0/2** interface.

Related Documentation

- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Configuring MAC Limiting \(CLI Procedure\)](#)
- [Configuring MAC Limiting on page 4758](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4760](#)
- [Configuring MAC Limiting \(J-Web Procedure\)](#)

Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks

In a rogue DHCP server attack, an attacker has introduced a rogue server into the network, allowing it to give IP address leases to the network's DHCP clients and to assign itself as the gateway device.

This example describes how to configure a DHCP server interface as untrusted to protect the switch from a rogue DHCP server:

- [Requirements on page 4722](#)
- [Overview and Topology on page 4723](#)
- [Configuration on page 4724](#)
- [Verification on page 4724](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or one QFX3500 switch
- Junos OS Release 9.0 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure an untrusted DHCP server interface to mitigate rogue DHCP server attacks, be sure you have:


- Connected the DHCP server to the switch.
- Enabled DHCP snooping on the VLAN.
- Configured a VLAN on the switch. See the task for your platform:
 - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches](#)
 - [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#) for the QFX Series

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from rogue DHCP server attacks.

This example shows how to explicitly configure an untrusted interface on an EX3200-24P switch and a QFX3500 switch. [Figure 180 on page 4723](#) illustrates the topology for this example.

Figure 180: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 374 on page 4723](#).

Table 374: Components of the Port Security Topology

Properties	Settings
Switch hardware	One EX3200-24P, 24 ports (8 PoE ports) or one QFX3500 switch
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is the subnet's broadcast address
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.

- DHCP snooping is enabled on the VLAN **employee-vlan**.
- The interface (port) where the rogue DHCP server has connected to the switch is currently trusted.

Configuration

To configure the DHCP server interface as untrusted because the interface is being used by a rogue DHCP server:

CLI Quick Configuration To quickly set the rogue DHCP server interface as untrusted, copy the following command and paste it into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]  
set interface ge-0/0/8 no-dhcp-trusted
```

Step-by-Step Procedure To set the DHCP server interface as untrusted:

- Specify the interface (port) from which DHCP responses are not allowed:

```
[edit ethernet-switching-options secure-access-port]  
user@switch# set interface ge-0/0/8 no-dhcp-trusted
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]  
user@switch# show  
interface ge-0/0/8.0 {  
    no-dhcp-trusted;  
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the DHCP Server Interface Is Untrusted

Purpose Verify that the DHCP server is untrusted.

- Action**
1. Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.
 2. Display the DHCP snooping information when the port on which the DHCP server connects to the switch is not trusted.

Meaning There is no output from the command because no entries are added to the DHCP snooping database.

Related Documentation

- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4768](#)
- [Enabling a Trusted DHCP Server \(J-Web Procedure\)](#)
- [secure-access-port](#)

- [secure-access-port on page 4826](#)
- [show dhcp snooping binding on page 4862](#)

Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch

You can configure DHCP snooping, dynamic ARP inspection (DAI), and MAC limiting on the access interfaces of switches to protect the switch and the Ethernet LAN against address spoofing and Layer 2 denial-of-service (DoS) attacks. To obtain those basic settings, you can use the switch's default configuration for port security, configure the MAC limit, and enable DHCP snooping and DAI on a VLAN. You can configure those features when the DHCP server is connected to a different switch from the one to which the DHCP clients (network devices) are connected.

This example describes how to configure port security features on a switch whose hosts obtain IP addresses and lease times from a DHCP server connected to a second switch:

- [Requirements on page 4725](#)
- [Overview and Topology on page 4726](#)
- [Configuring a VLAN, Interfaces, and Port Security Features on Switch 1 on page 4727](#)
- [Configuring a VLAN and Interfaces on Switch 2 on page 4729](#)
- [Verification on page 4730](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or QFX3500 switch—"Switch 1" in this example.
- An additional EX Series switch or QFX3500 switch—"Switch 2" in this example. You do not configure port security on this second switch.
- Junos OS Release 9.0 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series.
- A DHCP server connected to Switch 2. You use the server to provide IP addresses to network devices connected to Switch 1.
- At least two network devices (hosts) that you connect to access interfaces on Switch 1. These devices are DHCP clients.

Before you configure DHCP snooping, DAI, and MAC limiting port security features, be sure you have:

- Connected the DHCP server to Switch 2.
- Configured a VLAN on the switch. See the task for your platform:
 - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches](#)
 - ["Example: Setting Up Bridging with Multiple VLANs" on page 1971](#) for the QFX Series

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. To protect the devices from such attacks, you can configure:


- DHCP snooping to validate DHCP server messages
- DAI to protect against ARP spoofing
- MAC limiting to constrain the number of MAC addresses the switch adds to its MAC address cache

This example shows how to configure these port security features on Switch 1. Switch 1 is connected to another switch (Switch 2) that is not configured with port security features. That second switch is connected to a DHCP server. (See [Figure 181 on page 4726](#).) Network devices (hosts) that are connected to Switch 1 send requests for IP addresses (that is, the devices are DHCP clients). Those requests are transmitted from Switch 1 to Switch 2 and then to the DHCP server connected to Switch 2. Responses to the requests are transmitted along the reverse path of the one followed by the requests.

The setup for this example includes the VLAN **employee-vlan** on both switches.

[Figure 181 on page 4726](#) shows the network topology for the example.

Figure 181: Network Topology for Port Security Setup with Two Switches on the Same VLAN



The components of the topology for this example are shown in [Table 375 on page 4727](#).

Table 375: Components of Port Security Setup on Switch 1 with a DHCP Server Connected to Switch 2

Properties	Settings
Switch hardware	One EX Series switch or one QFX3500 switch (Switch 1), and an additional EX Series switch or QFX3500 switch (Switch 2)
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is subnet's broadcast address
Trunk interface on both switches	ge-0/0/11
Access interfaces on Switch 1	ge-0/0/1 , ge-0/0/2 , and ge-0/0/3
Access interface on Switch 2	ge-0/0/1
Interface for DHCP server	ge-0/0/1 on Switch 2

Switch 1 is initially configured with the default port security setup. In the default configuration on the switch:

- Secure port access is activated on the switch.
- The switch does not drop any packets, which is the default setting.
- DHCP snooping and dynamic ARP inspection (DAI) are disabled on all VLANs.
- All access interfaces are untrusted and trunk interfaces are trusted; these are the default settings.

In the configuration tasks for this example, you configure a VLAN on both switches.

In addition to configuring the VLAN, you enable DHCP snooping on Switch 1. In this example, you also enable DAI and a MAC limit of 5 on Switch 1.

Because the interface that connects Switch 2 to Switch 1 is a trunk interface, you do not have to configure this interface to be trusted. As noted above, trunk interfaces are automatically trusted, so DHCP messages coming from the DHCP server to Switch 2 and then on to Switch 1 are trusted.

Configuring a VLAN, Interfaces, and Port Security Features on Switch 1

CLI Quick Configuration

To quickly configure a VLAN, interfaces, and port security features, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options secure-access-port interface ge-0/0/1 mac-limit 5 action drop
clear ethernet-switching table interface ge-0/0/1
set ethernet-switching-options secure-access-port vlan employee-vlan arp-inspection
set ethernet-switching-options secure-access-port vlan employee-vlan examine-dhcp
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 20
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members 20
```

```
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members 20
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members 20
set vlans employee-vlan vlan-id 20
```

Step-by-Step Procedure To configure MAC limiting, a VLAN, and interfaces on Switch 1 and enable DAI and DHCP on the VLAN:

1. Configure the VLAN **employee-vlan** with VLAN ID **20**:

```
[edit vlans]
user@switch1# set employee-vlan vlan-id 20
```
2. Configure an interface on Switch 1 as a trunk interface:

```
[edit interfaces]
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
```
3. Associate the VLAN with interfaces **ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, and **ge-0/0/11**:

```
[edit interfaces]
user@switch1# set ge-0/0/1 unit 0 family ethernet-switching vlan members 20
user@switch1# set ge-0/0/2 unit 0 family ethernet-switching vlan members 20
user@switch1# set ge-0/0/3 unit 0 family ethernet-switching vlan members 20
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members 20
```
4. Enable DHCP snooping on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch1# set vlan employee-vlan examine-dhcp
```
5. Enable DAI on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch1# set vlan employee-vlan arp-inspection
```
6. Configure a MAC limit of **5** on **ge-0/0/1** and use the default action, drop (packets with new addresses are dropped if the limit has been exceeded):

```
[edit ethernet-switching-options secure-access-port]
user@switch1# set interface ge-0/0/1 mac-limit 5 drop
```
7. Clear the existing MAC address table entries from interface **ge-0/0/1**:

```
user@switch1# clear ethernet-switching table interface ge-0/0/1
```

Results Display the results of the configuration:

```
[edit]
user@switch1# show
ethernet-switching-options {
  secure-access-port {
    interface ge-0/0/1.0 {
      mac-limit 5 action drop;
    }
    vlan employee-vlan {
      arp-inspection;
      examine-dhcp;
    }
  }
}
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
```

```

        members 20;
    }
}
}
ge-0/0/2 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members 20;
            }
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members 20;
            }
        }
    }
}
ge-0/0/11 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 20;
            }
        }
    }
}
vlands {
    employee-vlan {
        vlan-id 20;
    }
}

```

Configuring a VLAN and Interfaces on Switch 2

To configure the VLAN and interfaces on Switch 2:

CLI Quick Configuration

To quickly configure the VLAN and interfaces on Switch 2, copy the following commands and paste them into the switch terminal window:

```

[edit]
set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set vlans employee-vlan vlan-id 20

```

Step-by-Step Procedure

To configure the VLAN and interfaces on Switch 2:

1. Configure an interface on Switch 2 as a trunk interface:
[edit interfaces]

- ```
user@switch2# set ge-0/0/11 unit 0 ethernet-switching port-mode trunk
```
2. Associate the VLAN with interfaces **ge-0/0/1** and **ge-0/0/11**:
- ```
[edit interfaces]
user@switch2# set ge-0/0/1 unit 0 family ethernet-switching vlan members 20
user@switch2# set ge-0/0/11 unit 0 family ethernet-switching vlan members 20
```

Results Display the results of the configuration:

```
[edit]
user@switch2# show
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members 20;
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 20;
        }
      }
    }
  }
}
vlans {
  employee-vlan {
    vlan-id 20;
  }
}
```

Verification

To confirm that the configuration is working properly.

- [Verifying That DHCP Snooping Is Working Correctly on Switch 1 on page 4730](#)
- [Verifying That DAI Is Working Correctly on Switch 1 on page 4731](#)
- [Verifying That MAC Limiting Is Working Correctly on Switch 1 on page 4731](#)

Verifying That DHCP Snooping Is Working Correctly on Switch 1

Purpose Verify that DHCP snooping is working on Switch 1.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface through which Switch 2 sends the DHCP server replies to clients connected to Switch 1 is trusted. The server has provided the IP addresses and leases:

```
user@switch1> show dhcp snooping binding
```

DHCP Snooping Information:

MAC Address	IP Address	Lease	Type	VLAN	Interface
00:05:85:3A:82:77	192.0.2.17	600	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:79	192.0.2.18	653	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:80	192.0.2.19	720	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:81	192.0.2.20	932	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:83	192.0.2.21	1230	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:90	192.0.2.20	932	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:3A:82:91	192.0.2.21	1230	dynamic	employee-vlan	ge-0/0/3.0

Meaning The output shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

Verifying That DAI Is Working Correctly on Switch 1

Purpose Verify that DAI is working on Switch 1.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch1> show arp inspection statistics
```

ARP inspection statistics:

Interface	Packets received	ARP inspection pass	ARP inspection failed
ge-0/0/1.0	7	5	2
ge-0/0/2.0	10	10	0
ge-0/0/3.0	18	15	3

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Verifying That MAC Limiting Is Working Correctly on Switch 1

Purpose Verify that MAC limiting is working on Switch 1.

Action Display the MAC addresses that are learned when DHCP requests are sent from hosts on **ge-0/0/1**:

```
user@switch1> show ethernet-switching table
```

```
Ethernet-switching table: 6 entries, 5 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:77	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:79	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:83	Learn	0	ge-0/0/1.0
employee-vlan	*	Flood	-	ge-0/0/1.0

Meaning The sample output shows that five MAC addresses have been learned for interface **ge-0/0/1**, which corresponds to the MAC limit of **5** set in the configuration. The last line of the output shows that a sixth MAC address request was dropped, as indicated by the asterisk (*) in the **MAC address** column.

- Related Documentation**
- [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Configuring Port Security \(CLI Procedure\) on page 4756](#)
 - [Configuring Port Security \(J-Web Procedure\)](#)
 - [secure-access-port](#)
 - [secure-access-port on page 4826](#)
 - [show arp inspection statistics on page 4861](#)
 - [show dhcp snooping binding on page 4862](#)
 - [show ethernet-switching table](#)
 - [show ethernet-switching table on page 2211](#)

Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks

In an ARP spoofing attack, the attacker associates its own MAC address with the IP address of a network device connected to the switch. Traffic intended for that IP address is now sent to the attacker instead of being sent to the intended destination. The attacker can send faked, or “spoofed,” ARP messages on the LAN.



NOTE: When dynamic ARP inspection (DAI) is enabled, the switch logs the number of invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. You can use these log messages to discover ARP spoofing on the network.

This example describes how to configure DHCP snooping and dynamic ARP inspection (DAI), two port security features, to protect the switch against ARP spoofing attacks:

- [Requirements on page 4733](#)
- [Overview and Topology on page 4733](#)
- [Configuration on page 4734](#)
- [Verification on page 4735](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or one QFX3500 switch
- Junos OS Release 11.4 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP snooping and DAI (two port security features) to mitigate ARP spoofing attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured a VLAN on the switch. See the task for your platform:
 - *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
 - [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#) for the QFX Series


Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch against one common type of attack, an ARP spoofing attack.

In an ARP spoofing attack, the attacker sends faked ARP messages, thus creating various types of problems on the LAN—for example, the attacker might launch a man-in-the middle attack.

This example shows how to configure port security features on a switch that is connected to a DHCP server. The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches* and [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#) for the QFX Series. That procedure is not repeated here. [Figure 182 on page 4734](#) illustrates the topology for this example.

Figure 182: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 376 on page 4734](#).

Table 376: Components of the Port Security Topology

Properties	Settings
Switch hardware	One EX3200-24P, 24 ports (8 PoE ports) or one QFX3500 switch
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is the subnet's broadcast address
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- DHCP snooping is disabled on the VLAN **employee-vlan**.
- All access ports are untrusted, which is the default setting.

Configuration

To configure DHCP snooping and dynamic ARP inspection (DAI) to protect the switch against ARP attacks:

CLI Quick Configuration To quickly configure DHCP snooping and dynamic ARP inspection (DAI), copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/8 dhcp-trusted
user@switch# set vlan employee-vlan examine-dhcp
user@switch# set vlan employee-vlan arp-inspection
```

Step-by-Step Procedure Configure DHCP snooping and dynamic ARP inspection (DAI) on the VLAN:

1. Set the **ge-0/0/8** interface as trusted:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/8 dhcp-trusted
```
2. Enable DHCP snooping on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan examine-dhcp
```
3. Enable DAI on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan arp-inspection
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/8.0 {
  dhcp-trusted;
}
vlan employee-vlan {
  arp-inspection;
  examine-dhcp;
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying That DHCP Snooping Is Working Correctly on the Switch on page 4735](#)
- [Verifying That DAI Is Working Correctly on the Switch on page 4736](#)

Verifying That DHCP Snooping Is Working Correctly on the Switch

Purpose Verify that DHCP snooping is working on the switch.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the port on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp-snooping binding
DHCP Snooping Information:
MAC Address      IP Address      Lease    Type    VLAN      Interface
-----
00:05:85:3A:82:77 192.0.2.17      600     dynamic employee-vlan ge-0/0/1.0
00:05:85:3A:82:79 192.0.2.18      653     dynamic employee-vlan ge-0/0/1.0
00:05:85:3A:82:80 192.0.2.19      720     dynamic employee-vlan ge-0/0/2.0
00:05:85:3A:82:81 192.0.2.20      932     dynamic employee-vlan ge-0/0/2.0
00:05:85:3A:82:83 192.0.2.21      1230    dynamic employee-vlan ge-0/0/2.0
00:05:85:27:32:88 192.0.2.22      3200    dynamic employee-vlan ge-0/0/3.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

Verifying That DAI Is Working Correctly on the Switch

Purpose Verify that DAI is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface      Packets received  ARP inspection pass  ARP inspection failed
-----
ge-0/0/1.0      7                 5                    2
ge-0/0/2.0     10                10                   0
ge-0/0/3.0     12                12                   0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Enabling DHCP Snooping \(CLI Procedure\) on page 4764](#)
 - [Enabling DHCP Snooping \(J-Web Procedure\)](#)
 - [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4766](#)

- [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)
- [secure-access-port](#)
- [secure-access-port on page 4826](#)
- [show arp inspection statistics on page 4861](#)
- [show dhcp snooping binding on page 4862](#)

Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks

In one type of attack on the DHCP snooping database, an intruder introduces a DHCP client on an untrusted access interface with a MAC address identical to that of a client on another untrusted interface. The intruder then acquires the DHCP lease of that other client, thus changing the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

This example describes how to configure allowed MAC addresses, a port security feature, to protect the switch from DHCP snooping database alteration attacks:

- [Requirements on page 4737](#)
- [Overview and Topology on page 4737](#)
- [Configuration on page 4739](#)
- [Verification on page 4739](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or one QFX3500 switch
- Junos OS Release 11.4 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure specific port security features to mitigate common access-interface attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured a VLAN on the switch. See the task for your platform:
 - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches](#)
 - [“Example: Setting Up Bridging with Multiple VLANs” on page 1971](#) for the QFX Series

Overview and Topology


Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from an attack on the DHCP snooping database that alters the MAC addresses assigned to some clients.

This example shows how to configure port security features on a switch that is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch.

[Figure 183 on page 4738](#) illustrates the topology for this example.

Figure 183: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 377 on page 4738](#).

Table 377: Components of the Port Security Topology

Properties	Settings
Switch hardware	One EX3200-24P, 24 ports (8 PoE ports) or one QFX3500 switch
VLAN name and ID	employee-vlan , tag 20
VLAN subnets	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is the subnet's broadcast address
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8
Interface for DHCP server	ge-0/0/8

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- DHCP snooping is enabled on the VLAN **employee-vlan**.
- All access ports are untrusted, which is the default setting.

Configuration

To configure allowed MAC addresses to protect the switch against DHCP snooping database alteration attacks:

CLI Quick Configuration To quickly configure some allowed MAC addresses on an interface, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Step-by-Step Procedure To configure some allowed MAC addresses on an interface:
Configure the five allowed MAC addresses on an interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/2.0 {
    allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83 00:05:85:
:3a:82:85 00:05:85:3a:82:88 ];
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying That Allowed MAC Addresses Are Working Correctly on the Switch on page 4739](#)

Verifying That Allowed MAC Addresses Are Working Correctly on the Switch

Purpose Verify that allowed MAC addresses are working on the switch.

Action Display the MAC cache information:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 6 entries, 5 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:88	Learn	0	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0

Meaning The output shows that the five MAC addresses configured as allowed MAC addresses have been learned and are displayed in the MAC cache. The last MAC address in the list, one that had not been configured as allowed, has not been added to the list of learned addresses.

- Related Documentation**
- [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Configuring MAC Limiting \(CLI Procedure\)](#)
 - [Configuring MAC Limiting \(J-Web Procedure\)](#)
 - [secure-access-port](#)
 - [secure-access-port on page 4826](#)
 - [show ethernet-switching table](#)
 - [show ethernet-switching table on page 2211](#)

Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

This example describes how to configure DHCP option 82 on a switch with DHCP clients, DHCP server, and switch all on the same VLAN:

- [Requirements on page 4740](#)
- [Overview and Topology on page 4741](#)
- [Configuration on page 4742](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch or one QFX3500 switch

- Junos OS Release 9.3 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP option 82 on the switch, be sure you have:

- Connected and configured the DHCP server.



NOTE: Your DHCP server must be configured to accept DHCP option 82. If it is not configured for DHCP option 82, it does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configured the **employee** VLAN on the switch and associated the interfaces on which the clients and the server connect to the switch with that VLAN. See the task for your platform:
 - *Configuring VLANs for EX Series Switches (CLI Procedure)*
 - “Configuring VLANs” on page 2048 for the QFX Series

Overview and Topology

If DHCP option 82 is enabled on the switch, then when a network device—a DHCP client—that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or other parameter for the client.


DHCP option 82 is enabled on an individual VLAN or on all VLANs on the switch.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.

Figure 184 on page 4742 illustrates the topology for this example.

Figure 184: Network Topology for Configuring DHCP Option 82 on a Switch That Is on the Same VLAN as the DHCP Clients and the DHCP Server



In this example, you configure DHCP option 82 on the switch. The switch connects to the DHCP server on interface `ge-0/0/8`. The DHCP clients connect to the switch on interfaces `ge-0/0/1`, `ge-0/0/2`, and `ge-0/0/3`. The switch, server, and clients are all members of the employee VLAN.

Configuration

CLI Quick Configuration

To quickly configure DHCP option 82, copy the following commands and paste them into the switch terminal window:

```
set ethernet-switching-options secure-access-port vlan employee dhcp-option82
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 circuit-id prefix
hostname
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 circuit-id
use-vlan-id
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id
prefix mac
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id
use-string employee-switch1
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 vendor-id
```

Step-by-Step Procedure

To configure DHCP option 82:

1. Specify DHCP option 82 for the **employee** VLAN:


```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82
```
2. Configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):


```
[edit ethernet-switching-options secure-access-port]
```

- ```
user@switch# set vlan employee dhcp-option82 circuit-id prefix hostname
```
3. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):
 

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id use-vlan-id
```
  4. Specify that the remote ID suboption be included in the DHCP option 82 information:
 

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id
```
  5. Configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):
 

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id prefix mac
```
  6. Specify that the remote ID suboption value contain a character string (here, the string is **employee-switch1**):
 

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id use-string employee-switch1
```
  7. Configure a vendor ID suboption value, and use the default value. To use the default value, do not type a character string after the **vendor-id** option keyword:
 

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 vendor-id
```

**Results** Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
vlan employee {
 dhcp-option82 {
 circuit-id {
 prefix hostname;
 use-vlan-id;
 }
 remote-id {
 prefix mac;
 use-string employee-switch1;
 }
 vendor-id;
 }
}
```

**Related Documentation**

- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4770](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.
- *secure-access-port*
- [secure-access-port on page 4826](#)

## Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

This example describes how to configure DHCP option 82 on a switch that is on the same VLAN with the DHCP clients but on a different VLAN from the DHCP server. In this example, the switch acts as a relay agent:

- [Requirements on page 4744](#)
- [Overview and Topology on page 4745](#)
- [Configuration on page 4745](#)

### Requirements

---

This example uses the following hardware and software components:

- One EX4200-24P switch or one QFX3500 switch
- Junos OS Release 9.3 or later for EX Series switches or Junos OS Release 12.1 or later for the QFX Series
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP option 82 on the switch, be sure you have:

- Connected and configured the DHCP server.



**NOTE:** Your DHCP server must be configured to accept DHCP option 82. If it is not configured for DHCP option 82, it does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configured the **employee** VLAN on the switch and associated the interfaces on which the clients connect to the switch with that VLAN. See the task for your platform:
  - [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)
  - [“Configuring VLANs” on page 2048](#) for the QFX Series
- Configured the **corporate** VLAN for the DHCP server.
- Configured the switch as a BOOTP relay agent. See *DHCP/BOOTP Relay for Switches Overview*.
- Configured the routed VLAN interface (RVI) to allow the switch to relay packets to the server and receive packets from the server. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or [“Configuring Routed VLAN Interfaces” on page 2046](#) for the QFX Series.

## Overview and Topology

If DHCP option 82 is enabled on the switch, then when a network device—a DHCP client—that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request (in this setting, it relays the request) to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or other parameter for the client.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.

In this example, you configure option 82 on the switch. The switch is configured as a BOOTP relay agent. The switch connects to the DHCP server through the routed VLAN interface (RVI) that you configured. The switch and clients are members of the **employee** VLAN. The DHCP server is a member of the **corporate** VLAN.

## Configuration

To configure DHCP option 82:

### CLI Quick Configuration

To quickly configure DHCP option 82, copy the following commands and paste them into the switch terminal window:

```
set forwarding-options helpers bootp dhcp-option82
set forwarding-options helpers bootp dhcp-option82 circuit-id prefix hostname
set forwarding-options helpers bootp dhcp-option82 circuit-id use-vlan-id
set forwarding-options helpers bootp dhcp-option82 remote-id
set forwarding-options helpers bootp dhcp-option82 remote-id prefix mac
set forwarding-options helpers bootp dhcp-option82 remote-id use-string employee-switch1
set forwarding-options helpers bootp dhcp-option82 vendor-id
```

### Step-by-Step Procedure

To configure DHCP option 82:

1. Specify DHCP option 82 for the **employee** VLAN:  

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82
```
2. Configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):  

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 circuit-id prefix hostname
```
3. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

- ```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 circuit-id use-vlan-id
```
4. Specify that the remote ID suboption be included in the DHCP option 82 information:
- ```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id
```
5. Configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):
- ```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id prefix mac
```
6. Specify that the remote ID suboption value contains a character string (here, the string is **employee-switch1**):
- ```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id use-string employee-switch1
```
7. Configure a vendor ID suboption value, and use the default value. To use the default value, do not type a character string after the **vendor-id** option keyword:
- ```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id
```

Results Check the results of the configuration:

```
[edit forwarding-options helpers bootp]
user@switch# show
dhcp-option82 {
  circuit-id {
    prefix hostname;
    use-vlan-id;
  }
  remote-id {
    prefix mac;
    use-string employee-switch1;
  }
  vendor-id;
}
```

**Related
Documentation**

- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740](#)
- [Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4773](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.
- *forwarding-options*

Firewall and Policer Configuration Tasks

- [Configuring Firewall Filters on page 4747](#)
- [Applying Firewall Filters to Interfaces on page 4750](#)
- [Assigning Forwarding Classes and Loss Priority on page 4751](#)

- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

Configuring Firewall Filters

You can configure firewall filters in a switch to control traffic that enters switch ports or enters and exits VLANs and Layer 3 (routed) interfaces. To use a firewall filter, you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

- [Configuring a Firewall Filter on page 4747](#)
- [Applying a Firewall Filter to a Port on page 4749](#)
- [Applying a Firewall Filter to a VLAN on page 4749](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 4749](#)

Configuring a Firewall Filter

To configure a firewall filter:

1. Configure the family address type, filter name, term name, and at least one match condition—for example, match on packets that contain a specific source address:

```
[edit]
user@switch# set firewall family ethernet-switching filter ingress-port-filter term term-one
from source-address 192.0.2.14
```

For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching**. For a firewall filter that is applied to a Layer 3 (routed) interface, specify the family address type **inet**.

The filter and term names can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. Each filter name must be unique. A filter can contain one or more terms, and each term name must be unique within a filter.

2. Configure additional match conditions. For example, match on packets that contain a specific source port:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one from]
user@switch# set source-port 80
```

You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term. The **from** statement is optional, but if included in a term, it cannot be empty. If you omit the **from** statement, all packets are considered to match.

3. If you want to apply a firewall filter to multiple interfaces and be able to see counters specific to each interface, configure the **interface-specific** option:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set interface-specific
```

4. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term. You can specify an action and action modifiers:

- To specify a filter action, for example, to discard packets that match the conditions of the filter term:

```
[edit firewall] family ethernet-switching filter ingress-port-filter term
term-one then]
```

```
user@switch# set discard
```

You can specify no more than one action (**accept**, **discard**, **reject**, **routing-instance**, or **vlan**) per term.

- To specify action modifiers, for example, to count and classify packets to a forwarding class:

```
[edit firewall] family ethernet-switching filter ingress-port-filter term
term-one then]
```

```
user@switch# set count counter-one
```

```
user@switch# set forwarding-class expedited-forwarding
```

```
user@switch# set loss-priority high
```

You can specify any of the following action modifiers in a **then** statement:

- **analyzer *analyzer-name***—Mirror port traffic to a specified analyzer, which you must configure at the **[ethernet-switching-options]** level.
- **count *counter-name***—Count the number of packets that pass this filter term.



NOTE: We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.



NOTE: On the QFX Series, filters automatically count packets that have been dropped on ingress because of cyclic redundancy check (CRC) errors.

- **forwarding-class *class***—Assign packets to a forwarding class.
- **log**—Log the packet header information in the Routing Engine.
- **loss-priority *priority***—Set the priority of dropping a packet.
- **policer *policer-name***—Apply rate-limiting to the traffic.
- **syslog**—Log an alert for this packet.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you should always explicitly configure an action in the **then** statement. You can include no more than one action statement, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.



NOTE: Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

Applying a Firewall Filter to a Port

To apply a firewall filter to an ingress port:

1. Provide a meaningful description of the firewall filter in the configuration of the port to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/0/6 description "filter to limit tcp traffic at trunk port for employee-vlan"
```

2. Apply the filter to the interface, specifying the unit number, family address type, the direction of the filter (for packets entering the port), and the filter name:

```
[edit]
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input ingress-port-filter
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.



NOTE: You can apply only one filter to a port for a given direction (ingress or egress).

Applying a Firewall Filter to a VLAN

To apply a firewall filter to a VLAN:

1. Provide a meaningful description of the firewall filter in the configuration of the VLAN to which the filter will be applied:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 description "filter to block rogue devices on employee-vlan"
```

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:

- To apply a filter to match packets that are entering the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter input ingress-vlan-rogue-block
```

- To apply a firewall filter to match packets that are exiting the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter output egress-vlan-filter
```



NOTE: You can apply only one filter to a VLAN for a given direction (ingress or egress).

Applying a Firewall Filter to a Layer 3 (Routed) Interface

To apply a firewall filter to a Layer 3 routed interface:

1. Provide a meaningful description of the firewall filter in the configuration of the interface to which the filter will be applied:

```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 description "filter to count and monitor traffic on layer 3 interface"
```

2. You can apply firewall filters to filter packets that enter or exit a Layer 3 routed interface:

- To apply a firewall filter to filter packets that enter a Layer 3 interface:

```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter input ingress-router-filter
```

- To apply a firewall filter to filter packets that exit a Layer 3 interface:

```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter output egress-router-filter
```

For firewall filters applied to Layer 3 routed interfaces, the family address type must be **inet**.



NOTE: You can apply only one filter to an interface for a given direction (ingress or egress).

Related Documentation

- [Overview of Firewall Filters on page 4635](#)
- [Firewall Filter Match Conditions and Actions on page 4644](#)
- [Verifying That Firewall Filters Are Operational on page 4848](#)
- [Monitoring Firewall Filter Traffic on page 4845](#)
- [Configuring Port Mirroring on page 4904](#)

Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
```

```
user@switch# set interface-name unit logical-unit-number family family-name filter (input | output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface. (You cannot apply an output filter to a loopback interface.)



NOTE: When you create a loopback interface, it is important to apply a filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface lo0, you include the **apply-groups** statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including lo0 and other loopback interfaces.

Related Documentation • [Configuring Firewall Filters on page 4747](#)

Assigning Forwarding Classes and Loss Priority

You can configure firewall filters to assign packet loss priority (PLP) and forwarding classes so that if congestion occurs, the marked packets can be dropped according to the priority you set. The valid match conditions are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. In other words, you can set the forwarding class and the PLP for each packet entering or an interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.



NOTE: Junos OS assigns forwarding classes and PLP on ingress only. Do not use a filter that assigns forwarding classes or PLP as an egress filter.

When tricolor marking is enabled, a switch supports four PLP designations: **low**, **medium-low**, **medium-high**, and **high**. You can also specify any of the forwarding classes listed in [Table 378 on page 4751](#)

Table 378: Unicast Forwarding Classes

Unicast Forwarding Class	For CoS Traffic Type
be	Best-effort traffic
no-loss	Guaranteed delivery for TCP traffic
fcoe	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic
nc	Network-control traffic

To assign forwarding classes in firewall filters:

1. Configure the family address type and filter name:

```
[edit]
user@switch# edit firewall family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter as appropriate, including the **forwarding-class** and **loss-priority** action modifiers. For example, each of the following terms in the filter examines various packet header fields and assigns the appropriate forwarding class and packet loss priority:

- The term **corp-traffic** matches all IPv4 packets with a **10.1.1.0/24** source address and assigns the packets to forwarding class **no-loss** with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term corp-traffic from source-address 10.1.1.0/24;
user@switch# set term corp-traffic then forwarding-class no-loss
user@switch# set term corp-traffic then loss-priority low
```

- The term **data-traffic** matches all IPv4 packets with a **10.1.2.0/24** source address and assigns the packets to forwarding class **be** (best effort) with a loss priority of **medium-high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from source-address 10.1.2.0/24;
user@switch# set term data-traffic then forwarding-class be
user@switch# set term data-traffic then loss-priority medium-high
```

- Because the loss of network-generated packets can jeopardize proper network operation, the delay of these packets is preferable to discarding these packets. The term **network-traffic** assigns the packets with an IP precedence of **net-control** to forwarding class **nc** (network control) with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class nc
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the packets to forwarding class **be** with a loss priority of **high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic then forwarding-class be
user@switch# set term accept-traffic then loss-priority high
```

3. Apply the filter **ingress-filter** to a port, VLAN, or Layer 3 interface. For information about applying the filter, see [“Configuring Firewall Filters” on page 4747](#). (Assigning forwarding classes and PLP is supported only on ingress filters.)

Related Documentation

- [Configuring Firewall Filters on page 4747](#)
- [Verifying That Firewall Filters Are Operational on page 4848](#)
- [Monitoring Firewall Filter Traffic on page 4845](#)
- [Overview of Policers on page 4664](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Forwarding Classes on page 5469](#)

Configuring Color-Blind Egress Policers for Medium-Low PLP

If you use color-blind mode and want to configure an egress policer that marks packets to have medium-low PLP, you must configure a single-rate two-color policer at the **[edit firewall policer *policer-name*]** hierarchy level, because color-blind mode does not support medium-low priority. For example:

1. Specify the name of the policer, the bandwidth limit in bits per second (bps) to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit]
user@switch# set firewall policer policer-name if-exceeding bandwidth-limit bytes
burst-size-limit bytes
```

2. Specify medium-low loss priority for matching packets:

```
[edit]
```

```
user@switch# set firewall policer policer-name then loss-priority medium-low;
```

3. Apply the filter to a port, VLAN, or Layer 3 interface.

Related Documentation

- [Overview of Policers on page 4664](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 4669](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 4671](#)
- [Configuring Firewall Filters on page 4747](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

You can rate-limit traffic by configuring a policer and specifying it as an action modifier for a term in a firewall filter. By default, if you specify the same policer in multiple terms, Junos OS creates a separate policer instance for each term and applies rate limiting separately for each instance. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, each policer instance enforces a 1-Gbps limit. In this case, the total bandwidth allowed by the filter is 3 Gbps.

You can also configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps.



NOTE: You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

1. [Configuring Two-Color Policers on page 4753](#)
2. [Configuring Three-Color Policers on page 4754](#)
3. [Specifying Policers in a Firewall Filter Configuration on page 4754](#)
4. [Applying a Firewall Filter That Includes a Policer on page 4755](#)

Configuring Two-Color Policers

To configure a two-color policer:

1. Specify the name of the policer, the bandwidth limit to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit firewall]
```

```
user@switch# set policer policer-name <filter-specific> if-exceeding bandwidth-limit bps  
burst-size-limit bytes
```

The policer name can contain letters, numbers, and hyphens (-) and can have as many as 64 characters.

The range for the bandwidth limit is 32000 (32k) through 102,300,000,000 (102300m) bps.

To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur and divide the result by 8:

maximum burst size = (interface bandwidth) X (allowable time for burst) / (8 bits/byte)

The range for the burst-size limit is 1 through 2,147,450,880 bytes.

2. Specify the policer action to discard or assign a loss priority to packets that exceed the rate limits:

```
[edit firewall policer policer-name]  
user@switch# set then (discard | loss-priority low | loss-priority high)
```

Configuring Three-Color Policers

To configure a three-color policer:

1. Specify the name of the policer and (optionally) whether to automatically discard packets with high loss priority (PLP):

```
[edit firewall]  
user@switch# set three-color-policer policer-name  
user@switch# set three-color-policer policer-name action loss-priority high then discard
```

2. Specify whether the three-color policer should be single-rate or two-rate and whether it should be color-aware or color-blind:

```
[edit firewall three-color-policer policer-name]  
user@switch# set (single-rate | two-rate) (color-aware | color-blind)
```

3. For single-rate three-color policers, configure the CIR, CBS, and EBS:

```
[edit firewall three-color-policer policer-name single-rate]  
user@switch# set committed-information-rate bps  
user@switch# set committed-burst-size bytes  
user@switch# set excess-burst-size bytes
```

4. For two-rate three-color policers, configure the CIR, CBS, PIR, and PBS:

```
[edit firewall three-color-policer policer-name single-rate]  
user@switch# set committed-information-rate bps  
user@switch# set committed-burst-size bytes  
user@switch# set peak-information-rate bps  
user@switch# set peak-burst-size bytes
```

Specifying Policers in a Firewall Filter Configuration

To use a two-color policer, configure a filter term that includes the action **policer**:

```
[edit firewall family family-name]  
user@switch# set filter filter-name term name then name
```

For example, the following commands apply a two-color policer to all packets sent from 192.0.2.0/24.

```
[edit firewall family family-name]  
user@switch# set filter limit—hosts term term1 from source-address 192.0.2.0/24  
user@switch# set filter limit—hosts term term1 then policer policer1
```

To use a three-color policer, configure a filter term that includes the action **three-color-policer**:

```
[edit firewall family name]
user@switch# set filter name term name from match-condition
user@switch# set filter name term name then three-color-policer (single-rate | two-rate) name
```

For example, the following commands apply a single-rate three-color policer to all packets received or sent by interface **ge-0/0/6** (depending on whether the filter is an ingress or egress filter).

```
[edit firewall family name]
user@switch# set filter srTCM term term-one from interface ge-0/0/6
user@switch# set filter srTCM term term-one then three-color-policer single-rate srTCM1-ca
```

You must specify whether the three-color policer is single-rate or two-rate, and this must match the policer itself. Otherwise, the configuration listing includes an error message indicating that the three-color policer you referenced in the filter does not exist.

Applying a Firewall Filter That Includes a Policer

A firewall filter that includes one or more policer action modifiers must be applied to a port, VLAN, or Layer 3 interface like any other filter. For information about applying firewall filters, see “Configuring Firewall Filters” on page 4747.



NOTE: You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

Related Documentation

- [Configuring Firewall Filters on page 4747](#)
- [Overview of Policers on page 4664](#)
- [Verifying That Two-Color Policers Are Operational on page 4856](#)
- [Verifying That Three-Color Policers Are Operational on page 4855](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752](#)

Port Security Configuration Tasks

- [Configuring Port Security \(CLI Procedure\) on page 4756](#)
- [Configuring MAC Limiting on page 4758](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4760](#)
- [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\) on page 4762](#)
- [Configuring the none Action to Override a MAC Limit Applied to All Interfaces \(CLI Procedure\) on page 4762](#)
- [Configuring Static ARP Entries on page 4763](#)
- [Configuring Static IP Addresses for DHCP Bindings on Access Ports \(CLI Procedure\) on page 4763](#)
- [Enabling DHCP Snooping \(CLI Procedure\) on page 4764](#)

- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4766](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4768](#)
- [Enabling a Trusted Port for DHCP on page 4769](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4770](#)
- [Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4773](#)

Configuring Port Security (CLI Procedure)

Ethernet LANs are vulnerable to attacks such as address spoofing and Layer 2 denial of service (DoS) on network devices. Port security features such as DHCP snooping, DAI (dynamic ARP inspection), MAC limiting, MAC move limiting, and persistent MAC learning, as well as trusted DHCP server, help protect the access ports on the switch against the losses of information and productivity that can result from such attacks.

Depending on the particular feature, you can configure the port security feature either on:

- VLANs—A specific VLAN or all VLANs
- Interfaces—A specific interface or all interfaces



NOTE: If you configure one of the port security features on all VLANs or all interfaces, the switch software enables that port security feature on all VLANs and all interfaces that are not explicitly configured with other port security features.

However, if you do explicitly configure one of the port security features on a specific VLAN or on a specific interface, you must explicitly configure any additional port security features that you want to apply to that VLAN or interface. Otherwise, the switch software automatically applies the default values for the feature.

For example, if you enable DHCP snooping on all VLANs and decide to explicitly enable IP source guard only on a specific VLAN, you must also explicitly enable DHCP snooping on that specific VLAN. Otherwise, the default value of no DHCP snooping applies to that VLAN.

To configure port security features using the CLI:

- [Enabling DHCP Snooping on page 4757](#)
- [Enabling Dynamic ARP Inspection \(DAI\) on page 4757](#)
- [Limiting Dynamic MAC Addresses on an Interface on page 4757](#)
- [Enabling Persistent MAC Learning on an Interface on page 4757](#)
- [Limiting MAC Address Movement on page 4758](#)
- [Configuring Trusted DHCP Servers on an Interface on page 4758](#)

Enabling DHCP Snooping

You can configure DHCP snooping to allow the device to monitor DHCP messages received, ensure that hosts only use the IP addresses assigned to them, and allow access only to authorized DHCP servers.

To enable DHCP snooping:

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan default examine-dhcp
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all examine-dhcp
```

Enabling Dynamic ARP Inspection (DAI)

You can enable DAI to protect against ARP snooping. To enable DAI:

- On a single VLAN (here, the VLAN is **employee-vlan**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan arp-inspection
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all arp-inspection
```

Limiting Dynamic MAC Addresses on an Interface

Limit the number of dynamic MAC addresses allowed on an interface and specify the action to take if the limit is exceeded—for example, set a MAC limit of **5** with an action of **drop**:

- On a single interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 mac-limit 5 action drop
```

- On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all mac-limit 5 action drop
```

You can also specify the actions **log** (do not drop the packet but generate an alarm, an SNMP trap, or a system log entry), **none** (no action), or **shutdown** (disable the interface and generate an alarm) to occur if the number of dynamic MAC addresses is exceeded.

Enabling Persistent MAC Learning on an Interface

You can configure learned MAC addresses to persist on an interface across restarts of the switch:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 persistent-learning
```

Limiting MAC Address Movement

You can limit the number of times a MAC address can move from its original interface in 1 second—for example, set a MAC move limit of **5** with an action of **drop** if the limit is exceeded:

- On a single VLAN (here, the VLAN is **employee-vlan**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan mac-move-limit 5 action drop
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all mac-move-limit 5 action drop
```

You can also specify the actions **log** (do not drop the packet but generate an alarm, an SNMP trap, or a system log entry), **none** (no action), or **shutdown** (disable the interface or VLAN and generate an alarm) to occur if the MAC address moves more than the specified number of times in 1 second.

Configuring Trusted DHCP Servers on an Interface

Configure a trusted DHCP server on an interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/1 dhcp-trusted
```

Related Documentation

- [Configuring Port Security \(J-Web Procedure\)](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Monitoring Port Security on page 4847](#)
- [Port Security Overview on page 4676](#)
- [secure-access-port](#)
- [secure-access-port on page 4826](#)

Configuring MAC Limiting

To configure MAC limiting on a specific interface or on all interfaces:

1. To limit the number of dynamic MAC addresses, set a MAC limit of **5**.

The action is not specified, so the switch performs the default action **drop** if the limit is exceeded:

- On a single interface (here, the interface is **xe-0/0/1**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface xe-0/0/1 mac-limit (Access Port Security) 5
```

- On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all mac-limit 5
```



CAUTION: Do not set the MAC limit to 1. The first learned MAC address is often inserted into the forwarding database automatically. (For instance, the first MAC address inserted into the forwarding database for routed VLAN interfaces is the MAC address of the RVI. For Aggregated Ethernet bundles using LACP, the first MAC address inserted into the forwarding database in the forwarding table is the source address of the protocol packet.) The switch therefore fails to learn MAC addresses other than the automatic addresses when the MAC limit is set to 1, and this causes problems with MAC learning and forwarding.

2. To specify allowed MAC addresses:

- On a single interface (here, the interface is `xe-0/0/2`):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface xe-0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface xe-0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface xe-0/0/2 allowed-mac 00:05:85:3A:82:83
```

- On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all allowed-mac 00:05:85:3A:82:80
user@switch# set interface all allowed-mac 00:05:85:3A:82:81
user@switch# set interface all allowed-mac 00:05:85:3A:82:83
```

Related Documentation

- [Port Security Overview on page 4676](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- [Overview of Access Port Protection on page 4674](#)
- [Verifying That MAC Limiting Is Working Correctly on page 4850](#)
- [Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 4718](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 4715](#)
- [no-allowed-mac-log on page 4821](#)

Configuring MAC Move Limiting (CLI Procedure)

When MAC move limiting is configured, MAC address movements are tracked by the switch and, if a MAC address changes more than the configured number of times within 1 second, the changes to MAC addresses are dropped, logged, ignored, or the interface is shut down.



NOTE: Although you enable this feature on VLANs, the MAC move limitation pertains to the number of movements for each individual MAC address rather than the total number of MAC address moves in the VLAN. For example, if the MAC move limit is set to 1, the switch allows an unlimited number of MAC address movements within the VLAN as long as the same MAC address does not change more than once.

You configure MAC move limiting per VLAN, not per interface (port). In the default configuration, the number of MAC moves permitted is unlimited.

You can choose to have one of the following actions performed when the MAC move limit is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interfaces in the VLAN and generate a system log entry. If you have configured the switch with the **port-error-disable** statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

To configure a MAC move limit for MAC addresses within a specific VLAN or for MAC addresses within all VLANs, using the CLI:

- On a single VLAN: To limit the number of MAC address movements that can be made by an individual MAC address within the VLAN **employee-vlan**, set a MAC move limit of 5:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan mac-move-limit 5
```

The action is not specified, so the switch performs the default action **drop** if it tracks that an individual MAC address within the **employee-vlan** has moved more than 5 times within one second.

- On all VLANs: To limit the number of MAC movements that can be made by individual MAC addresses within all VLANs, set a MAC move limit of 5:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all mac-move-limit 5
```

The action is not specified, so the switch performs the default action **drop** if it tracks that an individual MAC address within any of the VLANs has moved more than 5 times within 1 second.

Related Documentation

- *Configuring MAC Move Limiting (J-Web Procedure)*
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 4853](#)
- [Monitoring Port Security on page 4847](#)
- *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
- *Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches*
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- `clear ethernet-switching port-error`
- [clear ethernet-switching port-error on page 4859](#)
- [port-error-disable on page 4823](#)
- [port-error-disable on page 4823](#)
- `secure-access-port`
- [secure-access-port on page 4826](#)

Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)

An Ethernet access interface might shut down or be disabled as a result of one of the following configurations:

- MAC limiting—**mac-limit** statement is configured with action **shutdown**.
- MAC move limiting—**mac-move-limit** statement is configured with action **shutdown**.
- Storm control—**storm-control** statement is configured with the action **shutdown**.

You can configure a device to automatically restore the disabled interfaces to service after a specified period of time. Autorecovery applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.

To configure autorecovery from the disabled state due to MAC limiting, MAC move limiting, or storm control shutdown actions:

```
[edit ethernet-switching-options]  
user@switch# set port-error-disable disable-timeout seconds
```



NOTE: You must specify the disable timeout value—there is no default disable timeout period. If you do not specify a timeout value, you must use the [clear ethernet-switching port-error](#) command to clear the errors and restore the interfaces to service.

Related Documentation

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
- [Configuring MAC Limiting on page 4758](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4760](#)
- [Understanding Storm Control on page 4694](#)

Configuring the none Action to Override a MAC Limit Applied to All Interfaces (CLI Procedure)

If you set a MAC limit in your port security settings to apply to all interfaces, you can override that setting for a particular interface by specifying the action **none**.

To use the **none** action to override a MAC limit setting:

1. Set the MAC limit for all interfaces—for example, a limit of **5** with action **drop**:

```
[edit ethernet-switching-options secure-access-port]  
user@switch# set interface all mac-limit (Access Port Security) 5 action drop
```

2. Change the action for one interface with this command.

```
[edit ethernet-switching-options secure-access-port]  
user@switch# set interface xe-0/0/2 mac-limit action none
```

Related Documentation

- [Configuring MAC Limiting on page 4758](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)

- [Verifying That MAC Limiting Is Working Correctly on page 4850](#)
- [Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 4718](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)

Configuring Static ARP Entries

You can create static ARP table entries, which are explicit mappings between IP addresses and MAC addresses.

- To configure a static ARP entry:

```
[edit interfaces interface-name unit logical-unit-number family inet address
address]
user@switch# set arp ip-address (mac | multicast-mac) mac-address
```

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, use the **multicast-mac** statement.

Specify the MAC address as 6 hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

Related Documentation

- [Understanding Static ARP Entries on page 4693](#)
- [arp on page 2072](#)

Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)

You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database. These bindings are labeled as “static” in the database, while those bindings that have been added through the process of DHCP snooping are labeled “dynamic.”

To configure a static IP address/MAC address binding in the DHCP snooping database (replace **ge-0/0/2**, **10.0.10.12**, **data-vlan**, and **00:05:85:3A:82:80** with values for your configuration):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 static-ip 10.0.10.12 vlan data-vlan mac 00:05:85:3A:82:80
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- [Verifying That DHCP Snooping Is Working Correctly on page 4849](#)
- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [secure-access-port](#)

- [secure-access-port on page 4826](#)

Enabling DHCP Snooping (CLI Procedure)

DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. It builds and maintains a database of valid IP-address/MAC-address (IP-MAC) bindings called the DHCP snooping database.



NOTE: If you configure DHCP snooping for all VLANs and you enable a different port security feature on a specific VLAN, you must also explicitly enable DHCP snooping on that VLAN. Otherwise, the default value of no DHCP snooping applies to that VLAN.

This topic describes:

- [Enabling DHCP Snooping on page 4765](#)
- [Applying CoS Forwarding Classes to Prioritize Snooped Packets on page 4765](#)

Enabling DHCP Snooping

You configure DHCP snooping per VLAN, not per interface (port). By default, DHCP snooping is disabled for all VLANs. You can enable DHCP snooping on all VLANs or on specific VLANs.

To enable DHCP snooping on a VLAN or all VLANs:

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan vlan-name examine-dhcp
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan all examine-dhcp
```



TIP: By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.



TIP: For private VLANs (PVLANS), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

Applying CoS Forwarding Classes to Prioritize Snooped Packets

On EX Series switches you might need to use class of service (CoS) to protect packets from critical applications from being dropped during periods of network congestion and delay and you might also need the port security features of DHCP snooping on the same ports through which those critical packets are entering and leaving.



NOTE: This is not supported on the QFX Series switch.

To apply CoS forwarding classes and queues to snooped packets:

1. Create a user-defined forwarding class to be used for prioritizing snooped packets:

```
[edit class-of-service]
user@switch# set forwarding-classes class class-name queue queue-number
```

2. Enable DHCP snooping on a specific VLAN or on all VLANs and apply the desired forwarding class on the snooped packets:

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan vlan-name examine-dhcp forwarding-class class-name
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access port]  
user@switch# set vlan all examine-dhcp forwarding-class class-name
```

Related Documentation

- [Enabling DHCP Snooping \(J-Web Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Verifying That DHCP Snooping Is Working Correctly on page 4849](#)
- [Monitoring Port Security on page 4847](#)
- [Understanding DHCP Snooping for Port Security on page 4678](#)
- [class-of-service on page 5831](#)
- [secure-access-port](#)
- [secure-access-port on page 4826](#)

Enabling Dynamic ARP Inspection (CLI Procedure)

Dynamic ARP inspection (DAI) protects switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

This topic describes:

- [Enabling DAI on page 4767](#)
- [Applying CoS Forwarding Classes to Prioritize Inspected Packets on page 4767](#)

Enabling DAI

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

To enable DAI on a VLAN or all VLANs:

- On a single VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan vlan-name arp-inspection
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all arp-inspection
```

Applying CoS Forwarding Classes to Prioritize Inspected Packets

You might need to use class of service (CoS) to protect packets from critical applications from being dropped during periods of network congestion and delay and you might also need the port security features of DHCP snooping on the same ports through which those critical packets are entering and leaving.

To apply CoS forwarding classes and queues to DAI packets:

- Create a user-defined forwarding class to be used for prioritizing DAI packets:

```
[edit class-of-service]
user@switch# set forwarding-classes class class-name queue queue-number
```

- Enable DAI on a specific VLAN or on all VLANs and apply the desired forwarding class on the DAI packets:

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan vlan-name arp-inspection forwarding-class class-name
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan all arp-inspection forwarding-class class-name
```

Related Documentation

- [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Verifying That DAI Is Working Correctly on page 4848](#)
- [Monitoring Port Security on page 4847](#)

- [Understanding DAI for Port Security on page 4686](#)
- [Understanding DAI for Port Security on page 4686](#)
- [class-of-service on page 5831](#)
- [secure-access-port](#)
- [secure-access-port on page 4826](#)

Enabling a Trusted DHCP Server (CLI Procedure)

You can configure any interface on a switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

You configure a trusted DHCP server on an interface, not on a VLAN. By default, all access interfaces are untrusted, and all trunk interfaces are trusted.

To configure a trusted interface for a DHCP server by using the CLI (here, the interface is **ge-0/0/8**):

```
[edit ethernet-switching-options secure-access port]  
user@switch# set interface ge-0/0/8 dhcp-trusted
```

Related Documentation

- [Enabling a Trusted DHCP Server \(J-Web Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly on page 4854](#)
- [Monitoring Port Security on page 4847](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4690](#)
- [secure-access-port](#)
- [secure-access-port on page 4826](#)

Enabling a Trusted Port for DHCP

By default, all access ports are untrusted and all trunk ports are trusted with regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. Untrusted ports drop traffic from DHCP servers to prevent unauthorized servers from providing any configuration information to clients.

If you attach a DHCP server to an access port, you must configure it as trusted. You configure a trusted DHCP server on an interface, not on a VLAN.



NOTE: Before you attach a DHCP server to a trusted access port, ensure that the server is physically secure—that is, that access to the server is monitored and controlled.

- To configure a trusted interface for a DHCP server by using the CLI (here, the interface is **xe-0/0/8**):

```
[edit ethernet-switching-options secure-access port]
user@switch# set interface xe-0/0/8 dhcp-trusted
```

Related Documentation

- [Enabling a Trusted DHCP Server \(J-Web Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly on page 4854](#)
- [Monitoring Port Security on page 4847](#)
- [Understanding Trusted and Untrusted Ports on page 4690](#)

Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's replies to the clients. This topic describes this configuration.
- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as routed VLAN interfaces, or RVIs. The switch relays the clients' requests to the server and then forwards the server's replies to the clients. This configuration is described in [“Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\)”](#) on page 4773.

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.



.....

NOTE: Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

.....

- Configure a VLAN on the switch and associate the interfaces on which the clients and the server connect to the switch with that VLAN.

To configure DHCP option 82:



NOTE: Replace values displayed in *italics* with values for your configuration.

1. Specify DHCP option 82 for all VLANs associated with the switch or for a specified VLAN. (You can also configure the feature for a VLAN range.)

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all dhcp-option82
```

The remaining steps are optional.

2. To configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id prefix hostname
```

3. To specify that the circuit ID suboption value should contain the interface description rather than the interface name (the default):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id use-interface-description
```

4. To specify that the circuit ID suboption value should contain the VLAN ID rather than the VLAN name (the default):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id use-vlan-id
```

5. To specify that the remote ID suboption be included in the DHCP option 82 information:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id
```

6. To configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id prefix mac
```

7. To specify that the prefix for the remote ID suboption be the hostname of the switch rather than the MAC address of the switch (the default):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id prefix hostname
```

8. To specify that the remote ID suboption value should contain the interface description:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id use-interface-description
```

9. To specify that the remote ID suboption value should contain a character string:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id use-string mystring
```

10. To configure a vendor ID suboption and use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit ethernet-switching-options secure-access-port]
```

```
user@switch# set vlan employee dhcp-option82 vendor-id
```

11. To specify that the vendor ID suboption value should contain a character string value that you specify rather than **Juniper** (the default):

```
[edit ethernet-switching-options secure-access-port]
```

```
user@switch# set vlan employee dhcp-option82 vendor-id mystring
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740](#)
- *secure-access-port*
- [secure-access-port on page 4826](#)
- *Understanding DHCP Option 82 for Port Security on EX Series Switches*
- [Understanding DHCP Option 82 for Port Security on page 4691](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)

You can use DHCP option 82, also known as the DHCP relay agent information option, to help switches against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as routed VLAN interfaces, or RVIs. The switch relays the clients' requests to the server and then forwards the server's replies to the clients. This topic describes this configuration.
- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's replies to the clients. This configuration is described in [“Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)” on page 4770](#).

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.



NOTE: Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configure the VLAN on the switch and associate the interfaces on which the clients connect to the switch with that VLAN.
- Configure the routed VLAN interface (RVI) to allow the switch to relay packets to the server and receive packets from the server. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or [“Configuring Routed VLAN Interfaces” on page 2046](#) for the QFX Series.
- Configure the switch as a BOOTP relay agent. See *DHCP/BOOTP Relay for Switches Overview*.

To configure DHCP option 82:



NOTE: Replace values displayed in *italics* with values for your configuration.

1. Specify DHCP option 82 for the BOOTP server:

- On all interfaces that connect to the server:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82
```

- On a specific interface that connects to the server:

```
[edit forwarding-options helpers bootp]
user@switch# set interface ge-0/0/10 dhcp-option82
```

The remaining steps are optional. They show configurations for all interfaces; include the specific interface designation to configure any of the following options on a specific interface:

2. To configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 circuit-id prefix hostname
```

3. To specify that the circuit ID suboption value should contain the interface description rather than the interface name (the default):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 circuit-id use-interface-description
```

4. To specify that the circuit ID suboption value should contain the VLAN ID rather than the VLAN name (the default):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 circuit-id use-vlan-id
```

5. To specify that the remote ID suboption be included in the DHCP option 82 information:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id
```

6. To configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id prefix mac
```

7. To specify that the prefix for the remote ID suboption be the hostname of the switch rather than the MAC address of the switch (the default):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id prefix hostname
```

8. To specify that the remote ID suboption value should contain the interface description:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id use-interface-description
```

9. To specify that the remote ID suboption value should contain a character string:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id use-string mystring
```

10. To configure a vendor ID suboption and use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id
```

11. To specify that the vendor ID suboption value contains a character string value that you specify rather than **Juniper** (the default):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id mystring
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744](#)
- *[edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches*
- *Understanding DHCP Option 82 for Port Security on EX Series Switches*
- [Understanding DHCP Option 82 for Port Security on page 4691](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

Configuration Statements for Firewall Filters

- [family on page 4776](#)
- [filter on page 4777](#)
- [filter \(Layer 2 and Layer 3 Interfaces\) on page 4778](#)
- [filter \(VLANs\) on page 4779](#)
- [firewall on page 4780](#)
- [from on page 4781](#)
- [interface-specific on page 4782](#)
- [term on page 4782](#)
- [then \(Filters\) on page 4783](#)

family

Syntax family *family-name* {
 filter *filter-name* {
 interface-specific;
 term *term-name* {
 from {
 match-conditions;
 }
 then {
 action;
 action-modifiers;
 }
 }
 }
 }

Hierarchy Level [edit [firewall](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure a firewall filter for IP version 4 or IP version 6.

Options *family-name*—Type of addressing protocol:

- **ethernet-switching**—Filter Layer 2 Ethernet packets and Layer 3 (IP) packets (allows some Layer 3 filtering).
- **inet**—Filter Layer 3 IPv4 packets (provides additional Layer 3 filter options).
- **inet6**—Filter Layer 3 IPv6 packets (provides additional Layer 3 filter options).

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Firewall Filter Match Conditions and Actions on page 4644](#)
- [Configuring Firewall Filters on page 4747](#)
- [Overview of Firewall Filters on page 4635](#)

filter

Syntax	<pre> filter <i>filter-name</i> { <i>interface-specific</i>; term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } } } </pre>
Hierarchy Level	[edit firewall family <i>family-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure firewall filters.
Options	<p><i>filter-name</i>—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 4644 • Configuring Firewall Filters on page 4747 • Overview of Firewall Filters on page 4635

filter (Layer 2 and Layer 3 Interfaces)

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply a firewall filter to traffic transiting a port or Layer 3 interface.
Default	All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.
Options	<p><i>filter-name</i>—Name of a firewall filter defined at the [edit firewall family <i>family-name</i> filter] hierarchy level.</p> <p><i>input</i>—Apply a firewall filter to traffic entering the port or Layer 3 interface.</p> <p><i>output</i>—Apply a firewall filter to traffic exiting the port or Layer 3 interface.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Gigabit Ethernet Interfaces (CLI Procedure)• Configuring Firewall Filters on page 4747• Overview of Firewall Filters on page 4635

filter (VLANs)

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply a firewall filter to traffic ingressing or egressing a VLAN.
Default	All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.
Options	<p><i>filter-name</i>—Name of a firewall filter defined at the [edit firewall family <i>family-name</i> filter] hierarchy level.</p> <p>input—Apply a firewall filter to VLAN ingress traffic.</p> <p>output—Apply a firewall filter to VLAN egress traffic.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Firewall Filters on page 4747 • Overview of Firewall Filters on page 4635

firewall

```
Syntax  firewall {
        family family-name {
            filter filter-name {
                interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
            }
        }
        policer policer-name {
            filter-specific;
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
        three-color-policer policer-name {
            action {
                loss-priority high then discard;
            }
            single-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                excess-burst-size bytes;
            }
            two-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                peak-information-rate bps;
                peak-burst-size bytes;
            }
        }
    }
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 4644 • Configuring Firewall Filters on page 4747 • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Overview of Firewall Filters on page 4635

from

Syntax	<pre>from { match-conditions; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Match packet fields to values specified in a match condition. If the from statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the then statement are implemented.
Options	match-conditions —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the then statement to be implemented.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 4644 • Configuring Firewall Filters on page 4747 • Understanding Firewall Filter Match Conditions on page 4641

interface-specific

Syntax	interface-specific;
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure separate counters for each interface to which a filter is applied.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions on page 4644• Configuring Firewall Filters on page 4747• Overview of Firewall Filters on page 4635

term

Syntax	<pre>term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define a firewall filter term.
Options	<p><i>term-name</i>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions on page 4644• Configuring Firewall Filters on page 4747• Overview of Firewall Filters on page 4635

then (Filters)

Syntax	<pre>then { action; action-modifiers; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a firewall filter action.
Options	<p>action—Actions to accept, discard, or forward packets that match all conditions specified in a filter term.</p> <p>action-modifiers—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 4644 • Configuring Firewall Filters on page 4747 • Understanding Firewall Filter Match Conditions on page 4641

Configuration Statements for Policers

- [action on page 4784](#)
- [bandwidth-limit on page 4784](#)
- [burst-size-limit on page 4785](#)
- [color-aware on page 4786](#)
- [color-blind on page 4787](#)
- [committed-burst-size on page 4788](#)
- [committed-information-rate on page 4789](#)
- [excess-burst-size on page 4790](#)
- [filter-specific on page 4791](#)
- [firewall on page 4792](#)
- [if-exceeding on page 4793](#)
- [loss-priority high then discard \(Three-Color Policer\) on page 4794](#)
- [peak-burst-size on page 4795](#)
- [peak-information-rate on page 4796](#)
- [policer on page 4797](#)

- [single-rate on page 4798](#)
- [then \(Policers\) on page 4799](#)
- [three-color-policer on page 4800](#)
- [two-rate on page 4801](#)

action

Syntax	<code>action { loss-priority high then discard; }</code>
Hierarchy Level	[edit firewall three-color-policer name]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Discard traffic on a logical interface using tricolor marking policing.
Options	The statements are explained separately.
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.

bandwidth-limit

Syntax	<code>bandwidth-limit <i>bps</i>;</code>
Hierarchy Level	[edit firewall policer policer-name if-exceeding]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the traffic rate in bits per second.
Options	<code>bps</code> —Traffic rate in bits per second. Specify <code>bps</code> as a decimal value or as a decimal number followed by one of the abbreviation <code>k</code> (1000), <code>m</code> (1,000,000), or <code>g</code> (1,000,000,000). Range: 32000 bps (32 Kbps) through 10,000,000,000 bps (10 Gbps)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

burst-size-limit

Syntax	<code>burst-size-limit bytes;</code>
Hierarchy Level	[edit <code>firewall policer policer-name if-exceeding</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the maximum allowed burst size to control the amount of traffic bursting.
Options	bytes —Decimal value or a decimal number followed by k (thousand), m (million), or g (giga). Range: 1 through 2,147,450,880 bytes (2147 MB)
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664


color-aware

Syntax	color-aware;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the way preclassified packets are metered. In color-aware mode, the switch can assign a higher packet-loss priority, but cannot assign a lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded its committed information rate (CIR). The switch cannot change the PLP to low even if the packet conforms to the configured CIR of the appropriate interface. On the other hand, if an upstream device assigns low PLP to a packet but the packet exceeds the CIR and committed burst size (CBS) of the switch interface, the switch can increase the PLP to medium-high.
Default	If you omit the color-aware statement, the default behavior is color-aware mode.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of Policers on page 4664• Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 4670• Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 4672• color-blind on page 4787


color-blind

Syntax	color-blind;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the way preclassified packets are metered. In color-blind mode, the switch ignores any preclassification of packets and can assign a higher or lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded the CIR on the upstream device. The switch can change the PLP to low if the packet conforms to the CIR of the appropriate interface.
Default	If you omit the color-blind statement, the default behavior is color-aware mode.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Overview of Policers on page 4664 • Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 4669 • Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 4671 • Configuring Color-Blind Egress Policers for Medium-Low PLP on page 4752 • color-aware on page 4786


committed-burst-size

Syntax	<code>committed-burst-size bytes;</code>
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with low packet loss priority (green).
<hr/>	
<div> NOTE: When you include the <code>committed-burst-size</code> statement in the configuration, you must also include the <code>committed-information-rate</code> statement at the same hierarchy level.</div> <hr/>	
Options	bytes —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

committed-information-rate

Syntax	<code>committed-information-rate <i>bits-per-second</i>;</code>
Hierarchy Level	[edit <code>firewall three-color-policer <i>policer-name</i> single-rate</code>], [edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).
<div>  <p>NOTE: When you include the <code>committed-information-rate</code> statement in the configuration, you must also include the <code>committed-burst-size</code> statement at the same hierarchy level.</p> </div>	
Options	<p><i>bits-per-second</i>—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: 32,000 bps through 10,000,000,000 bps (10 gbps)</p>
Required Privilege Level	<p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Overview of Policers on page 4664

excess-burst-size

Syntax	<code>excess-burst-size bytes;</code>
Hierarchy Level	[edit <code>firewall three-color-policer policer-name</code> single-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).
<div> NOTE: When you include the <code>excess-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>committed-information-rate</code> statements at the same hierarchy level.</div>	
Options	bytes —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

filter-specific

Syntax	filter-specific;
Hierarchy Level	[edit firewall policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. If you use a filter-specific policer in multiple terms, both of the following are true:</p> <ul style="list-style-type: none">• Traffic is policed at the aggregate rate. For example, if you create a policer that has a bandwidth limit of 100 Mbps and use the policer in two terms, the total allowed bandwidth for both terms is 100 Mbps—not 100 Mbps for each term.• The implicit counter counts all the packets are that matched by any of the terms. For example, if you reference the same filter-specific policer in term1 and term2, and term1 matches 1000 packets and term2 matches 500 packets, the implicit counter shows 1500 matches for the policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

firewall

```
Syntax  firewall {  
        family family-name {  
            filter filter-name {  
                interface-specific;  
                term term-name {  
                    from {  
                        match-conditions;  
                    }  
                    then {  
                        action;  
                        action-modifiers;  
                    }  
                }  
            }  
        }  
        policer policer-name {  
            filter-specific;  
            if-exceeding {  
                bandwidth-limit bps;  
                burst-size-limit bytes;  
            }  
            then {  
                policer-action;  
            }  
        }  
        three-color-policer policer-name {  
            action {  
                loss-priority high then discard;  
            }  
            single-rate {  
                (color-aware | color-blind);  
                committed-information-rate bps;  
                committed-burst-size bytes;  
                excess-burst-size bytes;  
            }  
            two-rate {  
                (color-aware | color-blind);  
                committed-information-rate bps;  
                committed-burst-size bytes;  
                peak-information-rate bps;  
                peak-burst-size bytes;  
            }  
        }  
    }
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 4644 • Configuring Firewall Filters on page 4747 • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Overview of Firewall Filters on page 4635


if-exceeding

Syntax	<pre>if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; }</pre>
Hierarchy Level	[edit firewall policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure policer rate limits.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Overview of Policers on page 4664


loss-priority high then discard (Three-Color Policer)

Syntax	loss-priority high then discard;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> action]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>For packets with high loss priority, discard the packets. The loss priority setting is not configurable. Include this statement if you do not want the switch to forward packets that have high packet-loss priority.</p> <p>For single-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

peak-burst-size

Syntax	<code>peak-burst-size bytes;</code>
Hierarchy Level	[edit <code>firewall three-color-policer policer-name two-rate</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).
<div>  <p>NOTE: When you include the <code>peak-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>peak-information-rate</code> statements at the same hierarchy level.</p> </div>	
Options	<p>bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: 1500 bytes through 100,000,000,000 bytes (100 GB)</p>
Required Privilege Level	<p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4793 • Overview of Policers on page 4664

peak-information-rate

Syntax	<code>peak-information-rate <i>bits-per-second</i>;</code>
Hierarchy Level	[edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the maximum achievable rate. Packets that exceed the committed information rate (CIR) but are below the peak information rate (PIR) are marked with medium-high packet loss priority (yellow). Packets that exceed the PIR are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR.
<div> NOTE: When you include the <code>peak-information-rate</code> statement in the configuration, you must also include the <code>committed-information-rate</code> and <code>peak-burst-size</code> statements at the same hierarchy level.</div>	
Options	<i>bits-per-second</i> —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 32,000 bps through 10,000,000,000 bps (10 gbps)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664


policer

Syntax	<pre> policer <i>policer-name</i> { filter-specific; if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { <i>policer-action</i>; } } </pre>
Hierarchy Level	[edit firewall]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure policer rate limits and actions. To activate a policer, you must include the policer action modifier in the then statement in a firewall filter term.</p> <p>Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or across filters—the policer's implicit counter is used to count packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these approaches:</p> <ul style="list-style-type: none"> • Configure a unique policer for each term. • Configure only one policer, but use a unique, explicit counter in each term.
Options	<p><i>policer-name</i>—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Configuring Firewall Filters on page 4747 • Overview of Policers on page 4664

single-rate

Syntax	<pre>single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
Options	<i>policer-name</i> —Name of the three-color policer. Use this name when you apply the policer to an interface.
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

then (Policers)

Syntax	then { <i>policer-action</i> ; }
Hierarchy Level	[edit firewall policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a policer action.
Options	<i>policer-action</i> —Allowed policer actions are discard , loss-priority high , and loss-priority low . discard causes the system to drop traffic that exceeds the rate limits defined by the policer. Use loss-priority high to allow the system to forward matching traffic in some cases.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: If you specify a policer in an egress firewall filter, the only supported action is discard.</p> </div> </div>	
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753 • Configuring Firewall Filters on page 4747 • Overview of Policers on page 4664

three-color-policer

Syntax	<pre>three-color-policer <i>policer-name</i> { action { loss-priority high then discard; } single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; } two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; } }</pre>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a three-color policer.
Options	<p><i>policer-name</i>—Name of the three-color policer. Use this name when you apply the policer to an interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753• Overview of Policers on page 4664

two-rate

Syntax	<pre>two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>

Configuration Statements for Port Security

- [allowed-mac](#) on page 4803
- [arp-inspection](#) on page 4804
- [circuit-id](#) on page 4805
- [dhcp-trusted](#) on page 4806
- [dhcp-option82](#) on page 4807
- [dhcp-snooping-file](#) on page 4808
- [dhcp-trusted](#) on page 4808
- [disable-timeout \(Port Error Disable\)](#) on page 4809
- [ethernet-switching-options](#) on page 4810
- [examine-dhcp](#) on page 4812
- [examine-fip](#) on page 4813
- [fc-map](#) on page 4814
- [fcoe-trusted](#) on page 4815

- forwarding-class (for DHCP Snooping or DAI Packets) on page 4816
- interface (Secure Access Port) on page 4817
- location on page 4818
- mac on page 4818
- mac-limit on page 4819
- mac-move-limit on page 4820
- no-allowed-mac-log on page 4821
- no-dhcp-trusted on page 4821
- no-gratuitous-arp-request on page 4822
- persistent-learning on page 4822
- port-error-disable on page 4823
- prefix (Remote ID for Option 82) on page 4824
- remote-id on page 4825
- secure-access-port on page 4826
- static-ip on page 4827
- timeout (DHCP Snooping) on page 4828
- use-interface-description on page 4829
- use-string on page 4830
- use-vlan-id on page 4831
- vendor-id on page 4832
- vlan (Secure Access Port) on page 4833
- vlan (Static IP) on page 4834
- write-interval on page 4835

allowed-mac


Syntax	<code>allowed-mac <i>mac-address-list</i></code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify particular MAC addresses to be added to the MAC address cache.



NOTE: Although this configuration restricts the addresses that can be added to the MAC address cache, it does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the list of allowed MAC addresses. Control packets do not undergo the MAC address check, and they are therefore included in the statistics of packets received, though, they are not forwarded to another destination.

Default	Allowed MAC addresses take precedence over dynamic MAC values. For example, if the mac-limit statement is set to four and three allowed MACs are configured, only one dynamic MAC can be learned on that interface.
Options	<i>mac-address-list</i> —One or more MAC addresses configured as allowed MAC addresses for a specified interface or all interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Configuring MAC Limiting on page 4758 • Configuring MAC Move Limiting (CLI Procedure) on page 4760 • mac-limit on page 4819 • no-allowed-mac-log on page 4821

arp-inspection

Syntax	(arp-inspection no-arp-inspection) { forwarding-class (for DHCP Snooping or DAI Packets) <i>class-name</i> ; }
Hierarchy Level	[edit] ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Perform dynamic ARP inspection on all VLANs or on the specified VLAN. <ul style="list-style-type: none">• arp-inspection—Enable ARP inspection. <div> NOTE: When ARP inspection is enabled, the switch logs ARP request packets that it rejects.</div> <ul style="list-style-type: none">• no-arp-inspection—Disable ARP inspection.
Default	Disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Dynamic ARP Inspection (CLI Procedure) on page 4766• Example: Configuring Basic Port Security Features on page 4706Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725• Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732

circuit-id

Syntax	<pre>circuit-id { prefix hostname; use-interface-description; use-vlan-id; }</pre>
Hierarchy Level	<p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82]</p> <p>[edit forwarding-options helpers bootp dhcp-option82]</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the circuit-id suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (interface and/or VLAN) on which the DHCP request arrived.</p> <p>The format of the circuit-id information for Gigabit Ethernet interfaces that use VLANs is <i>interface-name:vlan-name</i> . On a Layer 3 interface, the format is just <i>interface-name</i> .</p> <p>The remaining statements are explained separately.</p>
Default	<p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i> .</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740 • Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770 • Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773 • [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches • RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046.

dhcp-trusted

Syntax	(dhcp-trusted no-dhcp-trusted);
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (Access Port Security) (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Allow or deny DHCP responses from the specified interfaces (ports) or all interfaces.</p> <ul style="list-style-type: none">• dhcp-trusted—Allow DHCP responses.• no-dhcp-trusted—Deny DHCP responses.
Default	Trusted for trunk ports, untrusted for access ports.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Overview of Access Port Protection on page 4674• Enabling a Trusted Port for DHCP on page 4769

dhcp-option82

Syntax	<pre> dhcp-option82 { circuit-id { prefix hostname; use-interface-description; use-vlan-id; } remote-id { prefix hostname mac none; use-interface-description; use-string <i>string</i>; } vendor-id <<i>string</i>>; } </pre>
Hierarchy Level	<p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)]</p> <p>[edit forwarding-options helpers bootp]</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>When the switch receives a DHCP request from a DHCP client connected on one of the switch's interfaces, have the switch insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header before it forwards or relays the request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from, in formulating the reply; the server does not, however, make any changes to the option 82 information in the packet header. The switch receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.</p> <p>The remaining statements are explained separately.</p>
Default	<p>Insertion of DHCP option 82 information is not enabled.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740 • Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770 • Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773 • [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches

- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.


dhcp-snooping-file

Syntax	<code>dhcp-snooping-file { location local_pathname remote_URL; timeout seconds; write-interval seconds; }</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify a local pathname or remote URL for the DHCP snooping database file to maintain persistence of IP-MAC bindings.</p> <p>The remaining statements are explained separately.</p>
Default	The IP-MAC bindings in the DHCP snooping database file are not persistent. If the switch is rebooted, the bindings are lost.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding DHCP Snooping for Port Security on page 4678

dhcp-trusted

Syntax	<code>(dhcp-trusted no-dhcp-trusted);</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (Access Port Security) (all interface-name)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Allow or deny DHCP responses from the specified interfaces (ports) or all interfaces.</p> <ul style="list-style-type: none">• dhcp-trusted—Allow DHCP responses.• no-dhcp-trusted—Deny DHCP responses.
Default	Trusted for trunk ports, untrusted for access ports.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Overview of Access Port Protection on page 4674• Enabling a Trusted Port for DHCP on page 4769

disable-timeout (Port Error Disable)

Syntax	<code>disable-timeout <i>timeout</i>;</code>
Hierarchy Level	[edit ethernet-switching-options port-error-disable]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify how long Ethernet switching interfaces remain in a disabled state due to MAC limiting, MAC move limiting, or storm control errors.
<div>  <p>NOTE: If you modify an existing timeout value, the new timeout value does not affect currently disabled interfaces are configured for automatic recovery. The new timeout value applies only to subsequent port errors. Run the clear ethernet-switching port-error command to restore currently disabled interfaces.</p> </div>	
Default	The disable timeout statement is not enabled.
Options	<p><i>timeout</i>—Time, in seconds, that an interface remains disabled. The disabled interface automatically returns to service when the specified time expires.</p> <p>Range: 10 through 3600 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Understanding Storm Control on page 4694 • Example: Configuring Storm Control on page 4713 • Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure) on page 4762 • action-shutdown on page 4836

ethernet-switching-options

```
Syntax ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
      output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds {
  }
  port-error-disable {
    disable-timeout timeout;
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      fcoe-trusted;
      mac-limit limit action action;
      no-allowed-mac-log;
    }
    vlan (all | vlan-name) {
      (arp-inspection | no-arp-inspection) [
        forwarding-class (for DHCP Snooping or DAI Packets) class-name;
      ]
    }
  }
}
```

```

dhcp-option82 {
  circuit-id {
    prefix (Circuit ID for Option 82) hostname;
    use-interface-description;
    use-vlan-id;
  }
  remote-id {
    prefix (Remote ID for Option 82) hostname | mac | none;
    use-interface-description;
    use-string string;
  }
  vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp) {
  forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
  examine-vn2vn {
    beacon-period milliseconds;
  }
  fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Ethernet switching options.


The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
 - [Overview of Access Port Protection on page 4674](#)
 - [Understanding Storm Control on page 4694](#)


examine-dhcp

Syntax	(examine-dhcp no-examine-dhcp);
Hierarchy Level	[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series
Description	<p>Enable DHCP snooping on all VLANs or on the specified VLAN.</p> <ul style="list-style-type: none">• examine-dhcp—Enable DHCP snooping. <div> NOTE: When DHCP snooping is enabled, the switch logs DHCPDISCOVER packets that it rejects.</div> <ul style="list-style-type: none">• no-examine-dhcp—Disable DHCP snooping.
Default	Disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Basic Port Security Features on page 4706• Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725• Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732• Enabling DHCP Snooping (CLI Procedure) on page 4764

examine-fip

Syntax	<pre>examine-fip { examine-vn2vn { beacon-period milliseconds; } fc-map fc-map-value; }</pre>
Hierarchy Level	[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)]
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement examine-vn2vn introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	<p>Enable FIP snooping on a specified VLAN. Ensure that the VLAN is a dedicated FCoE VLAN that transports only FCoE traffic.</p> <p>(QFX Series only) Enable VN_Port to VN_Port (VN2VN_Port) FIP snooping on a specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only VN2VN_Port traffic. One FCoE VLAN cannot support both VN_Port to VF_Port (VN2VF_Port) FIP snooping and VN2VN_Port FIP snooping. Configure separate, dedicated FCoE VLANs for VN2VN_Port FIP snooping and VN2VN_Port FIP snooping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>vlan</i> <i>Example: Configuring an FCoE Transit Switch</i> Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199

fc-map

Syntax	<code>fc-map <i>fc-map-value</i>;</code>
Hierarchy Level	<p>[edit ethernet-switching options secure-access-port vlan (all <i>vlan-name</i>) examine-fip]</p> <p>For the QFX Series only:</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).</p> <p>You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN_Port to VF_Port (VN2VF_Port) FIP snooping (0x0EFC00) than for VN_Port to VN_Port (VN2VN_Port) FIP snooping.</p> <p>The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN_Port MAC address.</p> <p>When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.</p>
	<p> NOTE: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.</p>
Options	<p><i>fc-map-value</i>—FC-MAP value, hexadecimal value preceded by “0x”.</p> <p>Range: 0x0EFC00 through 0x0EFCFF</p> <p>Default: 0x0EFC00 for VN2VF_Port FIP snooping 0x0EFD00 for VN2VN_Port FIP snooping</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • examine-fip on page 4813 • show fip snooping on page 5377 • <i>Example: Configuring an FCoE Transit Switch</i>

- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)

fcoe-trusted

Syntax	fcoe-trusted;
Hierarchy Level	<p>[edit ethernet-switching-options secure-access-port interface <i>interface-name</i>]</p> <p>For the QFX Series only:</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>(QFX Series only) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the fcoe-trusted configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show fip snooping on page 5377 • Example: Configuring an FCoE Transit Switch • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199

forwarding-class (for DHCP Snooping or DAI Packets)

Syntax	forwarding-class class <i>class-name</i> ;
Hierarchy Level	[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) (examine-dhcp arp-inspection)]
Release Information	Statement introduced in Junos OS Release 11.2 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Assign a user-defined or a predefined forwarding class to the packets that have been checked for DHCP snooping or dynamic ARP inspection (DAI).



NOTE: To assign a user-defined class, you must first configure the user-defined class by using the *forwarding-classes* configuration statement at the [edit *class-of-service*] hierarchy level.

Default	Disabled.
Options	<i>class-name</i> —Name of the forwarding class. The forwarding class can be one of the predefined forwarding classes (best-effort, assured-forwarding, expedited-forwarding, network-control) or it can be a user-defined forwarding class.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic</i>• <i>Understanding Junos OS CoS Components for EX Series Switches</i>• Understanding DHCP Snooping for Port Security on page 4678• Understanding DAI for Port Security on page 4686

interface (Secure Access Port)

Syntax	<pre> interface (all <i>interface-name</i>) { allowed-mac <i>mac-address-list</i>; (dhcp-trusted no-dhcp-trusted); mac-limit <i>limit</i> action <i>action</i>; no-allowed-mac-log; static-ip <i>ip-address</i> { vlan <i>vlan-name</i>; mac <i>mac-address</i>; } } </pre>
Hierarchy Level	[edit ethernet-switching-options secure-access-port]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Apply port security features to all interfaces or to the specified interface.</p> <p>The remaining statements are explained separately.</p>
Options	<p>all—Apply port security features to all interfaces. Does not apply to QFabric systems.</p> <p><i>interface-name</i>—Apply port security features to the specified interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of Access Port Protection on page 4674 • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Understanding Trusted and Untrusted Ports on page 4690 • Configuring MAC Limiting on page 4758 • Enabling a Trusted Port for DHCP on page 4769

location

Syntax	<code>location <i>local_pathname</i> <i>remote_URL</i>;</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port dhcp-snooping-file]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify either a local pathname or a remote URL as the location in which to store the DHCP snooping database.
Options	<p><i>local_pathname</i> <i>remote_URL</i> —Location for storing the DHCP snooping database.</p> <ul style="list-style-type: none">• <i>local_pathname</i> —Use <i>/path</i> to store the database on a local switch.• <i>remote_URL</i> —Use <code>ftp://ip-address</code> or <code>ftp://hostname/path</code> to store the database at a remote location.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

mac

Syntax	<code>mac <i>mac-address</i>;</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (Access Port Security) (all <i>interface-name</i>) static-ip <i>ip-address</i> vlan (DHCP Bindings on Access Ports) <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 on the QFX Series switches.
Description	Specify a media access control (MAC) address (hardware address) for the specified static IP address.
Options	<i>mac-address</i> —Value in hexadecimal format.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

mac-limit

Syntax	<code>mac-limit <i>limit</i> { <action <i>action</i>>; }</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the number of MAC addresses that can be dynamically added to the MAC address cache for this access interface (port) and the action to be taken if the limit is exceeded.
Default	The default action is drop .
Options	<p><i>limit</i>—Maximum number of MAC addresses.</p> <p><i>action action</i>—(Optional) Action to take when the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> • drop—Drop the packet and generate a system log entry. This is the default. • log—Do not drop the packet but generate a system log entry. • none—No action. • shutdown—Disable the interface and generate an alarm. If you configure the switch with the port-error-disable statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this statement is not configured, you can bring up the disabled interfaces by running the clear ethernet-switching port-error command.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Configuring MAC Limiting on page 4758 • allowed-mac on page 4803

mac-move-limit

Syntax	<code>mac-move-limit <i>limit</i> <fabric-limit <i>limit</i>> action <i>action</i>;</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port (all <i>vlan-name</i>)]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the number of times a MAC address can move to a new interface (port) in 1 second and the action to be taken by the switch if the MAC address move limit is exceeded.
Default	The default move limit is unlimited. The default action is drop .
Options	<p>fabric-limit—Specify the maximum number of moves in a QFabric system. If you do not specify a fabric limit, the value for mac-move-limit applies to the QFabric system.</p> <p>limit—Maximum number of moves to a new interface per second.</p> <p>action <i>action</i>—(Optional) Action to take when the MAC address move limit is reached:</p> <ul style="list-style-type: none">• drop—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.• log—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.• none—No action.• shutdown—Disable the interface and generate an alarm. If you have configured the switch with the port-error-disable statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the clear-ethernet-switch-port command.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• mac-limit on page 4819• Example: Configuring Basic Port Security Features on page 4706• Configuring MAC Move Limiting (CLI Procedure) on page 4760• Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)

no-allowed-mac-log

Syntax	no-allowed-mac-log;
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify that the switch should not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for allowed MAC addresses.
Default	The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular allowed (specific) MAC addresses.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Configuring MAC Limiting on page 4758 • allowed-mac on page 4803 • mac-limit on page 4819

no-dhcp-trusted

Syntax	(dhcp-trusted no-dhcp-trusted);
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (Access Port Security) (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Port security features, such as DHCP snooping and dynamic ARP inspection inspect packets only on untrusted interfaces.</p> <p>Allow or deny DHCP responses from the specified interfaces (ports) or all interfaces.</p> <ul style="list-style-type: none"> • dhcp-trusted—Allow DHCP responses. • no-dhcp-trusted—Deny DHCP responses.
Default	Trusted for trunk ports, untrusted for access ports.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Overview of Access Port Protection on page 4674 • Enabling a Trusted Port for DHCP on page 4769


no-gratuitous-arp-request

Syntax	no-gratuitous-arp-request;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-range</i> <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs).
Default	Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routed VLAN Interfaces on page 2046

persistent-learning

Syntax	persistent-learning;
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.4 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Basic Port Security Features on page 4706• Configuring Persistent MAC Learning (CLI Procedure)

port-error-disable

Syntax	<code>port-error-disable { disable-timeout <i>timeout</i>; }</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 11.1 on the QFX Series.
Description	Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and storm control, and allow the interface to recover automatically from the error condition after a specified period of time:
	<div>  <p>NOTE: The <code>port-error-disable</code> configuration does not apply to preexisting error conditions. It affects only error conditions that are detected after you enable and commit the <code>port-error-disable</code> statement. To clear a preexisting error condition and restore the interface to service, use the <code>clear ethernet-switching port-error</code> command.</p> </div>
	<ul style="list-style-type: none"> If you enable the <code>mac-limit</code> statement with the <code>shutdown</code> option and also enable the <code>port-error-disable</code> statement, the switch disables (rather than shuts down) the interface when the MAC address limit is reached. If you have enabled the <code>mac-move-limit</code> statement with the <code>shutdown</code> option and you enable the <code>port-error-disable</code> statement, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached. If you enable the <code>storm-control</code> statement with the <code>action-shutdown</code> option and you also enable <code>port-error-disable</code>, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceed the specified levels.
Default	Not enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 Understanding Storm Control on page 4694 Example: Configuring Storm Control on page 4713 Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure) on page 4762 action-shutdown on page 4836 disable-timeout on page 4809

- [clear ethernet-switching port-error on page 4859](#)

prefix (Remote ID for Option 82)

Syntax	prefix hostname mac none;
Hierarchy Level	[edit ethernet-switching-options secure-access-port vlan (all vlan-name) dhcp-option82 remote-id] [edit forwarding-options helpers bootp dhcp-option82 remote-id] [edit forwarding-options helpers bootp interface interface-name dhcp-option82 remote-id]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an optional prefix for the remote ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards the request or relays the request to a DHCP server.
Default	If prefix is not explicitly specified, no prefix is appended to the remote ID.
Options	hostname —Name of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server. mac —MAC address of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server. none —No prefix is applied to the remote ID.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740• Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744• Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770• Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773• [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches• RFC 3046, DHCP Relay Agent Information Option, at http://tools.ietf.org/html/rfc3046.

remote-id

Syntax	<pre>remote-id { prefix hostname mac none; use-interface-description; use-string string; }</pre>
Hierarchy Level	<p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82]</p> <p>[edit forwarding-options helpers bootp dhcp-option82]</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Insert the remote-id suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.</p> <p>The remaining statements are explained separately.</p>
Default	<p>If remote-id is not explicitly set, no remote ID value is inserted in the DHCP request packet header. If the remote-id option is specified but is not qualified by a keyword, the MAC address of the host device (the switch) is used as the remote ID.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740 • Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770 • Setting Up DHCP Option 82 on the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773 • [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches • RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046.

secure-access-port

```
Syntax  secure-access-port {
        deactivate;
        dhcp-snooping-file {
            location (local_pathname | remote_URL);
            timeout seconds;
            write-interval seconds;
        }
        interface (all | interface-name) {
            allowed-mac mac-address-list;
            (dhcp-trusted | no-dhcp-trusted);
            fcoe-trusted;
            mac-limit limit {
                <action action>;
            }
            no-allowed-mac-log;
            persistent-learning;
            static-ip ip-address {
                vlan vlan-name;
                mac mac-address;
            }
        }
        vlan (all | vlan-name) {
            (arp-inspection | no-arp-inspection) [
                forwarding-class (for DHCP Snooping or DAI Packets) class-name;
            ]
            dhcp-option82 {
                circuit-id {
                    prefix (Circuit ID for Option 82) hostname;
                    use-interface-description;
                    use-vlan-id;
                }
                remote-id {
                    prefix (Remote ID for Option 82) hostname | mac | none;
                    use-interface-description;
                    use-string string;
                }
                vendor-id <string>;
            }
            (examine-dhcp | no-examine-dhcp) {
                forwarding-class (for DHCP Snooping or DAI Packets) class-name;
            }
            examine-fip {
                examine-vn2vn {
                    beacon-period milliseconds;
                }
                fc-map fc-map-value;
            }
            mac-move-limit limit action action;
        }
    }
```

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure port security features, including MAC limiting and whether interfaces can receive DHCP responses, and apply dynamic ARP inspection, DHCP snooping, DHCP option 82, and MAC move limiting on no VLANs, specific VLANs, or all VLANs.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of Access Port Protection on page 4674 • Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688 • Understanding Trusted and Untrusted Ports on page 4690 • Configuring MAC Limiting on page 4758 • Enabling a Trusted Port for DHCP on page 4769

static-ip

Syntax	<pre>static-ip <i>ip-address</i>; mac <i>mac-address</i>; vlan <i>vlan-name</i>; }</pre>
Hierarchy Level	[edit ethernet-switching-optionssecure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 12.1 on the QFX Series.
Description	Bind a static IP address to a MAC address in the DHCP snooping database.
Options	<p><i>ip-address</i>—IP address assigned to the device connected on the specified interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 4763

timeout (DHCP Snooping)

Syntax	timeout <i>seconds</i> ;
Hierarchy Level	[edit ethernet-switching-options secure-access-port dhcp-snooping-file]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify a timeout value for remote read and write operations. This value determines the amount of time that the switch waits for a remote system to respond when the DHCP snooping database is stored on a remote FTP site.
Default	None
Options	<i>seconds</i> —Value in seconds. Range: 10 through 3600
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding DHCP Snooping for Port Security on page 4678

use-interface-description

Syntax	use-interface-description;
Hierarchy Level	<p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 circuit-id],</p> <p>[edit forwarding-options helpers bootp dhcp-option82 circuit-id],</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 circuit-id]</p> <p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 remote-id]</p> <p>[edit forwarding-options helpers bootp dhcp-option82 remote-id]</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 remote-id]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Use the interface description rather than the interface name (the default) in the circuit ID or remote ID value in the DHCP option 82 information.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740 • Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770 • Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773 • [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches • RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046.

use-string

Syntax	<code>use-string <i>string</i>;</code>
Hierarchy Level	<code>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 <i>remote-id</i>]</code> <code>[edit forwarding-options helpers bootp dhcp-option82 <i>remote-id</i>]</code> <code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 <i>remote-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Use a string rather than the MAC address of the host system (the default) in the remote ID value in the DHCP option 82 information.
Options	<i>string</i> —Character string used as the remote ID value. Range: 1–255 characters
Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740• Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744• Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770• Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773• [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches• RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046.

use-vlan-id

Syntax	use-vlan-id;
Hierarchy Level	[edit forwarding-options helpers bootp dhcp-option82-circuit-id] [edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82-circuit-id]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Use the VLAN ID rather than the VLAN name (the default) in the circuit ID value in the DHCP option 82 information.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740 • Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770 • Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773 • [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches • RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046.

vendor-id

Syntax	<code>vendor-id <string>;</code>
Hierarchy Level	<code>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82]</code> <code>[edit forwarding-options helpers bootp dhcp-option82]</code> <code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</code>
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.
Default	If vendor-id is not explicitly configured for DHCP option 82, no vendor ID is set.
Options	string —(Optional) A single string that designates the vendor ID. Range: 1–255 characters Default: If you specify vendor-id with no string value, the default vendor ID Juniper is configured.
Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server on page 4740• Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server on page 4744• Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4770• Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773• [edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches

vlan (Secure Access Port)

```
Syntax  vlan (all | vlan-name) {
    examine-fip {
        examine-vn2vn {
            beacon-period milliseconds;
        }
        fc-map fc-map-value;
    }
    dhcp-option82
    circuit-id {
        prefix (Circuit ID for Option 82) hostname;
        use-interface-description;
        use-vlan-id;
    }
    remote-id {
        prefix (Remote ID for Option 82) hostname | mac | none;
        use-interface-description;
        use-string string;
    }
    vendor-id <string>;
}
(arp-inspection | no-arp-inspection);
circuit-id {
    prefix (Circuit ID for Option 82) hostname;
    use-interface-description;
    use-vlan-id;
}
remote-id {
    prefix (Remote ID for Option 82) hostname | mac | none;
    use-interface-description;
    use-string string;
}
vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp);
mac-move-limit limit action action;
}
```

Hierarchy Level [edit [ethernet-switching-options secure-access-port](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Apply DHCP snooping, dynamic ARP inspection (DAI), DHCP option 82, and MAC move limiting.



TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

The remaining statements are explained separately.

Options	all —Apply DHCP snooping, DAI, DHCP option 82, and MAC move limiting to all VLANs. vlan-name —Apply DHCP snooping, DAI, DHCP option 82, and MAC move limiting to the specified VLAN.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of Access Port Protection on page 4674• Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688• Understanding Trusted and Untrusted Ports on page 4690• Configuring MAC Limiting on page 4758• Enabling a Trusted Port for DHCP on page 4769

vlan (Static IP)

Syntax	<code>vlan <i>vlan-name</i>;</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>) static-ip ip-address]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series switches.
Description	Associate a static IP address with the specified VLAN.
Options	vlan-name —Name of a VLAN associated with the specified interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 4763

write-interval

Syntax	<code>write-interval <i>seconds</i>;</code>
Hierarchy Level	[edit ethernet-switching-options secure-access-port dhcp-snooping-file]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify how frequently the switch writes the database entries from memory into the specified DHCP snooping database file.
Default	None
Options	<i>seconds</i> —Value in seconds. Range: 60 through 86400
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding DHCP Snooping for Port Security on page 4678

Configuration Statements for Device Security

- [action-shutdown on page 4836](#)
- [bandwidth on page 4837](#)
- [ethernet-switching-options on page 4838](#)
- [interface \(Storm Control\) on page 4840](#)
- [no-broadcast on page 4841](#)
- [no-multicast on page 4841](#)
- [no-unknown-unicast on page 4842](#)
- [storm-control on page 4843](#)

action-shutdown

Syntax	action-shutdown;
Hierarchy Level	[edit ethernet-switching-options storm-control]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Shut down or disable interfaces when the storm control level is exceeded, as follows:</p> <ul style="list-style-type: none">• If you set both the action-shutdown and the port-error-disable statements, the affected interfaces are disabled temporarily and recover automatically when the disable timeout expires.• If you set the action-shutdown statement and do not set the port-error-disable statement, the affected interfaces are shut down when the storm control level is exceeded, and they do not recover automatically. You must issue the clear ethernet-switching port-error command to clear the port error and restore the interfaces to service.
Default	The action-shutdown feature is disabled. If the storm control level is exceeded, the switch drops broadcast and unknown unicast messages on the specified interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 4694• Example: Configuring Storm Control on page 4713• port-error-disable on page 4823• disable-timeout on page 4809• clear ethernet-switching port-error on page 4859

bandwidth

Syntax	<code>bandwidth <i>bandwidth</i>;</code>
Hierarchy Level	[edit <code>ethernet-switching-options storm-control interface</code> (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces configured for storm control, configure the storm control level as the bandwidth in kilobits per second (Kbps). If the combination of broadcast and unknown unicast traffic exceeds this level, the switch performs the appropriate action.
Default	None.
Options	bandwidth —Broadcast and unknown unicast traffic rate in Kbps. Range: 100 through 10000000 Kbps Default: None



NOTE: When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually enforced. For example, if you configure a bandwidth limit of 150 Kbps, storm control enforces a bandwidth limit of 128 Kbps.



CAUTION: Junos OS allows you to configure a storm control value that exceeds the bandwidth of the interface. If you configure an interface with such a value, storm control does not drop broadcast or unknown unicast packets even if they consume all the available bandwidth.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none"> • Understanding Storm Control on page 4694 • Example: Configuring Storm Control on page 4713 • action-shutdown on page 4836 • port-error-disable on page 4823 • disable-timeout on page 4809 • clear ethernet-switching port-error on page 4859
------------------------------	--

ethernet-switching-options

```
Syntax ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
      output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds {
  }
  port-error-disable {
    disable-timeout timeout;
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      fcoe-trusted;
      mac-limit limit action action;
      no-allowed-mac-log;
    }
    vlan (all | vlan-name) {
      (arp-inspection | no-arp-inspection) [
        forwarding-class (for DHCP Snooping or DAI Packets) class-name;
      ]
    }
  }
}
```

```

dhcp-option82 {
  circuit-id {
    prefix (Circuit ID for Option 82) hostname;
    use-interface-description;
    use-vlan-id;
  }
  remote-id {
    prefix (Remote ID for Option 82) hostname | mac | none;
    use-interface-description;
    use-string string;
  }
  vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp) {
  forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
  examine-vn2vn {
    beacon-period milliseconds;
  }
  fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
}

```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
 - [Overview of Access Port Protection on page 4674](#)
 - [Understanding Storm Control on page 4694](#)

interface (Storm Control)

Syntax	<pre>interface (all <i>interface-name</i>) { bandwidth <i>bandwidth</i>; no-broadcast; no-multicast; no-unknown-unicast; }</pre>
Hierarchy Level	[edit ethernet-switching-options storm-control]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Apply storm control to all interfaces or to the specified interface.</p> <p>The remaining statement is explained separately.</p>
Default	Storm control is disabled.
Options	<p>all—Apply storm control to all interfaces.</p> <p><i>interface-name</i>—Apply storm control to the specified interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 4694• Example: Configuring Storm Control on page 4713

no-broadcast

Syntax	no-broadcast;
Hierarchy Level	[edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces configured for storm control, disable broadcast traffic storm control on the interface.
Default	When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Storm Control on page 4694 • Example: Configuring Storm Control on page 4713

no-multicast

Syntax	no-multicast;
Hierarchy Level	[edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.
Default	Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Storm Control on page 4694 • Example: Configuring Storm Control on page 4713

no-unknown-unicast

Syntax	no-unknown-unicast;
Hierarchy Level	[edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces configured for storm control, disable unknown unicast traffic storm control on the interface.
Default	When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 4694• Example: Configuring Storm Control on page 4713

storm-control

Syntax

```
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-multicast;
    no-unknown-unicast;
  }
}
```

Hierarchy Level [edit [ethernet-switching-options](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Apply storm control to all interfaces or to the specified interfaces.

The statements are explained separately.



NOTE: The **no-multicast** option is not supported on QFabric systems.

Default By default, storm control is enabled on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. You can change the storm control level by configuring it as a specific bandwidth value.

When you configure storm control bandwidth on an aggregated Ethernet interface, each member of the aggregated interface is assigned that bandwidth. For example, if you configure 7000000 Kbps on aggregated interface **ae1**, and **ae1** has two members, **xe-2:0/0/0** and **xe-2:0/0/1**, each member is allowed a bandwidth level of 7000000 Kbps. Thus, the storm control bandwidth on **ae1** could be as much as 14000000 Kbps of combined broadcast and unknown unicast traffic.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding Storm Control on page 4694](#)
- [Example: Configuring Storm Control on page 4713](#)
- [port-error-disable on page 4823](#)
- [disable-timeout on page 4809](#)
- [clear ethernet-switching port-error on page 4859](#)

CHAPTER 53

Administration

- [Routine Monitoring on page 4845](#)
- [Monitoring Commands on page 4856](#)

Routine Monitoring

- [Monitoring Firewall Filter Traffic on page 4845](#)
- [Monitoring Port Security on page 4847](#)
- [Verifying That Firewall Filters Are Operational on page 4848](#)
- [Verifying That DAI Is Working Correctly on page 4848](#)
- [Verifying That DHCP Snooping Is Working Correctly on page 4849](#)
- [Verifying That MAC Limiting Is Working Correctly on page 4850](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 4853](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 4854](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly on page 4854](#)
- [Verifying That Three-Color Policers Are Operational on page 4855](#)
- [Verifying That Two-Color Policers Are Operational on page 4856](#)

Monitoring Firewall Filter Traffic

You can use operational mode commands to monitor firewall filter traffic.

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on page 4845](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 4846](#)
- [Monitoring Traffic for a Specific Policer on page 4846](#)

Monitoring Traffic for All Firewall Filters and Policers That Are Configured

Purpose Monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

Action Use the **show firewall** operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
```

Name	Bytes	Packets
counter-employee-web	3348	27
Filter: ingress-port-limit-tcp-icmp		
Counters:		
Name	Bytes	Packets
icmp-counter	560	10
Policers:		
Name	Packets	
icmp-connection-policer	10	
tcp-connection-policer	0	
Filter: ingress-vlan-rogue-block		
Filter: ingress-vlan-limit-guest		

Meaning The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

Monitoring Traffic for a Specific Firewall Filter

Purpose Monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded policer rate limits.

Action Use the **show firewall filter *filter-name*** operational mode command:

```
user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                      Bytes      Packets
icmp-counter              560        10
```

Meaning The **show firewall filter *filter-name*** command limits the display information to the counters and policers that are defined for the specified filter.

Monitoring Traffic for a Specific Policer

Purpose Monitor the number of packets that exceeded the rate limits of a policer:

Action Use the **show firewall policer *policer-name*** operational mode command:

```
user@switch> show firewall policer icmp-connection-policer
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                      Packets
icmp-connection-policer  10
```

Meaning The **show firewall policer *policer-name*** command displays the number of packets that exceeded the rate limits for the specified policer.

Related Documentation

- [Configuring Firewall Filters on page 4747](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

- [Verifying That Firewall Filters Are Operational on page 4848](#)

Monitoring Port Security

Purpose Use the monitoring functionality to view these port security details:

- DHCP snooping database for a VLAN or all VLANs
- ARP inspection details for all interfaces

Action To monitor port security in the J-Web interface, select **Monitor > Security > Port Security**.

To monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp snooping binding**
- **clear dhcp snooping binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or MAC addresses.
- **show arp inspection statistics**
- **clear arp inspection statistics**

Meaning The J-Web Port Security Monitoring page comprises two sections:

- **DHCP Snooping**—Displays the DHCP snooping database for all the VLANs for which DHCP snooping is enabled. To view the DHCP snooping database for a specific VLAN, select the specific VLAN from the list.
- **ARP Inspection**—Displays the ARP inspection details for all interfaces. The information includes details of the number of packets that passed ARP inspection and the number of packets that failed the inspection. The pie chart graphically represents these statistics when you select an interface. To view ARP inspection statistics for a specific interface, select the interface from the list.

You have the following options on the page:

- **Clear ALL**—Clears the DHCP snooping database, either for all VLANs if the option **ALL** has been selected in the Select VLANs list or for the specific VLAN that has been selected in that list.
- **Clear**—Deletes a specific IP address from the DHCP snooping database.

To clear ARP statistics on the page, click **Clear All** in the ARP Statistics section.

Use the CLI commands to show and clear DHCP snooping database and ARP inspection statistics details.

Related Documentation

- [Configuring Port Security \(CLI Procedure\) on page 4756](#)
- [Configuring Port Security \(J-Web Procedure\)](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)

Verifying That Firewall Filters Are Operational

Purpose Verify that firewall filters are working properly after you apply them to ports, VLANs, or Layer 3 interfaces.

Action Use the **show firewall** operational mode command to verify that the firewall filters are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes      Packets
counter-employee-web               0           0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes      Packets
icmp-counter                       560         10
Policers:
Name                               Packets
icmp-connection-policer           10
tcp-connection-policer            0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. In the above example, the **icmp-counter** in the filter **ingress-port-limit-tcp-icmp** shows that the filter matched 10 packets. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits. The policer **icmp-connection-policer** shows that 10 ICMP packets were policed.

Related Documentation

- [Configuring Firewall Filters on page 4747](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)
- [Monitoring Firewall Filter Traffic on page 4845](#)

Verifying That DAI Is Working Correctly

Purpose Verify that dynamic ARP inspection (DAI) is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface      Packets received  ARP inspection pass  ARP inspection failed
-----
ge-0/0/1.0      7                 5                    2
ge-0/0/2.0     10                10                   0
ge-0/0/3.0     12                12                   0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4766](#)
 - [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)
 - [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
 - [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
 - [Monitoring Port Security on page 4847](#)

Verifying That DHCP Snooping Is Working Correctly

Purpose Verify that DHCP snooping is working on the switch and that the DHCP snooping database is correctly populated with both dynamic and static bindings.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC address	IP address	Lease (seconds)	Type	VLAN	Interface
00:05:85:3A:82:77	192.0.2.17	600	dynamic	employee	ge-0/0/1.0
00:05:85:3A:82:79	192.0.2.18	653	dynamic	employee	ge-0/0/1.0
00:05:85:3A:82:80	192.0.2.19	720	dynamic	employee	ge-0/0/2.0
00:05:85:3A:82:81	192.0.2.20	932	dynamic	employee	ge-0/0/2.0
00:05:85:3A:82:83	192.0.2.21	1230	dynamic	employee	ge-0/0/2.0
00:05:85:27:32:88	192.0.2.22	–	static	data	ge-0/0/4.0

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires. Static IP addresses have no assigned lease time. The statically configured entry never expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

Related Documentation

- [Enabling DHCP Snooping \(CLI Procedure\) on page 4764](#)
- [Enabling DHCP Snooping \(J-Web Procedure\)](#)
- [Configuring Static IP Addresses for DHCP Bindings on Access Ports \(CLI Procedure\) on page 4763](#)
- [Example: Configuring Basic Port Security Features on page 4706](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch on page 4725](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 4732](#)
- [Monitoring Port Security on page 4847](#)
- [Troubleshooting Port Security](#)

Verifying That MAC Limiting Is Working Correctly

MAC limiting protects against flooding of the Ethernet switching table by setting a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. The allowed MAC method binds MAC addresses to a VLAN so that the address is not registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

This topic includes the following tasks:

1. [Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly on page 4851](#)
2. [Verifying That Allowed MAC Addresses Are Working Correctly on page 4851](#)
3. [Verifying That Interfaces Are Shut Down on page 4852](#)
4. [Customizing the Ethernet Switching Table Display to View Information for a Specific Interface on page 4852](#)

Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly

Purpose Verify that MAC limiting for dynamic MAC addresses is working.

Action Display the MAC addresses that have been learned. The following sample output shows the results of sending two packets from hosts connected to **xe-1:0/0/1** and five packets from hosts connected to **xe-1:0/0/2**, with both interfaces configured with a MAC limit of 4 and the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	*	Flood	-	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:77	Learn	0	xe-1:0/0/1.0
employee-vlan	00:05:85:3A:82:79	Learn	0	xe-1:0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	xe-1:0/0/2.0

Meaning The output shows that the fifth packet received on the **xe-1:0/0/2** interface was dropped because it exceeded the MAC limit for that interface. The address was not learned, and thus an asterisk (*) rather than an address appears in the MAC address column in the first line of the sample output.

Verifying That Allowed MAC Addresses Are Working Correctly

Purpose Verify that allowed MAC addresses are working.

Action Display the MAC cache information after allowed MAC addresses have been configured on an interface. The following sample shows the MAC cache after four allowed MAC addresses had been configured on interface **xe-1:0/0/2** and a fifth MAC address appeared on the interface.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:80	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	xe-1:0/0/2.0
employee-vlan	*	Flood	-	xe-1:0/0/2.0

Meaning Because the fifth address was not allowed it was not learned, and an asterisk (*) rather than an address appears in the MAC address column in the last line of the sample output.

Verifying That Interfaces Are Shut Down

Purpose Verify that an interface is shut down when the MAC limit is exceeded.

Action For more information about interfaces that have been shut down because the MAC limit was exceeded, use the **show ethernet-switching interfaces** command.

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Tag	Tagging	Blocking
bme0.32770	down	mgmt	untagged	unblocked	
xe-0/0/0.0	down	v1	untagged	MAC limit exceeded	
xe-0/0/1.0	up	v1	untagged	unblocked	
xe-0/0/2.0	up	v1	untagged	unblocked	
me0.0	up	mgmt	untagged	unblocked	



NOTE: You can configure interfaces to recover automatically when the MAC limit has been exceeded by specifying the **port-error-disable** statement with a **disable timeout** value. The switch automatically restores the disabled interface to service when the disable timeout expires. The **port-error-disable** configuration does not apply to preexisting error conditions—it affects only error conditions that are detected after the **port-error-disable** statement has been enabled and the configuration has been committed. To clear a preexisting error condition and restore the interface to service, use the **clear ethernet-switching port-error** command.

Customizing the Ethernet Switching Table Display to View Information for a Specific Interface

Purpose You can use the **show ethernet-switching table** command to view information for a specific interface.

Action For example, to display the MAC addresses that have been learned on the **xe-0/0/2** interface, enter:

```
user@switch> show ethernet-switching table interface xe-0/0/2.0
```

Ethernet-switching table: 1 unicast entries

VLAN	MAC address	Type	Age	Interfaces
v1	*	Flood	-	All-members
v1	00:00:06:00:00:00	Learn	0	xe-0/0/2.0

Meaning The MAC limit value for the **xe-0/0/2** interface had been set to **1**, and the output shows that only one MAC address was learned and added to the MAC cache.

- Related Documentation**
- [Configuring MAC Limiting on page 4758](#)
 - [Monitoring Port Security on page 4847](#)
 - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
 - [Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 4737](#)
 - [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 4715](#)

Verifying That MAC Move Limiting Is Working Correctly

Purpose Verify that MAC move limiting is working on the switch.

Action Display the MAC addresses in the Ethernet switching table when MAC move limiting has been configured for a VLAN. The following sample shows the results after two of the hosts on **ge-0/0/2** sent packets after the MAC addresses for those hosts had moved to other interfaces more than five times in 1 second. The VLAN, **employee-vlan**, was set to a MAC move limit of **5** with the action **drop**:

```
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 7 entries, 4 learned
VLAN          MAC address      Type      Age      Interfaces
employee-vlan 00:05:85:3A:82:77 Learn      0      ge-0/0/1.0
employee-vlan 00:05:85:3A:82:79 Learn      0      ge-0/0/1.0
employee-vlan 00:05:85:3A:82:80 Learn      0      ge-0/0/2.0
employee-vlan 00:05:85:3A:82:81 Learn      0      ge-0/0/2.0
employee-vlan *              Flood     -      ge-0/0/2.0
employee-vlan *              Flood     -      ge-0/0/2.0
```

Meaning The last two lines of the sample output show that MAC addresses for two hosts on **ge-0/0/2** were not learned, because the hosts had been moved back and forth from the original interfaces more than five times in 1 second.

- Related Documentation**
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4760](#)
 - [Configuring MAC Move Limiting \(J-Web Procedure\)](#)
 - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
 - [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Monitoring Port Security on page 4847](#)

Verifying That the Port Error Disable Setting Is Working Correctly

Purpose Verify that the port error disable setting is working as expected for MAC limited and storm control interfaces.

Action Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
Interface      State    VLAN members    Blocking
xe-2:0/0/0.0   up       T1122           unblocked
xe-2:0/0/1.0   down     default         MAC limit exceeded
xe-2:0/0/2.0   down     default         Storm control in effect
xe-2:0/0/3.0   down     default         unblocked
xe-2:0/0/4.0   down     default         unblocked
xe-2:0/0/5.0   down     default         unblocked
xe-2:0/0/6.0   down     default         unblocked
```

Meaning For interfaces disabled by port security features, the sample output from the **show ethernet-switching interfaces** command specifies the reasons that the interfaces are disabled:

- **MAC limit exceeded**—The interface is temporarily disabled because of a **mac-limit** error. The disabled interface is automatically restored to service when the **disable-timeout (Port Error Disable)** expires.
- **MAC move limit exceeded**—The interface is temporarily disabled because of a **mac-move-limit** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.
- **Storm control in effect**—The interface is temporarily disabled because of a **storm-control** error. The disabled interface is automatically restored to service when the **disable-timeout (Port Error Disable)** expires.

- Related Documentation**
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 4688](#)
 - [port-error-disable on page 4823](#)
 - [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\) on page 4762](#)

Verifying That a Trusted DHCP Server Is Working Correctly

Purpose Verify that a DHCP trusted server is working on the switch. See what happens when the DHCP server is trusted and then untrusted.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC Address	IP Address	Lease	Type	VLAN	Interface
-----	-----	-----	----	----	-----
00:05:85:3A:82:77	192.0.2.17	600	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:79	192.0.2.18	653	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:80	192.0.2.19	720	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:3A:82:81	192.0.2.20	932	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:3A:82:83	192.0.2.21	1230	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:27:32:88	192.0.2.22	3200	dynamic	employee-vlan	ge-0/0/2.0

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

- Related Documentation**
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4768](#)
 - [Enabling a Trusted Port for DHCP on page 4769](#)
 - [Enabling a Trusted DHCP Server \(J-Web Procedure\)](#)
 - [Example: Configuring Basic Port Security Features on page 4706](#)
 - [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 4722](#)
 - [Monitoring Port Security on page 4847](#)
 - [Troubleshooting Port Security](#)

Verifying That Three-Color Policers Are Operational

Purpose Verify that three-color policers in firewall filter configurations are working properly.

Action Use the following operational mode commands to verify that a three-color policer is working properly:

- **show class-of-service forwarding-table classifiers**
- **show interfaces *interface-name* extensive**

- `show interfaces queue interface-name`

- Related Documentation**
- [Overview of Policers on page 4664](#)
 - [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

Verifying That Two-Color Policers Are Operational

Purpose Verify that two-color policers in firewall filter configurations are working properly.

Action Use the `show firewall policer` operational mode command to verify that the policers are working properly:

```
user@switch> show firewall policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                                     Packets
icmp-connection-policer                  10
tcp-connection-policer                   539
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The `show firewall policer` command displays the names of all firewall filters and policers that are configured. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

- Related Documentation**
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)
 - [Configuring Firewall Filters on page 4747](#)
 - [Monitoring Firewall Filter Traffic on page 4845](#)

Monitoring Commands

- `clear arp inspection statistics`
- `clear dhcp snooping binding`
- `clear ethernet-switching port-error`
- `clear firewall`
- `show arp inspection statistics`
- `show dhcp snooping binding`
- `show firewall`
- `show firewall policer`
- `show interfaces filters`

clear arp inspection statistics

Syntax	clear arp inspection statistics <interface <i>interface</i> >
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear ARP inspection statistics.
Options	none —Clears ARP statistics on all interfaces. interface <i>interface-names</i> —(Optional) Clear ARP statistics on one or more interfaces.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show arp inspection statistics on page 4861 • Example: Configuring Basic Port Security Features on page 4706 • Verifying That DAI Is Working Correctly on page 4848
List of Sample Output	clear arp inspection statistics on page 4857
Output Fields	This command produces no output.

Sample Output

clear arp inspection statistics

```
user@switch> clear arp inspection statistics
```

clear dhcp snooping binding

Syntax	<code>clear dhcp snooping binding</code> <code><mac (all <i>mac-address</i>)></code> <code><vlan (all <i>vlan-name</i>)></code> <code><vlan (all <i>vlan-name</i>) mac (all <i>mac-address</i>)></code>
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear the DHCP snooping database information.
Options	mac (all <i>mac-address</i>) —(Optional) Clear DHCP snooping information for the specified MAC address or all MAC addresses. vlan (all <i>vlan-name</i>) —(Optional) Clear DHCP snooping information for the specified VLAN or all VLANs.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Basic Port Security Features on page 4706• show dhcp snooping binding on page 4862
List of Sample Output	clear dhcp snooping binding on page 4858
Output Fields	This command produces no output.

Sample Output

clear dhcp snooping binding

```
user@switch> clear dhcp snooping binding
```

clear ethernet-switching port-error

Syntax	clear ethernet-switching port-error <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch or from the specified interface, and restore the interfaces or the specified interface to service.
Options	none —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore the interfaces to service. interface <i>interface-name</i> —(Optional) Clear all MAC limiting, MAC move limiting, and storm control errors from the specified interface and restore the interface to service.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• Configuring MAC Limiting on page 4758• Example: Configuring Storm Control on page 4713• Configuring Port Security (CLI Procedure) on page 4756• port-error-disable on page 4823• Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure) on page 4762
Output Fields	This command produces no output.

clear firewall

Syntax	<code>clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i>)</code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Clear statistics provided by firewall filters.</p> <p>When you clear the counters of a filter, this not only impacts the counters shown by the CLI, but also the ones tracked by SNMP 2.</p>
Options	<p>all—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p>counter <i>counter-name</i>—Clear the packet and byte counts for the specified firewall filter counter.</p> <p>filter <i>filter-name</i>—Clear the packet and byte counts for the specified firewall filter.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• Verifying That Firewall Filters Are Operational on page 4848• Verifying That Two-Color Policers Are Operational on page 4856• Overview of Firewall Filters on page 4635• Overview of Policers on page 4664

Sample Output

clear firewall all

```
user@switch> clear firewall all
```

clear firewall counter

```
user@switch> clear firewall counter port-filter-counter
```

clear firewall filter

```
user@switch> clear firewall filter ingress-port-filter
```


show arp inspection statistics

Syntax	show arp inspection statistics
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display ARP inspection statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear arp inspection statistics on page 4857 • Example: Configuring Basic Port Security Features on page 4706 • Verifying That DAI Is Working Correctly on page 4848
List of Sample Output	show arp inspection statistics on page 4861
Output Fields	Table 379 on page 4861 lists the output fields for the show arp inspection statistics command. Output fields are listed in the approximate order in which they appear.

Table 379: show arp inspection statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which ARP inspection has been applied.	All levels
Packets received	Total number of packets total that underwent ARP inspection.	All levels
ARP inspection pass	Total number of packets that passed ARP inspection.	All levels
ARP inspection failed	Total number of packets that failed ARP inspection.	All levels

Sample Output

show arp inspection statistics

```
user@switch> show arp inspection statistics
```

Interface	Packets received	ARP inspection pass	ARP inspection failed
-----	-----	-----	-----
ge-0/0/0	0	0	0
ge-0/0/1	0	0	0
ge-0/0/2	0	0	0
ge-0/0/3	0	0	0
ge-0/0/4	0	0	0
ge-0/0/5	0	0	0
ge-0/0/6	0	0	0
ge-0/0/7	703	701	2

show dhcp snooping binding

Syntax	show dhcp snooping binding <interface <i>interface-name</i>> <vlan <i>vlan-name</i>>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the DHCP snooping database information.
Options	interface <i>interface-name</i> —(Optional) Display the DHCP snooping database information for an interface. vlan <i>vlan-name</i> —(Optional) Display the DHCP snooping database information for a VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear dhcp snooping binding • Example: Configuring Basic Port Security Features on page 4706 • Verifying That DHCP Snooping Is Working Correctly on page 4849
List of Sample Output	show dhcp snooping binding on page 4862
Output Fields	Table 380 on page 4862 lists the output fields for the show dhcp snooping binding command. Output fields are listed in the approximate order in which they appear.

Table 380: show dhcp snooping binding Output Fields

Field Name	Field Description	Level of Output
MAC Address	MAC address of the network device; bound to the IP address.	All levels
IP Address	IP address of the network device; bound to the MAC address.	All levels
Lease	Lease granted to the IP address.	All levels
Type	How the MAC address was acquired.	All levels
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcp snooping binding

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC Address	IP Address	Lease	Type	VLAN	Interface
00:00:01:00:00:03	192.0.2.0	640	dynamic	guest	ge-0/0/12.0
00:00:01:00:00:04	192.0.2.1	720	dynamic	guest	ge-0/0/12.0
00:00:01:00:00:05	192.0.2.5	800	dynamic	guest	ge-0/0/13.0

show firewall

Syntax	<pre>show firewall <counter <i>counter-name</i>> <filter <i>filter-name</i>> <log <detail interface <i>interface-name</i>>> <terse></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display statistics about configured firewall filters.
Options	<p>counter <i>counter-name</i>—(Optional) Display statistics about a particular firewall filter counter.</p> <p>filter <i>filter-name</i>—(Optional) Display statistics about a particular firewall filter.</p> <p>log—(Optional) Display log entries for all firewall filter activity.</p> <p>terse—(Optional) Display firewall filter names only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying That Firewall Filters Are Operational on page 4848 • Verifying That Two-Color Policers Are Operational on page 4856 • Overview of Firewall Filters on page 4635 • Overview of Policers on page 4664
List of Sample Output	<p>show firewall on page 4865</p> <p>show firewall filter <i>filter-name</i> on page 4866</p> <p>show firewall counter <i>counter-name</i> on page 4866</p> <p>show firewall log on page 4866</p> <p>show firewall log detail on page 4866</p>
Output Fields	Table 381 on page 4864 lists the output fields for the show firewall command. Output fields are listed in the approximate order in which they appear.

Table 381: show firewall Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.	All levels

Table 381: show firewall Output Fields (*continued*)

Field Name	Field Description	Level of Output
Counters	Display filter counter information: <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the count firewall filter action modifier. • Bytes—Number of bytes that match the filter term where the count action modifier was specified. • Packets—Number of packets that matched the filter term where the count action modifier was specified. 	All levels
Policers	Display policer information: <ul style="list-style-type: none"> • Name—Name of the policer that is configured at the [edit firewall policer] hierarchy level. • Packets—Number of packets that matched the filter term where the policer action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies. 	All levels
Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard 	All levels
Interface	Interface on which the firewall filter is applied.	All levels
Protocol	Name of the packet protocol.	All levels
Packet Length	Length of the packet.	All levels
Src Addr	Source address of the packet.	All levels
Dest Addr	Destination address of the packet.	All levels

Sample Output

show firewall

```

user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes      Packets
counter-employee-web                0           0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes      Packets
icmp-counter                        560        10
Policers:
Name                               Packets
icmp-connection-policer            10
tcp-connection-policer              0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest

```

show firewall filter filter-name

```

user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                                     Bytes          Packets
icmp-counter                             560            10
Policers:
Name                                     Packets
icmp-connection-policer                  10
tcp-connection-policer                   0

```

show firewall counter counter-name

```

user@switch> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                     Bytes          Packets
icmp-counter                             560            10

```

show firewall log

```

user@switch> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
08:00:53  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:52  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:51  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:50  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:49  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:48  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:47  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4

```

show firewall log detail

```

user@switch> show firewall log detail
Log :

Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of

```

```
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
```

show firewall policer

Syntax	<code>show firewall policer</code> <code><policer-name></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display statistics about configured policers.
Options	<p>none—Display the count of policed packets for all configured policers.</p> <p>policer-name—(Optional) Display the count of policed packets for the specified policer.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying That Firewall Filters Are Operational on page 4848 • Verifying That Two-Color Policers Are Operational on page 4856 • Overview of Firewall Filters on page 4635 • Overview of Policers on page 4664
List of Sample Output	<p>show firewall policer on page 4868</p> <p>show firewall policer policer-name on page 4869</p>
Output Fields	Table 382 on page 4868 lists the output fields for the show firewall policer command. Output fields are listed in the approximate order in which they appear.

Table 382: show firewall policer Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the <code>[edit firewall family family-name filter]</code> hierarchy level.	All levels
Policers	Display policer information: <ul style="list-style-type: none"> • Filter—Name of filter that specifies the policer action modifier. • Name—Name of policer. • Packets—Number of packets that matched the filter term in which the policer action modifier is specified. This is the number of packets that exceed the rate limits that the policer specifies. 	All levels

Sample Output

show firewall policer

```

user@switch> show firewall policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                                     Packets

```


icmp-connection-policer	0
tcp-connection-policer	0
Filter: ingress-vlan-rogue-block	

show firewall policer policer-name

```
user@switch> show firewall policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                               Packets
tcp-connection-policer             0
```

show interfaces filters

Syntax	<code>show interfaces filters</code> <code><interface-name></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display firewall filters that are configured on each interface in a switch.
Options	none —Display firewall filter information about all interfaces. interface-name —(Optional) Display firewall filter information about a particular interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show firewall on page 4864
List of Sample Output	show interfaces filters on page 4870 show interfaces filters interface-name on page 4871
Output Fields	Table 383 on page 4870 lists the output fields for the show interfaces filters command. Output fields are listed in the approximate order in which they appear.

Table 383: show interfaces filters Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the physical interface.	All levels
Admin	Interface state: up or down .	All levels
Link	Link state: up or down .	All levels
Proto	Protocol that is configured on the interface.	All levels
Input Filter	Name of the firewall filter to be evaluated when packets are received on the interface.	All levels
Output Filter	Name of the firewall filter to be evaluated when packets are transmitted on the interface.	All levels

Sample Output

show interfaces filters

```

user@switch> show interfaces filters
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/6       up   up   eth-switch ingress-port-limit-tcp-icmp
ge-0/0/6.0     up   up   eth-switch ingress-port-limit-tcp-icmp
ge-0/0/7       up   down
ge-0/0/8       up   down

```

ge-0/0/9	up	down
ge-0/0/10	up	down
ge-0/0/10.0	up	down

show interfaces filters interface-name

```
user@switch> show interfaces filters ge-0/0/6
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch	ingress-port-limit-tcp-icmp	

Troubleshooting

- [Troubleshooting Procedures on page 4873](#)

Troubleshooting Procedures

- [Troubleshooting Firewall Filter Configuration on page 4873](#)
- [Troubleshooting Policer Configuration on page 4880](#)

Troubleshooting Firewall Filter Configuration

Use the following information to troubleshoot your firewall filter configuration.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 4873](#)
- [Filter Counts Previously Dropped Packet on page 4875](#)
- [Matching Packets Not Counted on page 4876](#)
- [Counter Reset When Editing Filter on page 4876](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 4876](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 4877](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 4877](#)
- [Egress Firewall Filters with Private VLANs on page 4877](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 4878](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 4878](#)
- [Invalid Statistics for Policer on page 4878](#)
- [Policers can Limit Egress Filters on page 4878](#)

Firewall Filter Configuration Returns a No Space Available in TCAM Message

Problem **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.  
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available

in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

Solution When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```



NOTE: The original filter is not deleted and is still available in the configuration.

Filter Counts Previously Dropped Packet

- Problem** **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
 - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

Solution This is expected behavior.

Matching Packets Not Counted

Problem **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet. For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **adminVLAN**, and interface xe-0/0/1 is a member of that VLAN.
- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

Solution This is expected behavior.

Counter Reset When Editing Filter

Problem **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

Solution This is expected behavior.

Cannot Include loss-priority and policer Actions in Same Term

Problem **Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

Solution This is expected behavior.

Cannot Egress Filter Certain Traffic Originating on QFX Switch

Problem **Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

Solution This is expected behavior.

Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

Problem **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

Solution This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

Egress Firewall Filters with Private VLANs

Problem **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

Solution These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

Egress Filtering of L2PT Traffic Not Supported

Problem **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

Solution This is expected behavior.

Cannot Drop BGP Packets in Certain Circumstances

Problem **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

Solution This is expected behavior.

Invalid Statistics for Policer

Problem **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

Solution This is expected behavior.

Policers can Limit Egress Filters

Problem **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume

two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
 - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
 - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

Solution You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

Related Documentation

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Configuring Firewall Filters on page 4747](#)
- [Verifying That Firewall Filters Are Operational on page 4848](#)

Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 4880](#)
- [Counter Reset When Editing Filter on page 4880](#)
- [Invalid Statistics for Policer on page 4880](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 4881](#)
- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 4882](#)
- [Policers Can Limit Egress Filters on page 4882](#)

Incomplete Count of Packet Drops

Problem **Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

Solution This is expected behavior.

Counter Reset When Editing Filter

Problem **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

Solution This is expected behavior.

Invalid Statistics for Policer

Problem **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

Solution This is expected behavior.

Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

Problem **Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

Solution This is expected behavior.

Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

Problem **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

Solution To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 4658](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

Policers Can Limit Egress Filters

Problem **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters. Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
 - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.

- Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

Solution You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

PART 18

Services

- [Overview on page 4887](#)
- [Configuration on page 4895](#)
- [Administration on page 4947](#)
- [Troubleshooting on page 4949](#)

CHAPTER 55

Overview

- [Port Mirroring on page 4887](#)
- [DHCP Relay on page 4892](#)

Port Mirroring

- [Understanding Port Mirroring on page 4887](#)
- [Understanding Layer 3 Logical Interfaces on page 4892](#)

Understanding Port Mirroring

- [Port Mirroring Overview on page 4887](#)
- [Port-Mirroring Terminology on page 4888](#)
- [Port Mirroring Constraints and Limitations on page 4889](#)

Port Mirroring Overview

Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is needed for traffic analysis on a switch because a switch normally sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to a local interface or a VLAN and run an analyzer application on a device connected to the interface or VLAN. You configure port mirroring by using the **analyzer** statement.

Keep performance in mind when configuring port mirroring. For example, If you mirror traffic from multiple ports, the mirrored traffic may exceed the capacity of the output interface. We recommend that you limit the amount of copied traffic by selecting specific interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter. Mirroring only the necessary packets reduces the possibility of a performance impact.

You can use port mirroring to copy any of the following:

- All packets entering or exiting an interface (in any combination)—For example, you can send copies of the packets entering some interfaces and the packets exiting other interfaces to the same local interface or VLAN. If you configure port mirroring to copy packets exiting an interface, traffic that originates on that switch or Node device (in a QFabric system) is not copied when it egresses. Only switched traffic is copied on egress. (See the limitation on egress mirroring below.)
- All packets entering a VLAN—You cannot use port mirroring to copy packets exiting a VLAN.
- Firewall-filtered sample—Sample of packets entering a port or VLAN. Configure a firewall filter to select certain packets for mirroring.



NOTE: Firewall filters are not supported on egress ports; therefore, you cannot specify policy-based sampling of packets exiting an interface.

Port-Mirroring Terminology

Table 384 on page 4888 lists the terms used in the documentation about port mirroring and provides definitions.

Table 384: Port Mirroring Terms and Definitions

Term	Description
Analyzer	Port-mirroring configuration. The analyzer includes a name, source interfaces or source VLAN, and a destination for mirrored packets (either a local access interface or a VLAN).
Output interface (also known as monitor interface)	<p>Access interface to which packet copies are sent and to which a device running an analyzer application is connected.</p> <p>The following limitations apply to an output interface:</p> <ul style="list-style-type: none"> • Cannot also be a source port. • Cannot be used for switching. • Cannot be an aggregated Ethernet interface (LAG). • Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP). • Loses any existing VLAN associations when you configure it as an analyzer output interface. <p>If the capacity of the output interface is insufficient to handle the traffic from the source ports, overflow packets are dropped.</p>

Table 384: Port Mirroring Terms and Definitions (*continued*)

Output IP address	<p>IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)</p> <ul style="list-style-type: none"> • An output IP address cannot be in the same subnet as any of the switch's management interfaces. • If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
Output VLAN (also known as monitor or analyzer VLAN)	<p>VLAN to which copies are sent and to which a device running an analyzer application is connected. The analyzer VLAN can span multiple switches.</p> <p>The following limitations apply to an output VLAN:</p> <ul style="list-style-type: none"> • Cannot be a private VLAN or VLAN range. • Cannot be shared by multiple analyzer statements. • An output VLAN interface cannot be a member of any other VLAN. • An output VLAN interface cannot be an aggregated Ethernet interface (LAG). • On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.
Input interface (also known as mirrored or monitored interface)	Interface that provides traffic to be mirrored. This traffic can be entering or exiting the interface. (Ingress or egress traffic can be mirrored.) An input interface cannot also be an output interface for an analyzer.
Monitoring station	Computer running an analyzer application.
Local port mirroring	Port-mirroring configuration in which the mirrored packets are sent to an interface on the same switch.
Remote port mirroring	Flooding mirrored packets to an analyzer VLAN that you create to receive mirror traffic or sending the mirrored packets to a remote IP address. (You cannot send mirrored packets to a remote IP address on a QFabric system.)
Policy-based mirroring	Mirroring of packets that match the match a firewall filter term. The action analyzer analyzer-name is used in the firewall filter to send the packets to the analyzer.

Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 4889](#)
- [Remote Port Mirroring Only on page 4891](#)

Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring with the QFX Series:

- You can create a total of four port-mirroring configurations on a QFX Series standalone switch.

- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
 - As many as four of the configurations can be for local port mirroring.
 - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
 - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
 - There can be no more than two configurations that mirror egress traffic.



NOTE: On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
 - **interface**
 - **ip-address**
 - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a QFX3500 device or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
 - Dedicated Virtual Chassis interfaces
 - Management interfaces (me0 or vme0)
 - Fibre Channel interfaces
 - Routed VLAN interfaces

- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

Remote Port Mirroring Only

The following constraints and limitations apply to remote port mirroring with the QFX Series:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

Related Documentation

- [Configuring Port Mirroring on page 4904](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 4900](#)
- [Troubleshooting Port Mirroring on page 4949](#)

Understanding Layer 3 Logical Interfaces

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks QFX3500 Switch to a Layer 2 switch. Only one physical connection is required between the switches. You can also use Layer 3 logical interfaces to provide alternative gateway addresses for smart DHCP relay.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series systems support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. Double-tagging, which is assigning more than one VLAN ID in the frame header, is not supported. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.

Related Documentation

- [Interfaces Overview on page 2415](#)
- [Configuring a Layer 3 Logical Interface on page 2537](#)
- [Configuring DHCP and BOOTP Relay on page 4907](#)
- *Junos® OS Network Interfaces*

DHCP Relay

- [DHCP and BOOTP Relay Overview on page 4892](#)

DHCP and BOOTP Relay Overview

You can configure a Juniper Networks switch to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that if the switch receives a broadcast DHCP or BOOTP request from a locally attached host (client), it relays the message to a specified DHCP or BOOTP server. You should configure the switch to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.

If you configure a switch to be a DHCP relay agent, you can also enable smart DHCP relay, which enables you to configure alternative gateway addresses for a DHCP server so that if the server fails to reply to the requests sent to the primary gateway address, the switch can resend the requests to the alternative gateway addresses. To use this

feature, you must configure a routed VLAN interface or Layer 3 subinterface with multiple IP addresses and configure that interface to be a relay agent.



NOTE: Because DHCP and BOOTP messages are broadcast and are not directed to a specific server, switch, or router, Juniper switches cannot function as both a DHCP server and a DHCP/BOOTP relay agent at the same time. The Junos operating system (Junos OS) generates a commit error if both options are configured at the same time, and the commit operation does not succeed until one of the options is removed.

**Related
Documentation**

- [Configuring DHCP and BOOTP Relay on page 4907](#)
- [bootp on page 4940](#)

CHAPTER 56

Configuration

- [Configuration Examples on page 4895](#)
- [Configuration Tasks on page 4904](#)
- [Configuration Statements for Port Mirroring on page 4910](#)
- [Configuration Statements for Encryption on page 4919](#)
- [Configuration Statements for DHCP Relay on page 4938](#)

Configuration Examples

- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 4900](#)

Example: Configuring Port Mirroring for Local Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies to a local interface for local monitoring.

This example describes how to configure port mirroring to copy traffic sent by employee computers to a switch to an access interface on the same switch.

- [Requirements on page 4895](#)
- [Overview and Topology on page 4896](#)
- [Mirroring All Employee Traffic for Local Analysis on page 4896](#)
- [Mirroring Employee-to-Web Traffic for Local Analysis on page 4897](#)
- [Verification on page 4899](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1
- A switch

Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering interfaces on the switch to an access interface on the same switch. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

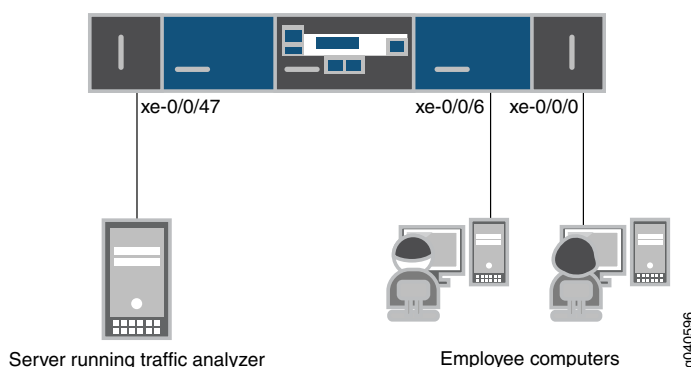
In this example, **xe-0/0/0** and **xe-0/0/6** serve as connections for employee computers. Interface **xe-0/0/47** is connected to a device running an analyzer application.



NOTE: Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Figure 185 on page 4896 shows the network topology for this example.

Figure 185: Network Topology for Local Port Mirroring Example



Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all traffic sent by employee computers for local analysis, perform the tasks explained in this section.

CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching
set interfaces xe-0/0/6 unit 0 family ethernet-switching
set interfaces xe-0/0/47 unit 0 family ethernet-switching
set ethernet-switching options analyzer employee-monitor input ingress interface xe-0/0/0.0
set ethernet-switching options analyzer employee-monitor input ingress interface xe-0/0/6.0
set ethernet-switching options analyzer employee-monitor output interface xe-0/0/47.0
```

Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the output interface:

1. Configure the interfaces connected to employee computers as input interfaces for the port-mirror analyzer **employee-monitor**:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/6.0
2. Configure the output analyzer interface for the employee-monitor analyzer. This will
be the destination interface for the mirrored packets:

[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output interface xe-0/0/47.0
```

Results Check the results of the configuration:

```
[edit]
user@switch# show ethernet-switching-options
analyzer employee-monitor {
  input {
    ingress {
      interface xe-0/0/0.0;
      interface xe-0/0/6.0;
    }
  }
  output {
    interface {
      xe-0/0/47.0;
    }
  }
}
```

Mirroring Employee-to-Web Traffic for Local Analysis

To mirror only traffic sent by employees to the Web for local analysis, perform the tasks explained in this section.

CLI Quick Configuration

To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:

```
[edit]
set ethernet-switching-options analyzer employee-web-monitor output interface xe-0/0/47.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer
employee-web-monitor
set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

Step-by-Step Procedure To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the output interface:

```
[edit interfaces]
user@switch# set xe-0/0/47 unit 0 family ethernet-switching
```

2. Configure the **employee-web-monitor** analyzer output. (Configure only the output—the input comes from the filter.)

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-web-monitor output interface xe-0/0/47.0
```

3. Configure a firewall filter called **watch-employee** that includes a term to match traffic sent to the Web and send it to the analyzer **employee-web-monitor**. Traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the analyzer:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from destination-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then analyzer employee-web-monitor
```

4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers):

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

Results Check the results of the configuration:

```
[edit]
user@switch# show ethernet-switching-options
  analyzer employee-web-monitor {
    output {
      interface xe-0/0/47.0;
    }
  }
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-web {
      from {
        destination-port 80;
      }
      then analyzer employee-web-monitor;
    }
  }
}
...
interfaces {
```

```

xe-0/0/0 {
  unit 0 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}
xe-0/0/6 {
  family ethernet-switching {
    filter {
      input watch-employee;
    }
  }
}
}

```

Verification

Verifying That the Analyzer Has Been Correctly Created

Purpose Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

Action You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```

user@switch> show analyzer
Analyzer name           : employee-monitor
Output interface        : xe-0/0/47.0
Mirror ratio            : 1
Loss priority           : Low
Ingress monitored interfaces : xe-0/0/0.0
Ingress monitored interfaces : xe-0/0/6.0
Egress monitored interfaces  : None

```

Meaning This output shows that the **employee-monitor** analyzer:

- Has a ratio of 1 (mirroring every packet, the default setting)
- Has a loss priority of low (set this option to high only when the analyzer output is to a VLAN)
- Is mirroring the traffic entering the **xe-0/0/0** and **xe-0/0/6** interfaces
- Is sending the mirrored traffic to the **xe-0/0/47** interface

Related Documentation

- [Understanding Port Mirroring on page 4887](#)
- [Configuring Port Mirroring on page 4904](#)

Example: Configuring Port Mirroring for Remote Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies either to a local interface for local monitoring or to a VLAN for remote monitoring. This example describes how to configure port mirroring for remote analysis.

- [Requirements on page 4900](#)
- [Overview and Topology on page 4900](#)
- [Mirroring All Employee Traffic for Remote Analysis on page 4901](#)
- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 4902](#)
- [Verification on page 4904](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 for the QFX Series
- A switch

Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to an analyzer VLAN so that you can perform analysis using a remote device. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

In this example:

- Interfaces **ge-0/0/0** and **ge-0/0/1** are Layer 2 interfaces that connect to employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to another switch.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.



NOTE: In addition to performing the configuration steps described here, you must also configure the analyzer VLAN (**remote-analyzer** in this example) on the other switches that are used to connect the source switch (the one in this configuration) to the one that the monitoring station is connected to.

Mirroring All Employee Traffic for Remote Analysis

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor output vlan remote-analyzer
```

Step-by-Step Procedure To configure basic remote port mirroring:

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):

```
[edit vlans]
user@switch# set vlans remote-analyzer vlan-id 999
```
2. Configure the interface connected to another switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```
4. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

Results Check the results of the configuration:

```
[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
```

Mirroring Employee-to-Web Traffic for Remote Analysis

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set ethernet-switching-options analyzer employee-web-monitor loss-priority high output vlan 999
set firewall family ethernet-switching filter watch-employee term employee-to-web from destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer employee-web-monitor
set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

Step-by-Step Procedure

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```
2. Configure an interface to associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-web-monitor** analyzer. (Configure only the output—the input comes from the filter.)

```
[edit ethernet-switching-options]
user@switch# set ethernet-switching-options analyzer employee-web-monitor output vlan 999
```
4. Configure a firewall filter called **watch-employee** to match traffic sent to the Web and send it to the analyzer **employee-web-monitor**:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then analyzer employee-web-monitor
```
5. Apply the firewall filter to the appropriate interfaces as an ingress filter:

```
[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```
6. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

Results Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  ...
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
```

```

        port-mode trunk;
        vlan {
            members remote-analyzer;
        }
    }
}
ge-0/0/0 {
    unit 0 {
        family ethernet-switching {
            filter {
                input watch-employee;
            }
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family ethernet-switching {
            filter {
                input watch-employee;
            }
        }
    }
}
...
firewall {
    family ethernet-switching {
        ...
        filter watch-employee {
            term employee-to-web {
                from {
                    destination-port 80;
                }
                then analyzer employee-web-monitor;
            }
        }
    }
}
ethernet-switching-options {
    analyzer employee-web-monitor {
        output {
            vlan {
                999;
            }
        }
    }
}
vpls {
    remote-analyzer {
        vlan-id 999;
    }
}

```

Verification

Verifying That the Analyzer Has Been Correctly Created

- Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.
- Action** You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.
- ```
user@switch> show analyzer
Analyzer name : employee-monitor
Output VLAN : remote-analyzer
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```
- Meaning** This output shows that the **employee-monitor** analyzer is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1** and is sending the mirror traffic to the analyzer **remote-analyzer**.
- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
  - [Configuring Port Mirroring on page 4904](#)
  - [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
  - [Overview of Firewall Filters on page 4635](#)

## Configuration Tasks

---

- [Configuring Port Mirroring on page 4904](#)
- [Configuring DHCP and BOOTP Relay on page 4907](#)

### Configuring Port Mirroring

You use port mirroring to copy packets and send the copies to a device running an application such as a network analyzer or intrusion detection application so that you can analyze traffic without delaying it. You can mirror traffic entering or exiting a port or entering a VLAN, and you can send the copies to a local access interface or to a VLAN through a trunk interface.

We recommend that you disable port mirroring when you are not using it. To avoid creating a performance issue If you do enable port mirroring, we recommend that you select specific input interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter.



**NOTE:** If you want to create additional analyzers without deleting an existing analyzer, first disable the existing analyzer using the **disable analyzer analyzer-name** command.

---



**NOTE:** You must configure port mirroring output interfaces as family `ethernet-switching`.

- [Configuring Port Mirroring for Local Analysis on page 4905](#)
- [Configuring Port Mirroring for Remote Analysis on page 4905](#)
- [Filtering the Traffic Entering an Analyzer on page 4906](#)

### Configuring Port Mirroring for Local Analysis

To mirror interface traffic to a local interface on the switch:

1. If you want to mirror traffic that is ingressing or egressing specific interfaces, choose a name for the port-mirroring configuration and configure what traffic should be mirrored by specifying the interfaces and direction of traffic:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some VLAN packets might contain incorrect VLAN IDs.



**NOTE:** If you configure mirroring for packets that egress an access interface, the original packets lose any VLAN tags when they exit the access interface, but the mirrored (copied) packets retain the VLAN tags when they are sent to the analyzer system.

2. If you want to specify that all traffic entering a VLAN should be mirrored, choose a name for the port-mirroring configuration and specify the VLAN:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name input ingress vlan vlan-name
```



**NOTE:** You cannot configure port mirroring to copy traffic that egresses a VLAN.

3. Configure the destination interface for the mirrored packets:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name output interface interface-name
```

### Configuring Port Mirroring for Remote Analysis

To mirror traffic to a VLAN for analysis at a remote location:

1. Configure a VLAN to carry the mirrored traffic:

```
[edit]
```

```
user@switch# set vlans vlan-name vlan-id number
```

2. Configure the interface that connects to another switch (the uplink interface) to trunk mode and associate it with the appropriate VLAN:

```
[edit]
```

```
user@switch# set interfaces interface-name unit 0 family ethernet-switching port-mode
trunk vlan members (vlan-name | vlan-id)
```

3. Configure the analyzer:

- a. Choose a name for the analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name
```

- b. Specify the interface to be mirrored and whether the traffic should be mirrored on ingress or egress:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```

- c. Specify the appropriate IP address or VLAN as the output (a VLAN is specified in this example:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name output vlan (vlan-name | vlan-id)
```

If you specify an IP address as the output, note the following constraints:

- The address cannot be in the same subnet as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (**inet.0** routing table).
- The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)

### Filtering the Traffic Entering an Analyzer

---

In addition to specifying which traffic to mirror by configuring an analyzer, you can also use a firewall filter to exercise more control over which packets are copied. For example, you might use a filter to specify that only traffic from certain applications be mirrored. The filter can use any of the available match conditions and must have an action of **analyzer *analyzer-name***. If you use the same analyzer in multiple filters or terms, the output packets are copied only once.



**NOTE:** You can include the action **analyzer** in ingress firewall filters only. You can apply ingress filters with this action to ports (Layer 2 interfaces), Layer 3 interfaces, and VLANs.

When you use a firewall filter as the input to an analyzer, you output the copied traffic to a local interface or a VLAN just as you do when a firewall is not involved.

To configure port mirroring with filters:

1. Configure an analyzer for local or remote analysis. Configure only the output. For example, for local analysis enter:

```
[edit ethernet-switching-options]
user@switch# set analyzer analyzer-name output interface interface-name
```



**NOTE:** Do not configure input to this analyzer.

2. Create a firewall filter using any of the available match conditions and specify the action as **analyzer *analyzer-name***.
3. Apply the firewall filter to the interfaces or VLAN that should provide the input to the analyzer:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching filter input
filter-name
[edit]
user@switch# set vlan (vlan-name | vlan-id) filter input filter-name
```

#### Related Documentation

- [Understanding Port Mirroring on page 4887](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 4900](#)
- [Overview of Firewall Filters on page 4635](#)

## Configuring DHCP and BOOTP Relay

You can configure the QFX Series to act as a Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) relay agent. This means that if a locally attached host can issue a DHCP or BOOTP request as a broadcast message and the switch relays the message to a specified DHCP or BOOTP server. You should configure a switch to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server.

If you configure a switch to be a DHCP relay agent, you can also enable smart DHCP relay, which allows you to configure alternative gateway addresses for a DHCP server so that if the server fails to reply to the requests sent using the primary gateway address, the switch can resend the requests via the alternative gateway addresses. To use this feature, you must configure a routed VLAN interface or Layer 3 logical interface with multiple IP addresses and configure that interface to be a relay agent.

- [Configuring a DHCP and BOOTP Relay Agent on page 4907](#)
- [Configuring DHCP Smart Relay on page 4909](#)

### Configuring a DHCP and BOOTP Relay Agent

To configure a switch to act as a DHCP and BOOTP relay agent, include the **bootp** statement at the **[edit forwarding-options helpers]** hierarchy level:

```
[edit forwarding-options helpers]
bootp {
 apply-secondary-as-giaddr text-description;
 client-response-ttl number;
 description text-description;
 interface (interface-name | interface-group) {
 client-response-ttl number;
 description text-description;
 maximum-hop-count number;
 minimum-wait-time seconds;
 no-listen;
 server address
 apply-secondary-as-giaddr
 }
 maximum-hop-count number;
 minimum-wait-time seconds;
 relay-agent-option;
 server server-identifier
}
```

To include a description of the BOOTP service, DHCP service, or interface, use the **description** statement.

To configure a logical interface or a group of logical interfaces with a specific DHCP relay or BOOTP configuration, include the **interface** statement.

To stop packets from being forwarded, include the **no-listen** statement.

To set the maximum allowed number in the hops field of the BOOTP message, include the **maximum-hop-count** statement. BOOTP messages that have a larger number in the hops field than the maximum allowed are not forwarded. If you omit the **maximum-hop-count** statement, the default maximum number of hops is four.

To set the minimum allowed number of seconds in the **secs** field of the BOOTP message, include the **minimum-wait-time** statement. This setting configures a minimum number of seconds since the client sent its first BOOTP request. BOOTP messages that have a smaller number in the **secs** field than the allowed minimum are not forwarded. The default value for the minimum wait time is zero (0).

To set the IP address that specify the DHCP or BOOTP server for the router, switch, or interface, include the **server** statement. You can include multiple **server** statements.

To set an IP time-to-live (TTL) value for DHCP response packets sent to a DHCP client, include the **client-response-ttl** statement.

The following example demonstrates a BOOTP relay agent configuration.

```
user@host# show forwarding-options
helpers {
 bootp {
 description "dhcp relay agent global parameters";
 server 192.168.55.44;
 server 172.16.0.3 routing-instance c3;
 maximum-hop-count 10;
 minimum-wait-time 8;
```



```

interface {
 xe-0/0/1 {
 description "use this info for this interface";
 server 10.10.10.10;
 server 192.168.14.14;
 maximum-hop-count 11;
 minimum-wait-time 3;
 }
 xe-0/0/2 {
 no-listen; ###ignore DHCPDISCOVER messages on this interface
 }
 all {
 description "globals apply to all other interfaces";
 }
}
}

```

### Configuring DHCP Smart Relay

You can use DHCP smart relay to provide redundancy and resiliency to your DHCP relay configuration. Smart relay provides additional relay functionality and requires all of the configuration settings required by DHCP relay. To use DHCP smart relay, you also need an interface with multiple IP addresses assigned to it. You can achieve this by doing either of the following tasks:

- Create a routed VLAN interface and assign at least two IP addresses to it. See [“Configuring Routed VLAN Interfaces” on page 2046](#) and [“Example: Configuring Routing Between VLANs on One Switch” on page 1976](#) for information about this approach.
- Create a Layer 3 logical interface (by using VLAN tagging) and assign at least two IP addresses to it. See [“Understanding Layer 3 Logical Interfaces” on page 2434](#) and [“Configuring a Layer 3 Logical Interface” on page 2537](#) for information about this approach.

Once you have created an interface with multiple IP addresses, complete the smart relay configuration by entering one of the following statements:

- **set forwarding-options helpers bootp smart-relay-global:** Use this statement to enable smart relay on all the interfaces that are configured as relay agents.
- **set forwarding-options helpers bootp interface *interface-name* smart-relay-agent:** Use this statement to enable smart relay on a specific interface.

When smart relay is configured for an interface, the switch initially sends DHCP request (discover) messages out of that interface using the primary address of the interface as the gateway IP address (in the giaddr field) for the DHCP message. If no DHCP offer message is received from a server in reply, the switch allows the client to send as many as three more discover messages using the same gateway IP address. If no DHCP offer message is received after three retries, the switch resends the discover message using the alternate IP address as the gateway IP address. If you configure more than two IP addresses on the relay agent interface, the switch repeats this process until a DHCP offer message is received or all of the IP addresses have been used without success.

## Configuration Statements for Port Mirroring

---

- [analyzer](#) on page 4911
- [egress](#) on page 4912
- [ethernet-switching-options](#) on page 4913
- [ingress \(ethernet-switching-options\)](#) on page 4915
- [input](#) on page 4916
- [interface \(Port Mirroring\)](#) on page 4917
- [ip-address \(Port Mirroring\)](#) on page 4918
- [output](#) on page 4918
- [vlan \(Port Mirroring\)](#) on page 4919

## analyzer

```
Syntax analyzer {
 name {
 input {
 egress {
 interface (all | interface-name);
 }
 ingress {
 interface (all | interface-name);
 vlan (vlan-id | vlan-name);
 }
 }
 output {
 interface interface-name;
 ip-address ip-address;
 vlan (vlan-id | vlan-name);
 }
 }
}
```

**Hierarchy Level** [edit [ethernet-switching-options](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Option **output** **vlan** added in Junos OS Release 12.1 for the QFX Series.  
Option **output** **ip-address** added in Junos OS Release 12.3 for the QFX Series.

**Description** Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:

- There can be no more than two configurations that mirror ingress traffic.
- There can be no more than two configurations that mirror egress traffic.

**Default** Port mirroring is disabled, and Junos OS creates no default analyzers.

**Options** **all**—Mirror all the access interfaces. Using this option does not cause the QSFP+ or management interfaces to be mirrored.



**CAUTION:** Configuring the **all** option in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

**name**—Name of the analyzer. The name can include as many as 125 characters; must begin with a letter; and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 4887](#)
- [Configuring Port Mirroring on page 4904](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)

---

## egress

---

**Syntax** egress {  
    **interface** (all | *interface-name*);  
}

**Hierarchy Level** [edit **ethernet-switching-options analyzer name input**]

**Release Information** Statement introduced in Junos OS Release 11.2 for the QFX Series.

**Description** Specify interfaces for which egressing traffic is mirrored.

The statement is explained separately.



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some of the mirrored packets might contain incorrect VLAN IDs.

---

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 4887](#)
- [Configuring Port Mirroring on page 4904](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)

## ethernet-switching-options

```

Syntax ethernet-switching-options {
 analyzer {
 name {
 input {
 egress {
 interface (all | interface-name);
 }
 ingress {
 interface (all | interface-name);
 vlan (vlan-id | vlan-name);
 }
 }
 output {
 interface interface-name;
 ip-address ip-address;
 vlan (vlan-id | vlan-name);
 }
 }
 }
 bpdv-block {
 interface (all | [interface-name]);
 disable-timeout timeout;
 }
 dot1q-tunneling {
 ether-type (0x8100 | 0x88a8 | 0x9100)
 }
 interfaces interface-name {
 no-mac-learning;
 }
 mac-table-aging-time seconds {
 }
 port-error-disable {
 disable-timeout timeout;
 }
 secure-access-port {
 dhcp-snooping-file {
 location local_pathname | remote_URL;
 timeout seconds;
 write-interval seconds;
 }
 interface (all | interface-name) {
 allowed-mac {
 mac-address-list;
 }
 (dhcp-trusted | no-dhcp-trusted);
 fcoe-trusted;
 mac-limit limit action action;
 no-allowed-mac-log;
 }
 vlan (all | vlan-name) {
 (arp-inspection | no-arp-inspection) [
 forwarding-class (for DHCP Snooping or DAI Packets) class-name;
]
 }
 }
}

```

```

dhcp-option82 {
 circuit-id {
 prefix (Circuit ID for Option 82) hostname;
 use-interface-description;
 use-vlan-id;
 }
 remote-id {
 prefix (Remote ID for Option 82) hostname | mac | none;
 use-interface-description;
 use-string string;
 }
 vendor-id <string>;
}
(examine-dhcp | no-examine-dhcp) {
 forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
 examine-vn2vn {
 beacon-period milliseconds;
 }
 fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action>;
}
}
storm-control {
 interface (all | interface-name) {
 bandwidth bandwidth;
 no-broadcast;
 no-multicast;
 no-unknown-unicast;
 }
}
traceoptions {
 file filename <files number> <no-stamp> <replace> <size size> <world-readable |
no-world-readable>;
 flag flag <disable>;
}
unknown-unicast-forwarding {
 vlan (all | vlan-name) {
 interface interface-name;
 }
}
}
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure Ethernet switching options.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
  - [Overview of Access Port Protection on page 4674](#)
  - [Understanding Storm Control on page 4694](#)

## ingress (ethernet-switching-options)

---

**Syntax**     ingress {  
                   [interface](#) (all | *interface-name*);  
                   [vlan](#) (*vlan-id* | *vlan-name*);  
                   }

**Hierarchy Level**     [edit [ethernet-switching-options analyzer name input](#)]

**Release Information**     Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**     Specify the interfaces or VLANs for which incoming traffic is mirrored as part of a port mirroring configuration.

The statements are explained separately.

**Required Privilege Level**     routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding Port Mirroring on page 4887](#)
  - [Configuring Port Mirroring on page 4904](#)
  - [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)

## input

---

**Syntax**    `input {  
              ingress {  
                  interface (all | interface-name);  
                  vlan (vlan-id | vlan-name);  
              }  
              egress {  
                  interface (all | interface-name);  
              }  
          }`

**Hierarchy Level**    [edit [ethernet-switching-options analyzer name](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Define the traffic to be mirrored. The definition can be a combination of traffic entering or exiting specific ports or VLANs.

                  The statements are explained separately.

**Default**    No default.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 4887](#)
- [Configuring Port Mirroring on page 4904](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)



## interface (Port Mirroring)

|                            |                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | interface (all   <i>interface-name</i> );                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit <a href="#">ethernet-switching-options analyzer name input egress</a> ],<br>[edit <a href="#">ethernet-switching-options analyzer name input ingress</a> ],<br>[edit <a href="#">ethernet-switching-options analyzer name output</a> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                            |
| <b>Description</b>         | Specify the interfaces for which ingressing traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).                                                                                     |
| <b>Options</b>             | all—Apply port mirroring to all interfaces on the switch (except the output interface).<br>Mirroring a high volume of traffic can cause performance issues, so you should generally select specific input interfaces.                        |



**CAUTION:** Configuring **all** in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

*interface-name*—Apply port mirroring to the specified interface only.

|                                 |                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 4887</a></li> <li>• <a href="#">Configuring Port Mirroring on page 4904</a></li> <li>• <a href="#">Example: Configuring Port Mirroring for Local Analysis on page 4895</a></li> </ul> |

## ip-address (Port Mirroring)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ip-address <i>ip-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">ethernet-switching-options analyzer name output</a> ]                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the IP address to which traffic should be mirrored (the IP address of the analyzer system). The device can be on a remote network. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.) This statement is not supported on QFabric systems. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 4887</a></li><li>• <a href="#">Configuring Port Mirroring on page 4904</a></li><li>• <a href="#">Example: Configuring Port Mirroring for Local Analysis on page 4895</a></li></ul>                                                                                                                                                             |

## output

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {<br/>  <a href="#">interface</a> <i>interface-name</i>;<br/>  <a href="#">ip-address</a> <i>ip-address</i>;<br/>  <a href="#">vlan</a> (<i>vlan-id</i>   <i>vlan-name</i>);<br/>}</pre>                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">ethernet-switching-options analyzer name</a> ]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.                                                                                                                           |
| <b>Description</b>              | Configure the destination for mirrored traffic, either an interface on the switch (for local monitoring) or a VLAN (for remote monitoring).<br><br>The statements are explained separately.                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 4887</a></li><li>• <a href="#">Configuring Port Mirroring on page 4904</a></li><li>• <a href="#">Example: Configuring Port Mirroring for Local Analysis on page 4895</a></li></ul> |

## vlan (Port Mirroring)

|                                 |                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan (<i>vlan-id</i>   <i>vlan-name</i>);</code>                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">ethernet-switching-options analyzer <i>name</i> input ingress</a> ],<br>[edit <a href="#">ethernet-switching-options analyzer <i>name</i> output</a> ]                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Option <code>output vlan</code> added in Junos OS Release 12.1 for the QFX Series.                                                                                                                         |
| <b>Description</b>              | Specify that traffic entering into a VLAN should be mirrored. Configure mirrored traffic to be sent to a VLAN for remote monitoring (output).                                                                                                                                   |
| <b>Options</b>                  | <i>vlan-id</i> —Numeric VLAN identifier.<br><br><i>vlan-name</i> —Name of the VLAN.                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 4887</a></li> <li>• <a href="#">Configuring Port Mirroring on page 4904</a></li> <li>• <a href="#">Example: Configuring Port Mirroring for Local Analysis on page 4895</a></li> </ul> |

## Configuration Statements for Encryption

- [authentication-key-chains on page 4920](#)
- [cache-size on page 4921](#)
- [cache-timeout-negative on page 4922](#)
- [ca-name on page 4922](#)
- [certificates on page 4923](#)
- [certification-authority on page 4924](#)
- [crl \(Encryption Interface\) on page 4924](#)
- [encoding on page 4925](#)
- [enrollment-retry on page 4925](#)
- [enrollment-url on page 4926](#)
- [file on page 4926](#)
- [key \(Authentication Keychain\) on page 4927](#)
- [key-chain \(Security\) on page 4928](#)
- [ldap-url on page 4929](#)
- [local on page 4930](#)
- [maximum-certificates on page 4931](#)

- [path-length on page 4931](#)
- [secret on page 4932](#)
- [security on page 4933](#)
- [ssh-known-hosts on page 4934](#)
- [start-time \(Authentication Key Transmission\) on page 4935](#)
- [traceoptions on page 4937](#)

---

## authentication-key-chains

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>authentication-key-chains {<br/>  key-chain key-chain-name {<br/>    description text-string;<br/>    key key {<br/>      algorithm (md5   hmac-sha-1);<br/>      options (basic   isis-enhanced);<br/>      secret secret-data;<br/>      start-time yyyy-mm-dd.hh:mm:ss;<br/>    }<br/>    tolerance seconds;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Hierarchy Level          | [edit security]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Release Information      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| Description              | <p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the <b>authentication-key-chains</b> statement is configured at the <b>[edit security]</b> hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the <b>[edit protocols]</b> hierarchy level or with the BFD protocol using the <b>bfd-liveness-detection</b> statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2848</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## cache-size

---

|                            |                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | cache-size <i>bytes</i> ;                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the cache size for digital certificates.                                                      |
| <b>Options</b>             | <b>bytes</b> —Cache size for digital certificates.<br><b>Range:</b> 64 through 4,294,967,295<br><b>Default:</b> 2 megabytes (MB)                                                               |



**NOTE:** We recommend that you limit your cache size to 4 MB.

---

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>      |

## cache-timeout-negative

---

|                            |                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | cache-timeout-negative <i>seconds</i> ;                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure a negative cache for digital certificates.                                                    |
| <b>Options</b>             | <b>seconds</b> —Negative time to cache digital certificates, in seconds.<br><b>Range:</b> 10 through 4,294,967,295<br><b>Default:</b> 20                                                       |



**CAUTION:** Configuring a large negative cache value can lead to a denial-of-service attack.

---

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>        |

## ca-name

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ca-name <i>ca-identity</i> ;                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> ]                                                                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the certificate authority (CA) identity to use in the certificate request.                      |
| <b>Options</b>                  | <b>ca-identity</b> —CA identity to use in the certificate request.                                                                                                                             |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                        |

## certificates

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> certificates {   cache-size bytes;   cache-timeout-negative seconds;   certification-authority ca-profile-name {     ca-name ca-identity;     crt file-name;     encoding (binary   pem);     enrollment-url url-name;     file certificate-filename;     ldap-url url-name;   }   enrollment-retry attempts;   local certificate-name {     certificate-key-string;     load-key-file URL filename;   }   maximum-certificates number;   path-length certificate-path-length; } </pre> |
| <b>Hierarchy Level</b>          | [edit security]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure the digital certificates for IPsec.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                       |

## certification-authority

---

|                                 |                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>certification-authority <i>ca-profile-name</i> {<br/>    <i>ca-name</i> <i>ca-identity</i>;<br/>    <i>crl</i> <i>file-name</i>;<br/>    <i>encoding</i> (binary   pem);<br/>    <i>enrollment-url</i> <i>url-name</i>;<br/>    <i>file</i> <i>certificate-filename</i>;<br/>    <i>ldap-url</i> <i>url-name</i>;<br/>}</code> |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                       |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure a certificate authority profile name.<br><br>The remaining statements are explained separately.                                                                                                                                     |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                                                                              |

## crl (Encryption Interface)

---

|                                 |                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>crl <i>file-name</i>;</code>                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                  |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis. |
| <b>Options</b>                  | <i>file-name</i> —Specify the file from which to read the CRL.                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                                                                         |



## encoding

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | encoding (binary   pem);                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit security ike policy <i>ike-peer-address</i> ],<br>[edit security certificates <b>certification-authority</b> <i>ca-profile-name</i> ]                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file format used for the <b>local-certificate</b> and <b>local-key-pair</b> statements.     |
| <b>Options</b>                  | <b>binary</b> —Binary file format.<br><br><b>pem</b> —Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.<br><b>Default:</b> binary                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> <li>• <i>Configuring an IKE Policy for Digital Certificates for an ES PIC</i></li> </ul>   |

## enrollment-retry

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | enrollment-retry <i>attempts</i> ;                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit security <b>certificates</b> ]                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                  |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify how many times a router or switch can resend a digital certificate request. |
| <b>Options</b>                  | <b>attempts</b> —Number of enrollment retries.<br><b>Range:</b> 0 through 100<br><b>Default:</b> 0                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                  |

## enrollment-url

---

|                                 |                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>enrollment-url <i>url-name</i>;</code>                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                          |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL). |
| <b>Options</b>                  | <i>url-name</i> —Certificate authority URL.                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                                                                                 |

## file

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file <i>certificate-filename</i>;</code>                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file from which to read the digital certificate.                                            |
| <b>Options</b>                  | <i>certificate-filename</i> —File from which to read the digital certificate.                                                                                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                        |

## key (Authentication Keychain)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>key key {   algorithm (md5   hmac-sha-1);   options (basic   isis-enhanced);   secret secret-data;   start-time yyyy-mm-dd.hh:mm:ss; }</pre>                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> ]                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> |
| <b>Description</b>              | Configure the authentication element.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>key</b>—Each key within a keychain is identified by a unique integer value.</p> <p><b>Range:</b> 0 through 63</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2848</a></li> <li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647</a></li> </ul>                                                                              |

## key-chain (Security)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>keychain <i>key-chain-name</i> {<br/>  description <i>text-string</i>;<br/>  key <i>key</i> {<br/>    algorithm (md5   hmac-sha-1);<br/>    options (basic   isis-enhanced);<br/>    secret <i>secret-data</i>;<br/>    start-time <i>yyyy-mm-dd.hh:mm:ss</i>;<br/>  }<br/>  tolerance <i>seconds</i>;<br/>}</pre>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains]                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| <b>Description</b>              | Create the key-chain configuration for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.                                                                                                                                                                                 |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">authentication-key-chains on page 4920</a></li><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2848</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647</a></li></ul>            |

---

## ldap-url

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <ldap-url <i>url-name</i> >;                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>(Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.       |
| <b>Options</b>                  | <i>url-name</i> —Name of the LDAP URL.                                                                                                                                                         |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>                                                                                        |

## local

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>local <i>certificate-name</i> {<br/>    <i>certificate-key-string</i>;<br/>    load-key-file <i>URL filename</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>certificate-namecertificate-key-string</i></b>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><b><i>certificate-name</i></b>—Name that uniquely identifies the certificate.</p> <p><b><i>load-key-file URL filename</i></b>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"><li>• Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)</li><li>• URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)</li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## maximum-certificates

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | maximum-certificates <i>number</i> ;                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the maximum number of peer digital certificates to be cached.                                 |
| <b>Options</b>                  | <b>number</b> —Maximum number of peer digital certificates to be cached.<br><b>Range:</b> 64 through 4,294,967,295 peer certificates<br><b>Default:</b> 1024 peer certificates                 |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                      |

## path-length

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | path-length <i>certificate-path-length</i> ;                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the digital certificate path length.                                                          |
| <b>Options</b>                  | <b>certificate-path-length</b> —Digital certificate path length.<br><b>Range:</b> 2 through 15 certificates<br><b>Default:</b> 15 certificates                                                 |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>                                                                                      |

## secret

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>secret <i>secret-data</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> |
| <b>Description</b>              | Specify a password in encrypted text or plain text format. The secret password always appears in encrypted format.                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b><i>secret-data</i></b> —Password to use; it can include spaces if the character string is enclosed in quotation marks.                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2848</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647</a></li></ul>                                                                                  |



## security

```
Syntax security {
 authentication-key-chains {
 key-chain key-chain-name {
 key key {
 secret secret-data;
 start-time yyyy-mm-dd.hh:mm:ss;
 }
 }
 }
 certificates {
 cache-size bytes;
 cache-timeout-negative seconds;
 certification-authority ca-profile-name {
 ca-name ca-identity;
 crl file-name;
 encoding (binary | pem);
 enrollment-url url-name;
 file certificate-filename;
 ldap-url url-name;
 }
 enrollment-retry attempts;
 local certificate-filename {
 certificate-key-string;
 load-key-file key-file-name;
 }
 maximum-certificates number;
 path-length certificate-path-length;
 }
 ssh-known-hosts {
 host {
 fetch-from-server host-name;
 load-key-file file-name;
 }
 }
 traceoptions {
 file filename <files number> <size size>;
 flag flag;
 level level;
 no-remote-trace
 }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure security services. Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

Required Privilege  
Level

Related  
Documentation

## ssh-known-hosts

---

**Syntax**    ssh-known-hosts {  
              host *host-name* {  
                  fetch-from-server *host-name*;  
                  load-key-file *file-name*;  
              }  
              }

**Hierarchy Level**    [edit security ssh-known-hosts]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure SSH support for known hosts and for administering SSH host key updates.

**Options**    **host *host-name***—Hostname of the SSH known host entry. This option has the following suboptions:

- **fetch-from-server *host-name***—Retrieve SSH public host key information from a specified server.
- **load-key-file *filename***—Import SSH host key information from the `/var/tmp/ssh-known-hosts` file.

**Required Privilege Level**    admin—To view this statement in the configuration.  
                                  admin-control—To add this statement to the configuration.

Related  
Documentation


## start-time (Authentication Key Transmission)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>start-time (now   yyyy-mm-dd.hh:mm:ss);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify a start time for key transmission. You do not need to specify an end time for the key. If a new key is present with a new start time, the keychain rolls over to the new one. The start time must be unique within the keychain.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>now</b>—Start time as the current year, month, day, hour, minute, and second.</p> <p><b>daydays</b>—Start time as the specified number of days after the current day. For example, if the current day is the 12th and you configure <b>start-time 2day</b>, the start time will be on the 14th, exactly two days after the configuration is entered.</p> <p><b>hourhours</b>—Start time as the specified number of hours after the current hour. For example, if the current hour is 9:00 and you configure <b>start-time 3hour</b>, the start time will be in 12:00, exactly three hours after the configuration is entered.</p> <p><b>minuteminutes</b>—Start time as the specified number of minutes after the current minute. For example, if the current minute is 27 minutes after the hour and you configure <b>start-time 5min</b>, the start time will be in 32 minutes after the hour, exactly five minutes after the configuration is entered.</p> <p><b>monthmonths</b>—Start time as the specified number of months after the current month. For example, if the current month is March and you configure <b>start-time 4month</b>, the start time will be in July, exactly four months after the configuration is entered.</p> <p><b>secondseconds</b>—Start time as the specified number of seconds after the current second. For example, if the current second is 10:20:40 and you configure <b>start-time 10seconds</b>, the start time will be 10:20:50, exactly 10 seconds after the configuration is entered.</p> <p><b>yearyears</b>—Start time as the specified number of years after the current year. For example, if the current year is 2011 and you configure <b>start-time 1year</b>, the start time will be in 2012, exactly one year after the configuration is entered.</p> <p><b>yyyy-mm-dd.hh:mm:ss</b>—Start time in UTC (Coordinated Universal Time). The start time must be unique within the keychain.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Related  
Documentation**

- *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*
- [Example: Configuring BFD Authentication for Static Routes on page 2848](#)
- [Example: Configuring BFD Authentication for Static Routes on page 2848](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3647](#)

## traceoptions

|                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                             | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt;;     flag all;     flag certificates;     flag database;     flag general;     flag ike;     flag parse;     flag policy-manager;     flag routing-socket;     flag timer;     level     no-remote-trace } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                    | <p>[edit security],<br/>[edit services ipsec-vpn]</p> <p>Trace options can be configured at either the [edit security] or the [edit services ipsec-vpn] hierarchy level, but not at both levels.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                                                                                                                                                                                                                                | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                        | <p>Configure security trace options.</p> <p>To specify more than one trace option, include multiple <b>flag</b> statements. Trace option output is recorded in the <code>/var/log/kmd</code> file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <div style="display: flex; align-items: center;">  <p><b>NOTE:</b> The <code>traceoptions</code> statement is not supported on QFabric systems.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                                                            | <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file (for example, <code>kmd</code>) reaches its maximum size, it is renamed <code>kmd.0</code>, then <code>kmd.1</code>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 0 files</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, <code>kmd</code>) reaches this size, it is renamed, <code>kmd.0</code>, then <code>kmd.1</code> and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Default:</b> 1024 KB</p> |

**flag *flag***—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- **all**—Trace all security events.
- **certificates**—Trace certificate events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

**level *level***—(Optional) Set traceoptions level.

- **all**—match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**no-remote-trace**—(Optional) Disable remote tracing

|                           |                                                           |
|---------------------------|-----------------------------------------------------------|
| <b>Required Privilege</b> | admin—To view the configuration.                          |
| <b>Level</b>              | admin-control—To add this statement to the configuration. |

|                              |                                                               |
|------------------------------|---------------------------------------------------------------|
| <b>Related Documentation</b> | • <i>Configuring Tracing Operations for Security Services</i> |
|------------------------------|---------------------------------------------------------------|

---

## Configuration Statements for DHCP Relay

---

- [apply-secondary-as-giaddr](#) on page 4939
- [bootp](#) on page 4940
- [broadcast](#) on page 4941
- [client-response-ttl](#) on page 4941
- [description \(Forwarding Options\)](#) on page 4942
- [interface \(BOOTP\)](#) on page 4943
- [maximum-hop-count](#) on page 4944

- [minimum-wait-time](#) on page 4944
- [no-listen](#) on page 4945
- [server \(DHCP and BOOTP Relay Agent\)](#) on page 4945

## [apply-secondary-as-giaddr](#)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>apply-secondary-as-giaddr;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers <a href="#">bootp</a> ],<br>[edit forwarding-options helpers <a href="#">bootp interface</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configures the interfaces on a switch that are DHCP relay agents to be enabled for smart DHCP relay:</p> <ul style="list-style-type: none"> <li>• When you configure this statement directly under the <b>bootp</b> statement, it enables smart relay on all the interfaces that are relay agents.</li> <li>• When you configure this statement under the <b>interface</b> statement, it enables smart relay on the specified interface.</li> </ul> <p>Smart relay requires the interfaces to be routed VLAN interfaces or Layer 3 logical interfaces that have multiple IP addresses.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DHCP and BOOTP Relay</a> on page 4907</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## bootp

---

**Syntax**    bootp {  
              client-response-ttl *number*;  
              description *text-description*;  
              apply-secondary-as-giaddr  
              interface (*interface-name* | *interface-group*) {  
                  broadcast *number*;  
                  client-response-ttl *number*;  
                  description *text-description*;  
                  maximum-hop-count *number*;  
                  minimum-wait-time *seconds*;  
                  no-listen;  
                  server *address* ;  
                  apply-secondary-as-giaddr  
              }  
              maximum-hop-count *number*;  
              minimum-wait-time *seconds*;  
              server *address* {  
              }  
              }  
              }

**Hierarchy Level**    [edit forwarding-options helpers]

**Release Information**    Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Configure a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent.

**Options**    The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring DHCP and BOOTP Relay on page 4907](#)



## broadcast

---

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>broadcast <i>number</i>;</code>                                                                                   |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp <a href="#">interface</a> ( <i>interface-name</i>   <i>interface-group</i> )]    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                                  |
| <b>Description</b>              | If the specified interface is unavailable, broadcast DHCP and BOOTP packets.                                            |
| <b>Options</b>                  | None                                                                                                                    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DHCP and BOOTP Relay on page 4907</a></li> </ul>       |

## client-response-ttl

---

|                                 |                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>client-response-ttl <i>number</i>;</code>                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp],<br>[edit forwarding-options helpers bootp <a href="#">interface</a> ( <i>interface-name</i>   <i>interface-group</i> )] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                          |
| <b>Description</b>              | Set the IP time-to-live (TTL) value in DHCP response packets sent to a DHCP client.                                                                              |
| <b>Options</b>                  | <i>number</i> —Decrement amount.<br><b>Default:</b> None                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</a></li> </ul>                 |

## description (Forwarding Options)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text-description;</code>                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp],<br>[edit forwarding-options helpers bootpinterface (interface-name   interface-group)],<br>[edit forwarding-options helpers domain],<br>[edit forwarding-options helpers domain interface interface-name],<br>[edit forwarding-options helpers tftp],<br>[edit forwarding-options helpers tftpinterface interface-name] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                                                                                                                                              |
| <b>Description</b>              | Describe a BOOTP, DHCP, Domain Name System (DNS), or Trivial File Transfer Protocol (TFTP) service, or an interface that is configured for the service.                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring DNS and TFTP Packet Forwarding</i></li><li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li></ul>                                                                                                                                                                |

## interface (BOOTP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface (<i>interface-name</i>   <i>interface-group</i>) {   broadcast;   client-response-ttl <i>number</i>;   description <i>text-description</i>;   maximum-hop-count <i>number</i>;   minimum-wait-time <i>seconds</i>;   no-listen;   server <i>address</i> {     logical-system <i>logical-system-name</i> &lt;routing-instance [ &lt;default&gt;       <i>routing-instance-names</i> ]&gt;;     routing-instance [ &lt;default&gt; <i>routing-instance-names</i> ];   }   apply-secondary-as-giaddr (QFX platforms only) }</pre> |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify the interface for a DHCP and BOOTP relay agent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>interface-group</i></b>—Sets a logical interface or group of logical interfaces with a specific DHCP relay configuration.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</a></li> <li>• <a href="#">Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4773</a></li> </ul>                                                                                                                                                                                                                                            |

## maximum-hop-count

---

|                                 |                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | maximum-hop-count <i>number</i> ;                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp],<br>[edit forwarding-options helpers bootp <code>interface</code> ( <i>interface-name</i>   <i>interface-group</i> )]                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches. |
| <b>Description</b>              | Specify the maximum number of hops allowed.                                                                                                                                                         |
| <b>Options</b>                  | <i>number</i> —Maximum number of hops.<br><b>Default:</b> 4 hops                                                                                                                                    |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li></ul>                                                               |

## minimum-wait-time

---

|                                 |                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-wait-time <i>seconds</i> ;                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp],<br>[edit forwarding-options helpers bootp <code>interface</code> ( <i>interface-name</i>   <i>interface-group</i> )]                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches. |
| <b>Description</b>              | Specify the minimum time allowed.                                                                                                                                                                   |
| <b>Options</b>                  | <i>seconds</i> —Minimum time.<br><b>Default:</b> 0 seconds                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li></ul>                                                               |

## no-listen

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-listen;                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp <b>interface</b> ( <i>interface-name</i>   <i>interface-group</i> )],<br>[edit forwarding-options helpers domain interface <i>interface-name</i> ],<br>[edit forwarding-options helpers tftp interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                                                   |
| <b>Description</b>              | Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring DNS and TFTP Packet Forwarding</i></li> <li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li> </ul>                                                                  |

## server (DHCP and BOOTP Relay Agent)

|                                 |                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>server address {     logical-system <i>logical-system-name</i> &lt;routing-instance [ &lt;default&gt;         <i>routing-instance-names</i> ]&gt;;     routing-instance [ &lt;default&gt; <i>routing-instance-names</i> ]; }</pre>                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit forwarding-options helpers bootp],<br>[edit forwarding-options helpers bootp <b>interface</b> ( <i>interface-name</i>   <i>interface-group</i> )]                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                                                                                                                                          |
| <b>Description</b>              | Configure the router or switch to act as a DHCP and BOOTP relay agent.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li>• <b>address</b>—One or more addresses of the server.</li> <li>• <b>logical-system <i>logical-system-name</i></b>—(Optional) Logical system of the server.</li> <li>• <b>routing-instance <i>routing-instance-names</i></b>—(Optional) Routing instance name that belong to the DHCP or BOOTP relay agent.</li> </ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li> </ul>                                                                                                                                                                                                                      |



## CHAPTER 57

# Administration

- [Monitoring Commands for Port Mirroring on page 4947](#)

### Monitoring Commands for Port Mirroring

---

- `show analyzer`

## show analyzer

|                                 |                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show analyzer</b> < <i>analyzer-name</i> >                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                  |
| <b>Description</b>              | Display information about port mirroring.                                                                                                                                        |
| <b>Options</b>                  | <i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer (port-mirroring configuration).                                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show analyzer on page 4948</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 385 on page 4948</a> describes the output fields for the <b>show analyzer</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 385: show analyzer Output Fields**

| Field Name                          | Field Description                                                                                                                 |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Analyzer name</b>                | Name of the analyzer.                                                                                                             |
| <b>Output interface</b>             | Local interface to which mirror packets are sent. If you configure an output interface, you cannot also configure an output VLAN. |
| <b>Output VLAN</b>                  | VLAN to which mirror packets are sent. If you configure an output VLAN, you cannot also configure an output interface.            |
| <b>Egress monitored interfaces</b>  | Interfaces for which egress traffic is mirrored.                                                                                  |
| <b>Ingress monitored interfaces</b> | Interfaces for which ingress traffic is mirrored.                                                                                 |
| <b>Ingress monitored VLANs</b>      | VLANs for which ingress traffic is mirrored.                                                                                      |

## Sample Output

### show analyzer

```

user@switch> show analyzer
Analyzer name : employee-monitor
Output interface : ge-0/0/10.0
Output VLAN : remote-analyzer
Egress monitored interfaces : ge-0/0/7.0
Ingress monitored interfaces : ge-0/0/8.0
Ingress monitored interfaces : ge-0/0/9.0

```



## CHAPTER 58

# Troubleshooting

- [Troubleshooting Procedures on page 4949](#)

## Troubleshooting Procedures

---

- [Troubleshooting Port Mirroring on page 4949](#)

## Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 4949](#)
- [Egress Port Mirroring with VLAN Translation on page 4951](#)
- [Egress Port Mirroring with Private VLANs on page 4951](#)

## Port Mirroring Constraints and Limitations

---

- [Local and Remote Port Mirroring on page 4949](#)
- [Remote Port Mirroring Only on page 4951](#)

### ***Local and Remote Port Mirroring***

The following constraints and limitations apply to local and remote port mirroring with the QFX Series:

- You can create a total of four port-mirroring configurations on a QFX Series standalone switch.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a QFX3500 device or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### ***Remote Port Mirroring Only***

The following constraints and limitations apply to remote port mirroring with the QFX Series:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

### **Egress Port Mirroring with VLAN Translation**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

### **Egress Port Mirroring with Private VLANs**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Port Mirroring on page 4887](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 4900](#)

## PART 19

# Storage

- [Overview on page 4955](#)
- [Configuration on page 5099](#)
- [Administration on page 5269](#)
- [Troubleshooting on page 5403](#)



## CHAPTER 59

# Overview

- [Software Features Overview on page 4955](#)
- [Fibre Channel, FCoE, and FIP on page 4961](#)
- [QFabric Specific on page 5084](#)

### **Software Features Overview**

---

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Overview of FIP on page 4960](#)

## Overview of Fibre Channel on the QFX Series

Fibre Channel (FC) is a high-speed network technology that interconnects network elements and allows them to communicate with one another. The International Committee for Information Technology Standards (INCITS) T11 Technical Committee sets FC standards.

FC networks provide high-performance characteristics such as lossless transport combined with flexible network topology. FC is primarily used in storage area networks (SANs) because it provides reliable, lossless, in-order frame transport between initiators and targets. FC components include initiators, targets, and FC-capable switches that interconnect FC devices and may also interconnect FC devices with Fibre Channel over Ethernet (FCoE) devices. Initiators originate I/O commands. Targets receive I/O commands. For example, a server can initiate an I/O request to a storage device target.

Juniper Networks QFX Series can function as an FCoE-FC gateway or as an FCoE transit switch.

FCoE transports native FC frames over an Ethernet network by encapsulating the unmodified frames in Ethernet. It also provides protocol extensions to discover FCoE devices through the Ethernet network. FCoE requires that the Ethernet network support data center bridging (DCB) extensions that ensure lossless transport and allow the Layer 2 Ethernet domain to meet the requirements of FC transport.

The FCoE-FC gateway functionality is a licensed feature on the QFX Series. As an FCoE-FC gateway, the switch connects FCoE devices on an Ethernet network to a SAN FC switch.

You do not need a license to use the switch as an FCoE transit switch. As an FCoE transit switch, the switch:

- Is a Layer 2 data center bridging (DCB) switch that can transport FCoE frames.
- Implements FCoE Initialization Protocol (FIP) snooping.
- Connects multiple FCoE endpoints to the FC network.



**NOTE:** The QFX3600 device does not support FC or FCoE features.

---

This topic describes:

- [Fibre Channel Transport Protocol on page 4956](#)
- [How FC Works on the QFX Series on page 4957](#)
- [Supported FC Features and Functions on page 4959](#)
- [Lossless Transport Support on page 4959](#)

### Fibre Channel Transport Protocol

The Fibre Channel Protocol is a transport protocol that consists of five layers as shown in [Table 386 on page 4957](#):



**Table 386: Fibre Channel Protocol Layers**

| FC Protocol Layer | Description                                |
|-------------------|--------------------------------------------|
| FC-0              | Physical (cabling, connectors, and so on)  |
| FC-1              | Data link layer                            |
| FC-2              | Network layer (defines the main protocols) |
| FC-3              | Common services                            |
| FC-4              | Protocol mapping                           |

The FC protocol layers are generally split into three groups:

- FC-0 and FC-1 are the physical layers.
- FC-2 is the protocol layer, similar to OSI Layer 3.
- FC-3 and FC-4 are the services layers.

The FCoE-FC gateway operates the physical layers and the protocol layer, and provides FIP and service redirection at the services layer.

#### How FC Works on the QFX Series

The switch connects devices that support FC and Ethernet (such as FCoE servers on an Ethernet network) to an FC SAN, thus converging the Ethernet and FC networks on a single physical network infrastructure. The switch provides the class-of-service (CoS) features needed to handle the different types of traffic appropriately.

To converge FC and Ethernet networks, you can configure the switch as an:

- [FCoE-FC Gateway on page 4957](#)
- [FCoE Transit Switch on page 4958](#)
- [FCoE VLANs on page 4958](#)

#### **FCoE-FC Gateway**

When the switch functions as an FCoE-FC gateway, the switch aggregates FCoE traffic and performs the encapsulation and de-encapsulation of native FC frames in Ethernet as it transports the frames between FCoE devices in the Ethernet network and the FC switch. In effect, the switch translates Ethernet to FC and FC to Ethernet.

The gateway receives FC frames encapsulated in Ethernet from FCoE devices through an FCoE VLAN interface composed of one or more 10-Gigabit Ethernet interfaces. The gateway removes the Ethernet encapsulation from the FC frames, and then sends the native FC frames to the FC switch through a native FC interface.

The gateway receives native FC frames from the FC switch on the gateway's native FC interfaces. The gateway encapsulates the native FC frames in Ethernet, and then sends the encapsulated frames to the appropriate FCoE device through the FCoE VLAN interface.

To FCoE devices, the gateway behaves like an FC switch and can present multiple virtual F\_Ports (VF\_Ports) on a single interface. To an FC switch, the gateway behaves like an FC node that is doing N\_Port ID virtualization (NPIV).

### **FCoE Transit Switch**

When the switch functions as an FCoE transit switch, it forwards traffic (including FCoE traffic) based on Layer 2 media access control (MAC) forwarding and is a normal DCB-enabled Layer 2 switch that also performs FIP snooping. The switch aggregates FCoE traffic and passes it through to an FCF. The switch does not remove the Ethernet encapsulation from the FC frames, but it does preserve the class of service (CoS) required to transport FC frames.

The switch inspects (snoops) FIP information in order to create filters that permit only valid FCoE traffic to flow through the switch between FCoE devices and the FCF. The switch does not use native FC ports because the FC frames are encapsulated in Ethernet when they flow between the FCoE devices and the FCF. Virtual point-to-point links between each FCoE device and the FCF pass transparently through the switch, so the switch is not seen as a terminating point or an intermediate point by FCoE devices or by the FCF.

### **FCoE VLANs**

On the QFX Series, all FCoE traffic must travel in a VLAN dedicated to transporting only FCoE traffic. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE or FIP traffic must travel in a different VLAN.



**NOTE:** On a QFX3500 or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.

---



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features. IGMP snooping is enabled by default on all VLANs; be sure to disable IGMP snooping on FCoE VLANs.

---

You can configure more than one FCoE VLAN, but any given virtual link must be in only one FCoE VLAN.



**NOTE:** All 10-Gigabit Ethernet interfaces that connect to FCoE devices must have a native VLAN configured in order to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.

---



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols.

---

Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

---

## Supported FC Features and Functions

---

QFX Series supports the following features and functionality:

- As an FCoE-FC gateway:
  - DCB, including Data Center Bridging Capability Exchange protocol (DCBX), priority-based flow control (PFC), enhanced transmission service (ETS), and 10-Gigabit Ethernet interfaces
  - FCoE Initialization Protocol (FIP)
  - Proxy for FCoE devices when communicating with FC switches and acts as a proxy for FC switches when communicating with FCoE devices
  - Up to 12 native FC interfaces per QFX3500 switch (each interface can be configured as a 2-Gigabit, 4-Gigabit, or 8-Gigabit Ethernet interface)
- As an FCoE transit switch:
  - DCB functions
  - FIP snooping
  - Transparent Layer 2 MAC forwarding of FCoE frames

## Lossless Transport Support

---

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the QFX3500 or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.

### Related Documentation

- [Understanding Fibre Channel on page 4961](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE Transit Switch Functionality on page 4972](#)
- [Understanding FCoE on page 4968](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Overview of FIP on page 4960](#)

- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

## Overview of FIP

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) is a Layer 2 protocol that establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices such as server FCoE Nodes (ENodes) and FC switches. FIP can also establish and maintain virtual links between FCoE devices and an FCoE-FC gateway such as the QFX Series, where the gateway acts on behalf of the FC switch.

FIP enables FCoE devices to discover one another and to initialize and maintain virtual links over a physical Ethernet network. This allows FCoE devices in the Ethernet network to access storage devices in the FC storage area network (SAN).

FIP solves the problem presented by the FC requirement for point-to-point connections (FC does not permit point-to-multipoint connections) by creating a unique virtual link for each connection between an ENode VN\_Port and an FC switch VF\_Port. Multiple virtual links can use a single physical link and virtual links can traverse Ethernet transit (passthrough) switches while appearing to be direct point-to-point connections to the FC switch.

FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic. FIP operations occur on a per-VLAN basis.

For more details about FIP, see the Technical Committee T11 organization document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* available at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf>.

### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding Fibre Channel on page 4961](#)
- [Understanding FIP Functions on page 4984](#)
- [Understanding FIP Implementation on page 4988](#)
- [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)
- [Understanding Fibre Channel Virtual Links on page 4995](#)
- [Understanding FCoE on page 4968](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

---

## Fibre Channel, FCoE, and FIP

---

- [Understanding Fibre Channel on page 4961](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding FCoE on page 4968](#)
- [Understanding FCoE Transit Switch Functionality on page 4972](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
- [Understanding FCoE and FIP Session High Availability on page 4981](#)
- [Understanding FIP Functions on page 4984](#)
- [Understanding FIP Implementation on page 4988](#)
- [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)
- [Understanding Fibre Channel Virtual Links on page 4995](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on page 5024](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Understanding MC-LAGs on an FCoE Transit Switch on page 5047](#)
- [Understanding DCBX on page 5051](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

### Understanding Fibre Channel

Fibre Channel (FC) is a serial I/O interconnect network technology capable of supporting multiple protocols. It is used primarily for storage area networks (SANs). The committee standardizing FC is the International Committee for Information Technology Standards (INCITS).

When configured as a Fibre Channel over Ethernet (FCoE)-FC gateway, the QFX Series supports the transport of native FC traffic between FC switches and the gateway's native FC interfaces.

FC concepts include:

- [FC Fabrics on page 4962](#)
- [FC Port Types on page 4962](#)
- [FC Switches on page 4962](#)
- [Adapters on page 4963](#)

- [N\\_Port ID Virtualization \(NPIV\) on page 4963](#)
- [FC Services on page 4963](#)

---

## FC Fabrics

An FC fabric is a switched network topology that interconnects FC devices using FC switches, usually to create a SAN. An FC switch is a Layer 3 network switch that is compatible with the FC protocol, forwards FC traffic, and provides FC services to the components of the FC fabric. FC devices are usually servers or storage devices such as disk arrays.

Switches called FCoE forwarders (FCFs) perform a subset of FC switch functions. An FCF is a Layer 3 network switch that is compatible with the FC protocol and forwards FC traffic, but does not provide network services.

When configured as an FCoE-FC gateway, the QFX Series acts as a proxy for the FCF functionality of an FC switch. The gateway provides FCoE devices on the Ethernet network access to the FC network without requiring the FC switches in the SAN to support Ethernet interfaces. The gateway is not an FCF and does not provide FC services.

FC network design often uses two fabrics (dual-rail topology) for redundancy. The two fabrics connect to edge devices but are otherwise unconnected, so that if one fabric goes down, the other fabric can continue to provide connectivity.

---

## FC Port Types

The QFX Series supports the following FC port types:

- **N\_Port**—An N\_Port is a port on the node of an FC device such as a server or a storage device and is also known as a node port.
- **F\_Port**—An F\_Port is a port on an FC switch that connects to an FC device N\_Port in a point-to-point connection. F\_Ports are also known as fabric ports.

These port types are a subset of the existing FC port types that can be supported in an FC fabric.

---

## FC Switches

FC switches provide FC services to the FC network. FC switches forward Layer 3 traffic. They may transport a combination of native FC traffic and other traffic, such as Internet Small Computer Systems Interface (iSCSI) or FCoE, or they may transport only native FC traffic. When an FC switch supports FCoE, it combines FCoE termination functions with the FC stack on an FC switching element. This is also known as a dual-stack switch.

When FC switches support FCoE, they present virtual FC interfaces in the form of virtual F\_Ports (VF\_Ports) to the FCoE nodes (ENodes) on FCoE devices. A VF\_Port is an endpoint in a virtual point-to-point connection with an ENode virtual N\_Port (VN\_Port). A VF\_Port emulates a native FC F\_Port and performs similar functions. A VF\_Port is an intermediate port in a connection between an FCoE device such as a server in the Ethernet network and a storage device in the FC SAN.

FC switches that support FCoE contain at least one lossless Ethernet media access controller (MAC) paired with an FCoE controller. The lossless Ethernet MAC implements Ethernet extensions to avoid frame loss due to congestion. The FCoE controller instantiates and terminates virtual port instances as they are needed. Each VF\_Port instance has one unique virtual link to an ENode VN\_Port.

FCoE support also requires one FCoE Link End Point (LEP) for each VF\_Port connection. An FCoE LEP is a virtual FC interface mapped onto the physical Ethernet interface. It transmits and receives FCoE frames on the virtual link, and handles FC frame encapsulation for traffic going from the FC switch to the FCoE device and frame de-encapsulation of traffic received from the FCoE device.

When you configure the QFX Series as an FCoE-FC gateway, the gateway performs these FC-to-Ethernet and Ethernet-to-FC conversion functions so that the FC switch does not need Ethernet (FCoE) ports.

---

### Adapters

FC host bus adapters (HBAs) in FC switches and devices perform functions similar to those of Ethernet adapters in Ethernet switches and devices. Switches that perform FCoE functions and FCoE devices have converged network adapters (CNAs) that support both native FC and Ethernet functionality.

---

### N\_Port ID Virtualization (NPIV)

FC requires a unique point-to-point link between the FC switch (F\_Port) and each host N\_Port. In order to avoid using one physical link for each F\_Port to N\_Port connection, the port connections must be virtualized so that they can share a physical link while maintaining logical separation.

FC accomplishes this by enabling you to create an independent virtual link for each FC session by mapping each session to a virtualized N\_Port. This process is called N\_Port ID virtualization (NPIV).

NPIV makes each virtual link look like a dedicated point-to-point link. In this way, multiple FC devices and multiple applications or virtual machines (VMs) on a single FC device can connect to an FC switch using one physical port instead of using a physical port for each connection. The virtual link creates a secure boundary between traffic from different sources on a single physical connection.

NPIV works by creating a unique virtual port identifier for each logical connection on a physical port. Conceptually, this is similar to splitting a single physical interface into multiple logical interfaces or subinterfaces. A virtual port identifier consists of the port's unique worldwide name (WWN) combined with a Fibre Channel ID (FCID) that the FC switch assigns to the virtual connection. This creates a virtual host bus adapter (HBA) for each virtual link that uniquely identifies the link to the FC switch.

---

### FC Services

When you configure the QFX Series as an FCoE-FC gateway, the gateway connects FCoE devices in the Ethernet network to the FC fabric. The gateway does not provide FC services

directly. The gateway logs in to the FC fabric and obtains FC services from the FC fabric, including:

- Management servers
  - Zone server—Defines which devices can connect to each other in the FC fabric.
  - Fabric configuration server—Discovers FC fabric topology and attributes.
  - Policy server—Distributes the rules for administering, managing, and controlling access to FC fabric resources.
  - HBA management server—Registers HBA information with the FC fabric.
- Domain manager—Allocates domain IDs to virtual switches.
- Fabric login server—Provides login services to the gateway so that the native FC ports on the gateway can perform initial fabric login (FLOGI) to the FC fabric and subsequent fabric discovery (FDISC) logins for the physical and virtual ports on the FCoE devices in the Ethernet network. This includes allocating Fibre Channel IDs (FCIDs) to ports.
- Name server—Discovers, registers, and unregisters N\_Port attributes, including the attributes of the native FC ports on the gateway that connect to the FC fabric.
- Event server—Validates incoming events to ensure transaction integrity.
- Time server—Maintains a common time for devices in the FC fabric.
- Fabric controller
  - Fabric Shortest Path First (FSPF)—The FC fabric provides link-state path selection to the gateway.
  - State change notification (SCN) / registered state change notification server (RSCN)—Notifies the appropriate nodes when new devices come online, when other nodes fail, or when changes on an online node affect system operation.

**Related  
Documentation**

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding FCoE on page 4968](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding Fibre Channel Terminology on page 5074](#)



## Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series supports the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 4965](#)
- [ETS on page 4966](#)
- [DCBX on page 4967](#)

### Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 4966](#)
- [Buffer Management on page 4966](#)
- [Physical Interfaces on page 4966](#)

### **PFC**

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

### **Buffer Management**

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

### **Physical Interfaces**

The switch supports 10-Gbps, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps (or faster) Ethernet interfaces.

---

### **ETS**

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).
- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.

- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict priority flows.

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

## DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding FCoE on page 4968](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding DCBX on page 5051](#)
- [Understanding Fibre Channel Terminology on page 5074](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)

## Understanding FCoE

Fibre Channel over Ethernet (FCoE) is a method of supporting converged Fibre Channel (FC) and Ethernet traffic on a data center bridging (DCB) network. FCoE encapsulates unmodified FC frames in Ethernet to transport the FC frames over a physical Ethernet network. The T11 Technical Committee, which is the International Committee for Information Technology Standards (INCITS) committee responsible for FC interfaces, developed the FCoE standard to provide a method for transporting FC frames over a DCB network. The T11 document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf> provides details about the FCoE version 1 standard.



**NOTE:** The switch does not support T11 Annex F *FCoE Pre-FIP Virtual Link Instantiation Protocol*.

To the Ethernet network, an FCoE frame is the same as any other Ethernet frame because the Ethernet encapsulation provides the header information needed to forward the frames. However, to achieve the lossless behavior that FC transport requires, the Ethernet network must conform to DCB standards.

DCB standards create an environment over which FCoE can transport native FC traffic encapsulated in Ethernet while preserving the mandatory class of service (CoS) and other characteristics that FC traffic requires.

Supporting FCoE in a DCB network requires that the FCoE devices in the Ethernet network and the FC switches at the edge of the SAN network handle both Ethernet and native FC traffic. To handle Ethernet traffic, an FC switch does one of two things:

- Incorporates FCoE interfaces.
- Uses an FCoE-FC gateway such as a QFX Series to de-encapsulate FCoE traffic from FCoE devices into native FC and to encapsulate native FC traffic from the FC switch into FCoE and forward it to FCoE devices through the Ethernet network.

FCoE concepts include:

- [FCoE Devices on page 4968](#)
- [FCoE Frames on page 4970](#)
- [Virtual Links on page 4970](#)
- [FCoE VLANs on page 4971](#)

---

### FCoE Devices

Each FCoE device has a converged network adapter (CNA) that combines the functions of an FC host bus adapter (HBA) and a lossless Ethernet network interface card (NIC) with 10-Gbps Ethernet ports. The portion of the CNA that handles FCoE traffic is called an FCoE Node (ENode). An ENode combines FCoE termination functions and the client part of the FC stack on the CNA.

ENodes present virtual FC interfaces to FC switches in the form of virtual N\_Ports (VN\_Ports). A VN\_Port is an endpoint in a virtual point-to-point connection called a virtual link. The other endpoint of the virtual link is an FC switch (or FCF) port. A VN\_Port emulates a native FC N\_Port and performs similar functions: handling the creation, detection, and flow of messages to and from the FC switch. A single ENode can host multiple VN\_Ports. Each VN\_Port has a separate, unique virtual link with a FC switch.

ENodes contain at least one lossless Ethernet media access controller (MAC). Each Ethernet MAC is paired with an FCoE controller. The lossless Ethernet MAC is a full-duplex Ethernet MAC that implements Ethernet extensions to avoid frame loss due to congestion and supports frames of at least 2500 bytes. The FCoE controller instantiates and terminates VN\_Port instances dynamically as they are needed for FCoE sessions. Each VN\_Port instance has a unique virtual link to an FC switch.



**NOTE:** A *session* is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.

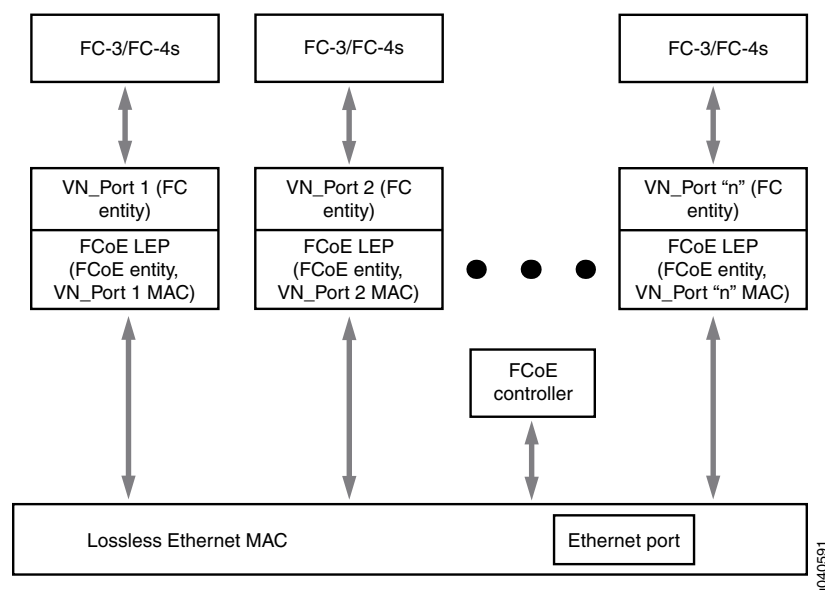
ENodes also contain one FCoE link end point (LEP) for each VN\_Port connection. An FCoE LEP is a virtual FC interface mapped onto the physical Ethernet interface.

An FCoE LEP:

- Transmits and receives FCoE frames on the virtual link.
- Handles FC frame encapsulation for traffic going from the server to the FC switch.
- Performs frame de-encapsulation of traffic received from the FC switch.

Figure 186 on page 4969 shows a block diagram of the major ENode components.

**Figure 186: ENode Components**



## FCoE Frames

---

The FCoE protocol specification replaces the FC0 and FC1 layers of the FC stack with Ethernet, but retains the FC frame header. Retaining the FC frame header enables the FC frame to pass directly to a native FC SAN after de-encapsulation. The FCoE header carries the FC start of file (SOF) bits and end of file (EOF) bits in an encoded format. FCoE supports two frame types, control frames and data frames. FCoE Initialization Protocol (FIP) carries all of the discovery and fabric login frames.

FIP control frames handle FCoE device discovery, initializing communication, and maintaining communication. They do not carry a data payload. FIP has its own EtherType (0x8914) to distinguish FIP traffic from FCoE traffic and other Ethernet traffic. To establish communication, the ENode uses the globally unique MAC address assigned to it by the CNA manufacturer.

After FIP establishes a connection between FCoE devices, the FCoE data frames handle the transport of the FC frames encapsulated in Ethernet. FCoE also has its own EtherType (0x8906) to distinguish FCoE frames from other Ethernet traffic and ensure the in-order frame handling that FC requires. FCoE frames include:

- 2112 bytes FC payload
- 24 bytes FC header
- 14 bytes standard Ethernet header
- 14 bytes FCoE header
- 8 bytes cyclic redundancy check (CRC) plus EOF
- 4 bytes VLAN header
- 4 bytes frame check sequence (FCS)

The payload, headers, and checks add up to 2180 bytes. Therefore, interfaces that carry FCoE traffic should have a configured maximum transmission unit (MTU) of 2180 or larger. An MTU size of 2180 bytes is the minimum size; some network administrators prefer an MTU of 2240 or 2500 bytes.

## Virtual Links

---

Native FC uses point-to-point physical links between FC devices. In FCoE, virtual links replace the physical links. A virtual link emulates a point-to-point link between two FCoE device endpoints, such as a server VN\_Port and an FC switch (or FCF) VF\_Port.

Each FCoE interface can support multiple virtual links. The MAC addresses of the FCoE endpoints (the VN\_Port and the VF\_Port) uniquely identify each virtual link and allow traffic for multiple virtual links to share the same physical link while maintaining data separation and security.

A virtual link exists in one FCoE VLAN and cannot belong to more than one VLAN. Although the FC switch and the FCoE device detect a virtual link as a point-to-point connection, virtual links do not need to be direct connections between a VF\_Port and a VN\_Port. A virtual link can traverse one or more transit switches, also known as passthrough switches.

A transit switch can transparently aggregate virtual links while still appearing and functioning as a point-to-point connection to the FCoE devices. However, a virtual link must remain within a single Layer 2 domain.

### FCoE VLANs

On the QFX Series, all FCoE traffic must travel in a VLAN dedicated to transporting only FCoE traffic. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE or FIP traffic must travel in a different VLAN.



**NOTE:** On a QFX3500 or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features. IGMP snooping is enabled by default on all VLANs; be sure to disable IGMP snooping on FCoE VLANs.

You can configure more than one FCoE VLAN, but any given virtual link must be in only one FCoE VLAN.



**NOTE:** All 10-Gigabit Ethernet interfaces that connect to FCoE devices must have a native VLAN configured in order to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols. Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

#### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding Fibre Channel on page 4961](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding FCoE Transit Switch Functionality on page 4972](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding Fibre Channel Terminology on page 5074](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)

## Understanding FCoE Transit Switch Functionality

You can use the QFX Series as a Fibre Channel over Ethernet (FCoE) transit switch. An FCoE transit switch is a Layer 2 data center bridging (DCB) switch that can transport FCoE frames and implements FCoE Initialization Protocol (FIP) snooping. A DCB switch transports both FCoE and Ethernet LAN traffic over the same network infrastructure while preserving the class of service (CoS) that Fibre Channel (FC) traffic requires.

An FCoE transit switch does not encapsulate or de-encapsulate FC frames in Ethernet. It is an access switch that transports FC frames that have already been encapsulated in Ethernet between FCoE initiators such as servers and a storage area network (SAN) FC switch that supports both Ethernet and native FC traffic on its interfaces. The transit switch acts as a passthrough switch and is transparent to the FC switch, which detects each connection to an FCoE device as a direct point-to-point link.

When a QFX Series acts as a transit switch, the VLANs you configure for FCoE traffic can use any of the switch ports on the device (QFX3500 switch or QFabric system Node device) because the traffic in both directions is standard Ethernet traffic, not native FC traffic.



**NOTE:** The Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets. It is a good practice to keep the native VLAN separate from the VLANs that carry FCoE traffic. FCoE VLANs should carry only FCoE traffic, but other types of untagged traffic might use the native VLAN.

FCoE traffic should use a VLAN dedicated only to FCoE traffic. Do not mix FCoE traffic with standard Ethernet traffic on a VLAN on the switch.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features. IGMP snooping is enabled by default on all VLANs; be sure to disable IGMP snooping on FCoE VLANs.



**NOTE:** On a QFX3500 switch or on a QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode. If you configure both a transit switch and an FCoE-FC gateway on the same QFX3500 switch or QFabric system Node device, configure different FCoE VLANs for the transit switch and the FCoE-FC gateway.

Transit switch architecture differs from FCoE-FC gateway architecture. As an FCoE-FC gateway, the system transports traffic to the FC SAN as native FC frames, and the VLAN must use an FCoE VLAN interface and native FC interfaces to transport that traffic. As a transit switch, the system forwards Ethernet traffic, and requires DCB configuration for



lossless transport of that traffic and FIP snooping at FCoE device access ports, but not the FCoE-FC gateway features necessary for transporting FC traffic.

The QFX Series complies with DCB standards for ensuring lossless transport and low latency, and provides 10-Gbps ports for FCoE traffic. For lossless transport to function correctly, you must use priority-based flow control (PFC, described in IEEE 802.1Qbb) to create bandwidth reservations and ensure proper CoS for FCoE traffic.

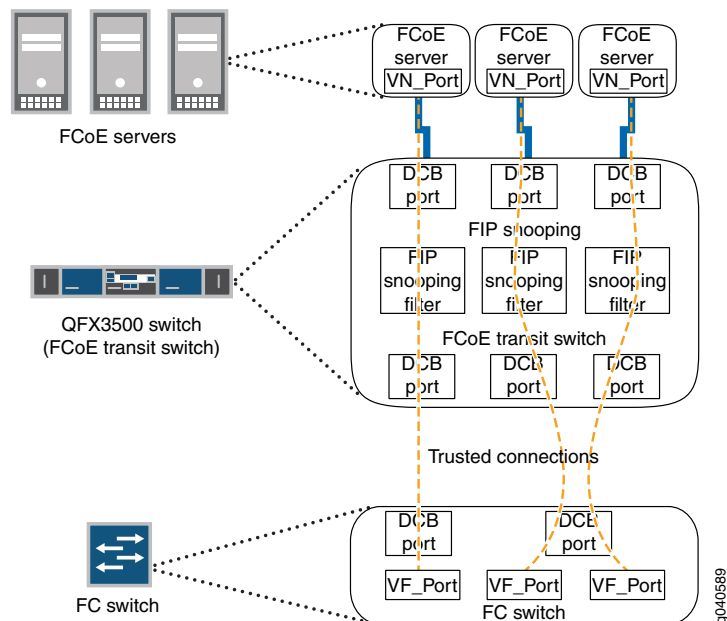
FIP snooping adds security by filtering access so that only traffic from servers that have successfully logged in to the FC network passes through the transit switch and reaches the FC network. The Technical Committee T11 organization specifications describe two types of FIP snooping:

- The FC-BB-5 specification describes VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network and FCF or FC switch VF\_Ports.
- The FC-BB-6 specification describes VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network.

To accommodate the larger size of Ethernet-encapsulated frames, FCoE interfaces should be configured with a maximum transmission unit (MTU) size of at least 2180 bytes.

The transit switch transparently connects FCoE-capable devices such as servers in an Ethernet LAN to an FC switch or to a gateway switch (hereafter referred to as the FC switch), as shown in [Figure 187 on page 4973](#). The transit switch acts as a transparent DCB access layer between FCoE servers and the FC switch.

**Figure 187: FCoE Transit Switch Connecting FCoE Devices to an FC Switch**



The transit switch performs FIP snooping at the ports connected to the FCoE devices. For VN2VF\_Port FIP snooping, at the SAN edge, the FC switch must be able to convert the FCoE traffic to native FC traffic. (VN2VN\_Port FIP snooping switches traffic between VN\_Ports directly through the transit switch, without going through the FC switch, so no conversion of FCoE traffic to native FC traffic is needed.)

Encapsulated FCoE traffic flows through the transit switch to the FCoE ports on the FC switch. The FC switch removes the Ethernet encapsulation from the FCoE frames to restore the native FC frames. Native FC traffic travels out native FC ports to storage devices in the FC SAN.

Native FC traffic from storage devices flows to the FC switch FC ports, and the FC switch encapsulates that traffic in Ethernet as FCoE traffic. The FCoE traffic flows through the transit switch to the appropriate FCoE device.



**NOTE:** The FC switch and FC fabric apply appropriate zoning checks on traffic to and from each ENode and provide FC services (for example, name server, fabric login server, or event server).

---



**NOTE:** The QFX3500 switch supports VN\_Port to VN\_Port FIP snooping to allow FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch. An FCoE VLAN can support either VN2VF\_Port FIP snooping (FC-BB-5) or VN2VN\_Port FIP snooping (FC-BB-6), but not both. The same QFX3500 switch can have multiple FCoE VLANs configured, some FCoE VLANs for VN2VF FIP snooping traffic and others for VN2VN FIP snooping traffic.

---

For load balancing, increasing available bandwidth, and port failover protection, you can configure the 10-Gigabit Ethernet interfaces that belong to an FCoE VLAN as a link aggregation group (LAG). In addition, creating a LAG prevents spanning tree algorithms from blocking physical links and wasting bandwidth.


#### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE on page 4968](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

## Understanding an FCoE-FC Gateway

A Fibre Channel over Ethernet (FCoE)-Fibre Channel (FC) gateway connects FCoE devices on an Ethernet network to an FC switch in an FC storage area network (SAN) as shown in [Figure 188 on page 4975](#). To FCoE devices such as servers, the FCoE-FC gateway presents virtual fabric ports (VF\_Ports) and appears to be an FCoE forwarder (FCF). To the FC switch, the FCoE-FC gateway presents a proxy node port (NP\_Port) and appears to be an FC device.

**Figure 188: FCoE-FC Gateway Topology**



The FCoE-FC gateway handles FCoE Initialization Protocol (FIP) and FCoE traffic on the interfaces connected to FCoE devices. The gateway forwards native FC traffic on the interfaces to the FC switch. The gateway does not provide FC services (such as fabric login server or name server). It is a proxy for an FCF, not an FCF or an FC switch. The gateway transparently substitutes for the FC switch when communicating with FCoE devices and transparently substitutes for FCoE devices when communicating with the FC switch.

The gateway does not use an FC domain ID, so it extends the SAN fabric while saving domain resources. Using the gateway also means that the FC switch does not have to handle FCoE traffic (and therefore requires no FCoE blades or ports). The gateway converges Ethernet and FC backbones to leverage existing resources.

- [Gateway FC Fabric on page 4976](#)
- [Fabric Services on page 4977](#)
- [FCoE-FC Gateway Traffic Switching on page 4977](#)

## Gateway FC Fabric

---

A gateway FC fabric is a QFX Series configuration construct. It is not the same thing as an FC fabric in the SAN; the gateway FC fabric is local to the switch. It creates associations that connect FCoE devices with converged network adapters (CNAs) on the Ethernet network to an FC switch on the Fibre Channel network. A gateway FC fabric consists of:

- A unique fabric name.
- A unique fabric ID.
- At least one dedicated VLAN for FCoE traffic. VLANs that carry FCoE traffic should not carry any other type of traffic.



**NOTE:** On a QFX3500 or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.

- At least one FCoE VLAN interface (Layer 3 VLAN interface) that includes one or more 10-Gigabit Ethernet interfaces connected to FCoE devices. The FCoE VLANs transport traffic between the FCoE servers and the FCoE-FC gateway. Each FCoE VLAN must carry only FCoE traffic. You cannot mix FCoE traffic and standard Ethernet traffic on the same VLAN.

The 10-Gigabit Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic because FIP VLAN discovery and notification frames are exchanged as untagged packets.

Each FCoE VLAN interface can present multiple VF\_Port interfaces to the FCoE network.



**NOTE:** Storm control must be disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped.

- One or more native FC interfaces. The native FC interfaces transport traffic between the gateway and the FC switch.



**TIP:** If the network does not use a dual-rail architecture for redundancy, configure more than one native FC interface for each FC fabric to create redundant connections between the FCoE devices and the FC switch. If one physical link goes down, any sessions it carried can log in again and connect to the FC switch on a different interface. Even in dual-rail architecture networks, creating redundant connections between the QFabric system and the FC switch is the best practice.

You can also configure FIP parameters for the fabric or accept the default FIP parameters. VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping is automatically enabled on all server-facing ports because all ports are untrusted by default. You can disable VN2VF\_Port FIP snooping on a port-by-port basis by marking a port as an FCoE trusted interface. You

can disable VN2VF\_Port FIP snooping on all Ethernet ports in an FC fabric by configuring the fabric as FCoE trusted.

Because the switch has 12 native FC ports and each FC fabric requires a minimum of one native FC port, the switch supports a maximum of 12 FC fabrics. However, as a best practice for redundancy, we recommend that you assign at least two native FC interfaces to each FC fabric.

On a QFabric system, all of the FC and FCoE traffic that belongs to a particular gateway FC fabric must ingress and egress the same gateway Node device. Gateway FC fabrics do not span across Node devices. All of the native FC interfaces and the Ethernet interfaces that belong to the FCoE VLAN must reside on the same gateway Node device to be included in an FC fabric on that Node device.

Traffic from FC and FCoE devices that are not in the same FC fabric remain separate and cannot communicate with each other through the gateway.

### **Fabric Services**

---


The FC switch provides all FC services (domain manager, name server, fabric login server, and so on) except FIP to the FCoE devices. The FC switch assigns all FCIDs (through N\_Port ID virtualization) and fabric attributes to FCoE device VN\_Ports.

The FCoE-FC gateway does not provide FC services (except FIP). The gateway relays communication between the FC switch and the FCoE devices, encapsulates and de-encapsulates native FC frames, converges Ethernet and FC backbones, and aggregates FCoE device VN\_Port sessions.

### **FCoE-FC Gateway Traffic Switching**

---

All traffic that flows through the gateway FC fabric is switched through the FC switch. Even if two hosts on the Ethernet FCoE network connect directly to the gateway, FCoE communication between them goes through the FC switch, as shown in [Figure 189 on page 4978](#).

**Figure 189: Traffic Switching Between FCoE Hosts Connected to the FC Network by an FCoE-FC Gateway**

For example, FCoE host server *Host1* sends frames destined for FCoE host server *Host2*. Both *Host1* and *Host2* are directly connected to the gateway. The communication path looks like this:

1. *Host1* sends FCoE frames destined for *Host2* to the gateway.
2. The gateway de-encapsulates the FCoE frames from *Host1* into native FC frames and switches them to the FC switch.
3. The FC switch processes the native FC frames and sends them back to the gateway destined for *Host2*.
4. The gateway encapsulates the FC frames in Ethernet and sends the resulting FCoE frames to *Host2*.
5. When *Host2* replies, the FCoE reply goes to the gateway. The gateway de-encapsulates the reply and switches it to the FC switch for processing. The FC switch then sends it back to the gateway, which encapsulates the FC frames and sends them to *Host1*.

#### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding Fibre Channel on page 4961](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
- [Overview of FIP on page 4960](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway on page 5189](#)

- [Understanding Fibre Channel Terminology on page 5074](#)

## Understanding FCoE-FC Gateway Functions

When it functions as a Fibre Channel over Ethernet (FCoE)-Fibre Channel (FC) gateway, the QFX Series provides the following functions:

- [Login and Logout on page 4979](#)
- [FCoE and FC Frame Handling on page 4979](#)
- [Data Center Bridging on page 4979](#)
- [Disabling the Fabric WWN Verification Check on page 4980](#)
- [Load Balancing on page 4981](#)

### Login and Logout

Each of the native FC interfaces on the gateway performs a fabric login (FLOGI) to the FC switch when each interface initializes. This establishes the link between each gateway FC interface and the FC switch.

When FCoE devices on the Ethernet network send an FCoE Initialization Protocol (FIP) login (FIP FLOGI) or FIP discovery (FIP FDISC) request to the gateway, the gateway acts on behalf of those devices and converts their FIP FLOGI and FIP FDISC requests to FC FDISC requests. The gateway then sends the FC FDISC requests to the FC switch. When the FC switch responds to an FDISC request, the gateway converts the FC response into a FIP response and sends it to the appropriate FCoE device.

The gateway also converts FIP logout (LOGO) requests from FCoE devices into FC LOGO requests to the FC switch, and converts the FC switch response into a FIP response for the FCoE device.

### FCoE and FC Frame Handling

When it receives FCoE frames from FCoE devices, the gateway strips away the Ethernet encapsulation from the FC frame before sending the native FC frame to the FC switch.

When it receives native FC frames from the FC switch, the gateway encapsulates the native FC frames in Ethernet before sending the resulting FCoE frames to the appropriate VN\_Port.

### Data Center Bridging

The Ethernet ports connected to the FCoE devices are 10-Gbps Ethernet ports and support data center bridging (DCB) specifications:

- Priority-based flow control (PFC, described in IEEE 802.1Qbb)
- Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB)
- Enhanced transmission selection (ETS, described in IEEE 802.1Qaz)
- 10-Gigabit Ethernet ports

### Disabling the Fabric WWN Verification Check

---

The gateway connects to a SAN fabric using the gateway NP\_Ports (native FC ports). When the NP\_Ports initialize, each port sends a FLOGI to the FC switch to which it is connected in the SAN fabric. The FC switch sends a FLOGI accept (FLOGI-ACC) message back to each NP\_Port. The FLOGI-ACC message includes the SAN fabric worldwide name (WWN). The gateway uses the SAN fabric WWN in the multicast discovery advertisement (MDA) that the gateway sends to the ENodes in the FCoE network.

Some FC switches substitute their own WWN (often the FC switch's virtual WWN) for the SAN fabric WWN in the FLOGI-ACC message. When the FC switch substitutes its own WWN for the fabric WWN, gateway NP\_Ports that log in to the same SAN fabric might receive different fabric WWNs in the FLOGI-ACC messages if the NP\_Ports are connected to different FC switches in that SAN fabric. This creates a problem, because different fabric WWNs indicate different SAN fabrics. But in this scenario, the different fabric WWNs come from different FC switches in the same SAN fabric.

If the gateway receives different fabric WWNs on NP\_Ports that are connected to the same SAN fabric, the gateway uses the first fabric WWN it receives in the MDA it sends to the ENodes. The gateway isolates the NP\_Ports connected to that fabric that receive a different fabric WWN in the FLOGI-ACC message. No ENode sessions are assigned to the isolated NP\_Ports. FC traffic is assigned only to NP\_Ports that receive a fabric WWN that matches the fabric WWN received by the first NP\_Port to log in to the FC fabric. (If an NP\_Port receives a fabric WWN that does not match the fabric WWN received by the first NP\_Port to log in to the FC fabric, it does not carry traffic to the SAN fabric.)

In summary, the scenario is:

1. The gateway has multiple NP\_Ports connected to more than one FC switch in a SAN fabric.
2. When the NP\_Ports initialize, each NP\_Port sends a FLOGI to the FC switch to which it is connected.
3. The FC switches substitute their own WWNs for the fabric WWN in the FLOGI-ACC message, so different NP\_Ports receive different fabric WWNs.
4. In the MDA the gateway sends to FCoE devices, the gateway uses the fabric WWN that the first NP\_Port to log in to the fabric receives in the FLOGI-ACC message. If other NP\_Ports receive a different fabric WWN from other FC switches in the SAN fabric, that fabric WWN is not advertised.
5. NP\_Ports that receive a fabric WWN that does not match the first received fabric WWN are isolated, and the ENode sessions cannot use those ports.

To prevent this from happening, you can disable the gateway fabric WWN verification check so that all NP\_Ports connected to a SAN fabric are used to carry traffic between the gateway and the FC switch, regardless of the fabric WWN the NP\_Port receives in the FLOGI-ACC message.





**NOTE:** Disabling or enabling the fabric WWN verification check logs out all FCoE sessions.

### Load Balancing

The switch performs automatic link load balancing for the connections between the gateway and the FC SAN and can also perform load balancing for the connections between the gateway and the FCoE devices in the Ethernet network. On the native FC links (NP\_Ports) between the gateway and the FC SAN, the gateway can use one of the following two load-balancing algorithms:

- Simple load balancing—Each ENode FLOGI session and VN\_Port FDISC session is assigned to the least-loaded link. This is the default load-balancing algorithm.
- ENode-based load balancing—When an ENode logs in to the fabric, all subsequent DISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link.



**NOTE:** Changing the load-balancing algorithm when FCoE sessions are running forces the FCoE sessions to log out, then log in again.

You can balance the load on the Ethernet ports facing the FCoE devices by configuring those ports as a link aggregation group (LAG).

#### Related Documentation

- [Understanding Fibre Channel on page 4961](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Disabling the Fabric WWN Verification Check on page 5180](#)
- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

### Understanding FCoE and FIP Session High Availability

The QFX Series provides high availability features to maintain storage network sessions when a system process is terminated and during certain types of upgrades:

- [High Availability for Fibre Channel Process Termination \(FCoE-FC Gateway Mode\) on page 4982](#)
- [High Availability for FIP Snooping on page 4982](#)
- [Nonstop Software Upgrade \(QFabric Systems\) on page 4983](#)

### High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode)

In FCoE-FC gateway mode, the QFX3500 switch provides high availability to restore the FCoE sessions running on the switch in case the Fibre Channel (FC) process is terminated. A session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric, not an end-to-end server-to-storage session.

The switch stores FCoE session data in a persistent storage module. If the FC process terminates, the switch restores the existing FCoE sessions on the same interfaces that they were on before the FC process terminated. Data traffic for existing sessions is not affected during session restoration.

For a brief time, the system does not process control traffic because of the FC process restart and session restoration. During this brief time, no new FCoE sessions can be established, and no existing sessions can log out.



.....

**NOTE:** During the restoration process, if the FC process does not receive an *interface up* notification from a particular interface within a certain time, the switch times out the restore operation and discards the data on that interface. The previously existing FCoE sessions on that interface are not restored, and the ENodes must log in again.

.....



.....

**NOTE:** An FC process restart and session restoration resets the Fibre Channel statistics.

.....

If the FC process terminates repeatedly, the operating system disables the process until you manually restart it. To restart the FC process manually, issue the **restart fibre-channel** command.

### High Availability for FIP Snooping

You can configure the system to perform FIP snooping on Ethernet interfaces that are connected to FCoE devices that have ENodes. The QFX Series provides high availability to restore running FIP snooping sessions in case the Ethernet switching process is terminated.

The Ethernet switching process stores the FIP snooping state in a persistent storage module. If the Ethernet switching process terminates, the QFX Series restores the existing FIP snooping sessions on the same interfaces that they were on before the Ethernet switching process terminated. The high availability features preserve:

- Logged in ENodes
- Discovered FCFs
- Existing sessions
- Existing FIP snooping filters

The complete restoration process, including reconciling all valid states, takes a maximum of 8 seconds. During the restoration process, the switch can learn a new FCF or a new FC switch, and new ENodes can log in to the FC network. However, FDISC messages from an ENode that is already logged in to the network might be dropped if the ENode has not yet been restored.

When the Ethernet switching process terminates ungracefully, the FIP keepalive timer is reset to the normal initial value, not the value at the time of the Ethernet switching process termination.

In the event of an Ethernet switching process termination, ENodes remain logged in, and existing sessions are not interrupted.



**NOTE:** An Ethernet switching process restart and session restoration resets the FIP snooping statistics.

### Nonstop Software Upgrade (QFabric Systems)

On QFabric system Node groups that have more than one Node device, nonstop software upgrade (NSSU) enables you to upgrade the Node devices with minimal packet loss and maximum uptime. NSSU automates software upgrades on the QFabric system components in an orderly and consistent manner to maximize system uptime.

The system upgrades components with redundant architectures, such as redundant server Node groups and network Node groups that have two or more members, in stages. While the system upgrades one component, the redundant component continues to function.

For example, while one member of a redundant server Node group is upgraded, the other member continues to forward traffic. When the first Node group member completes the upgrade, it comes online while the system upgrades the second member.

NSSU provides high availability for the lossless traffic forwarding required to support storage networks. If your system design includes redundancy (redundant Node devices in Node groups, LAGs, and so on) so that an alternate traffic path is available, when you upgrade a Node device, traffic is not impacted.

In fully redundant topologies, NSSU preserves FIP session, FIP snooping filter, VN2VF\_Port session, and VN2VN\_Port session information and prevents traffic loss in most cases. An exception is that Node devices that are directly connected to ENodes experience momentary traffic loss when the Node device reboots.

#### Related Documentation

- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE on page 4968](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 57](#)
- [Performing a Nonstop Software Upgrade on the QFabric System on page 105](#)

## Understanding FIP Functions

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) performs four major functions:

- FIP VLAN discovery: FCoE device FCoE nodes (ENodes) discover the FCoE VLANs on which to transmit and receive FIP and FCoE traffic.
- FIP discovery: FCoE devices discover Fibre Channel (FC) switches to which they can connect.
- Initialization: FCoE devices perform fabric login (FLOGI) and fabric discovery (FDISC) to create a virtual link with an FC switch.
- Maintenance: The QFX Series ensures that the virtual link between the FCoE device and the FC switch remains valid, and also that the link termination logout (LOGO) functions properly.

When you configure the switch as an FCoE-FC gateway, it converts FIP requests and information from FCoE devices into FC requests and information and relays them to the FC switch. To FCoE devices, the gateway appears to be an FCoE forwarder (FCF) and presents virtual fabric port (VF\_Port) interfaces to the server ENode. To FC switches, the gateway appears to be an FC device that supports N\_Port ID virtualization (NPIV) and presents an N\_Port interface to the FC switch F\_Port interface. When you configure the QFX Series as an FCoE transit switch, you do not configure FIP parameters on the switch.

FIP FLOGI, FDISC, and LOGO are similar to the same processes in the native FC protocol.

This topic describes:

- [FIP VLAN Discovery on page 4984](#)
- [FIP Discovery on page 4985](#)
- [FIP FLOGI on page 4986](#)
- [FIP FDISC on page 4986](#)
- [FIP Maintenance \(Keepalive Messages\) on page 4987](#)
- [FIP LOGO on page 4987](#)

---

### FIP VLAN Discovery

The gateway supports FIP VLAN discovery. Host ENodes use FIP VLAN discovery to discover the FCoE VLANs on which they will send and receive FIP and FCoE traffic and on which they will establish a virtual link with the FC switch. This means FCoE devices do not need manually configured FCoE VLANs.

FIP VLAN discovery and notification takes place on the native VLAN that the FCoE device uses for Ethernet traffic:

1. The ENode sends a FIP VLAN discovery request to a multicast address called *ALL-FCF-MACs* to which all FC switches and FCFs on the VLAN listen.
2. The FC switches and FCFs respond on the native VLAN with a list of the FCoE VLANs that are available for login.

3. The ENode selects an FCoE VLAN and continues the FIP process on that VLAN.

Except for FIP VLAN discovery, all other FIP and FCoE traffic runs on an FCoE VLAN.



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols. Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

### FIP Discovery

The FIP discovery process allows an FCoE device ENode MAC to locate (discover) the FC switches in the FCoE VLAN to which it belongs. The ENode selects an FC switch to log in to from the available FC switches. Either the ENode MAC or the FC switch can initiate the FIP discovery process.

Server ENode MACs initiate FIP discovery:

1. When an ENode MAC comes online, it sends a multicast discovery solicitation message on its FCoE VLAN to a multicast address called *ALL-FCF-MACs* to which all FCFs (including the FCF functionality of FC switches) on the VLAN listen. The discovery solicitation message includes the ENode's addressing mode and the maximum protocol data unit (PDU) size the ENode MAC uses for FCoE traffic.

The ENode uses the globally unique ENode MAC address assigned to it by the converged network adapter (CNA) manufacturer as an identifier in the FIP frame header.

2. The FCFs on the VLAN that have a similar supported addressing mode, match the maximum FCoE size, and can accept a login from the ENode reply to the discovery solicitation message by sending a solicited unicast discovery advertisement message to the soliciting ENode MAC.
3. The ENode MAC compiles a list of FCFs that are available for login, selects an FCF (the FCF with the highest priority setting), and is then ready to log in to the FCF.

The FIP discovery process is similar when the FC switch or FCF initiates discovery:

1. FCF MACs periodically send unsolicited multicast discovery advertisements on the FCoE VLAN to the *ALL-ENode-MACs* multicast address, to which all ENode MACs on the VLAN listen. The FIP keepalive advertisement period timer (FKA\_ADV\_PERIOD) controls the interval between multicast discovery advertisements. The multicast discovery advertisements inform ENodes on the VLAN that FCF VF\_Ports are available for establishing virtual links with ENode VN\_Ports.
2. ENodes on the FCoE VLAN create an entry for the FCF-MAC in their FCF-MAC lists.
3. An ENode can respond to the unsolicited multicast discovery advertisement with a unicast discovery solicitation message to the FCF.
4. Upon receiving the ENode's unicast discovery solicitation, the FCF replies with a unicast discovery advertisement sent to the ENode MAC.

After the ENode MAC selects an FCF to log in to, FIP initialization begins. To proceed from discovery to initialization, the server ENode addressing mode must match the FCF addressing mode and maximum FCoE size. In addition, the FCF must be configured to allow FIP FLOGI from that ENode.

### FIP FLOGI

---

FIP initialization is the server ENode login process to the FCF after the ENode discovers the FCFs (including FC switches) on the FCoE VLAN:

1. The ENode sends a fabric login (FLOGI) request message to the FCF.
2. The FCF replies to confirm the ENode login and provides the ENode a locally unique MAC address to use for FCoE frame transactions. The locally unique MAC address identifies the VN\_Port interface of the ENode for the session the login establishes. (The ENode continues to use the globally unique ENode MAC address for FIP frame transactions.)

The locally unique ENode MAC address for FCoE operations depends on whether the ENode address mode is configured as a fabric-provided MAC address (FPMA) or as a server-provided MAC address (SPMA; the gateway does not support ENodes in SPMA mode and rejects login attempts from ENodes in SPMA mode):

- For FPMA mode, the FCF provides a MAC address to the ENode during the FIP FLOGI exchange. The FPMA MAC address is a 48-bit value that is unique to the local fabric and consists of a 24-bit FCoE mapped address prefix (FC-MAP) and a 24-bit FC identifier (FCID). You can configure the FC-MAP value on the FCF or use the default value of 0EFC00h. The FCoE device must use the same FC-MAP value as the FCF, or else discovery and login fail.
- For SPMA mode, the server provides its MAC address to the FCF. The FCF compares the server MAC address to a list of addresses approved for FCoE access. The gateway does not support ENodes in SPMA mode.

Successful login instantiates a secure virtual link between the ENode and the FCF and terminates the FIP virtual link instantiation phase. The initiating server behind the ENode can exchange FC payloads with storage devices in the FC SAN by sending FCoE frames over the virtual link.

### FIP FDISC

---

After an ENode successfully logs in to an FCF and establishes a virtual link, the ENode can request more virtual links (sessions) over the same physical link by sending a FIP fabric discovery (FDISC) request. FDISC allows the creation of multiple separate secure VN\_Port virtual links on one physical link. Each virtual link receives a locally unique identifier from the FCF to enable security and separation between the VN\_Port virtual links sharing a physical ENode port. This is called N\_Port ID virtualization (NPIV).

FDISC is similar to FLOGI in that it requests a login and a unique ID from the FCF. The difference is that FLOGI obtains the initial login and ID for the physical link, whereas FDISC obtains additional logins and IDs so that multiple virtual links can share one physical link securely.

After a VN\_Port FDISC is complete, the application using that VN\_Port can send FCoE frames over the virtual link.

### FIP Maintenance (Keepalive Messages)

Although FCoE protocol handles the payload communication between the initiating ENode and the target FC device, FIP continues to run in the background. FIP constantly updates ENode FCF lists by listening to the periodic FCF multicast discovery advertisements, and it verifies the ability to reach the FCF by transmitting periodic FIP keepalive advertisements.

The ENode sends periodic ENode FIP keepalive advertisements to the FCF with the ENode MAC address as the identifier. The ENode also sends periodic VN\_Port FIP keepalive advertisements on behalf of each VN\_Port on the ENode, using the VN\_Port MAC address as the source MAC. The VN\_Port FIP keepalive advertisements occur every 90 seconds. The keepalive advertisements reset the session timer for the virtual link connection to the FCF. If the FCF does not receive a keepalive advertisement for a logged-in ENode or VN\_Port before the session timer expires, the virtual link is terminated.

The periodic unsolicited multicast discovery advertisements the FCF sends to the *ALL-ENode-MACs* address continuously verify that the FCF is still reachable. The ENode and the FCF periodic unsolicited multicast discovery advertisements occur at the configured FIP keepalive advertisement period interval (FKA\_ADV\_PERIOD) plus or minus a random offset to prevent a flood of simultaneous keepalive advertisements.

If the FCF does not receive the ENode keepalive advertisements before the FCF's FIP keepalive timer expires, the FCF considers the virtual link to the ENode as "down" and terminates the virtual link to the ENode. The keepalive timer expires in 2.5 times the configured timer value. This also terminates any VN\_Port virtual links instantiated by that ENode.

If the FCF does not receive a VN\_Port keepalive advertisement before the FCF's FIP keepalive timer expires, the FCF considers the virtual link to the VN\_Port as "down" and terminates the virtual link to that VN\_Port. The VN\_Port keepalive timer expires in 2.5 times the configured timer value.

If the ENode does not receive the FCF unsolicited multicast discovery advertisement before the ENode's FIP keepalive timer expires, the ENode considers the virtual link to the FCF as "down" and all of the VN\_Port virtual links to that FCF on the ENode are terminated.

### FIP LOGO

FIP handles ENode and VN\_Port logout when a session is finished.

#### Related Documentation

- [Overview of FIP on page 4960](#)
- [Understanding FIP Implementation on page 4988](#)
- [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)
- [Understanding Fibre Channel Virtual Links on page 4995](#)
- [Understanding FCoE on page 4968](#)

## Understanding FIP Implementation

In a network that converges Fibre Channel (FC) and Ethernet traffic, when you configure the QFX Series as a Fibre Channel over Ethernet (FCoE)-FC gateway, it translates FCoE Initialization Protocol (FIP) frames from FCoE nodes (ENodes) into native FC frames for FC switches and translates native FC frames from FC switches into FIP frames for ENodes. To an FCoE device, the gateway appears to be an FCoE forwarder (FCF) and presents a fabric port (F\_Port) interface to the FCoE device ENode. To an FC switch, the gateway appears to be an FC host capable of N\_Port ID virtualization (NPIV) and presents a node port (N\_Port) interface to the FC switch F\_Port interface.



**NOTE:** The N\_Ports that the gateway presents to the FC switch are called proxy N\_Ports (NP\_Ports). To the FC switch, the gateway NP\_Ports appear to be native FC N\_Ports that are capable of performing NPIV. The NP\_Ports are proxies for the FCoE devices in the Ethernet network. The NP\_Ports convert FCoE traffic from the FCoE devices into native FC traffic for the FC switch. The NP\_Ports also convert native FC traffic from the FC switch into FCoE traffic for the FCoE devices on the Ethernet network.

- [FIP Basics on page 4988](#)
- [Fabric Login and FIP Login Overview on page 4989](#)
- [Proxy FIP Discovery on page 4990](#)
- [Proxy FIP Initialization on page 4990](#)
- [Proxy FIP Maintenance on page 4991](#)
- [Proxy FIP Logout on page 4991](#)

---

### FIP Basics

FIP is enabled by default on all VLAN interfaces that belong to each FC fabric configured on the gateway. You can configure FIP parameters at a global level or on an individual interface. When you configure a parameter on an interface, it overrides the global configuration only for that interface. If you do not explicitly configure a FIP parameter, the gateway uses the default value.

In order for the gateway to connect FCoE devices with FCFs, the FIP parameters you configure on the gateway must be compatible with the parameters configured on the FC switch (for example, the FC-MAP values of the FC switch and of the FC fabric FIP configuration on the gateway must match, or the FC switch drops the frames).

When the NP\_Ports on the gateway come up, they perform an FC FLOGI to the connected FC switch. Successful login establishes communication between the gateway and the FC switch, and gateway NP\_Ports are marked for sending FDISC messages. Successful login also creates a next-hop entry in the gateway for the FC switch. If the FC switch rejects the FLOGI request, no link is established. The gateway maintains a list of valid FCF-MACs with which ENodes can connect.




After establishing communication with an FC switch, the gateway can connect FCoE devices in the Ethernet network to the FC switch. All of the subsequent connections the gateway makes with FC switches as a proxy for ENodes (on behalf of ENodes) are virtualized (NPIV) connections.

### Fabric Login and FIP Login Overview

Figure 190 on page 4989 provides a brief overview of the FCoE-FC gateway fabric login to the FC switch and the FCoE device FIP login to the gateway.

Figure 190: FCoE-FC Gateway Fabric Login and FIP Login



The numbers in the following list correspond to the numbers in Figure 190 on page 4989 and briefly describe each step of the login process:

1. The FCoE-FC gateway NP\_Port sends an FC fabric login (FLOGI) request to the FC switch F\_Port.
2. The FC switch accepts the gateway FLOGI.
3. The gateway sends FIP multicast discovery advertisements on the FCoE VLAN (the gateway F\_Port interface) to all connected FCoE device ENodes.
4. The FCoE device ENode sends a discovery solicitation message to the gateway.
5. The gateway responds with a unicast discovery advertisement to the ENode.



**NOTE:** The gateway limits the number of discovery solicitations it accepts from FCoE devices to a maximum of 100 outstanding requests at any given time. If the gateway has 100 discovery solicitations outstanding, the gateway does not respond to new discovery solicitations. Instead, the gateway drops new discovery solicitations and reports the number of dropped discovery solicitations in the **Dropped** field of the `show fibre-channel fip statistics` command output. When there are fewer than 100 outstanding discovery solicitations, the system responds to new requests as usual with a discovery advertisement.

6. The FCoE device sends a FIP FLOGI or FIP FDISC message to the gateway.
7. The gateway converts the FIP FLOGI or FIP FDISC to an FC FDISC and forwards it to the FC switch to obtain a login for the FCoE device.
8. The FC switch responds to the FC FDISC by sending a new ID for the NPIV session to the gateway.
9. The gateway converts the FC FDISC response from the FC switch to a FIP FDISC response and forwards it to the FCoE device.

The following sections describe some of these steps in greater detail.

### Proxy FIP Discovery

---

After the gateway establishes a connection with an FC switch:

1. The gateway sends periodic FIP multicast discovery advertisements on the FCoE VLAN so that ENodes can add the gateway to their FCF lists.
2. The ENode initializes and sends a multicast discovery solicitation message on the FCoE VLAN. If the ENode has already initialized and has a list of FCFs, it can send a unicast discovery solicitation message to a particular FCF such as the gateway.



.....

**NOTE:** The gateway limits the number of discovery solicitations it accepts from FCoE devices to a maximum of 100 outstanding requests at any given time. If the gateway has 100 discovery solicitations outstanding, the gateway does not accept new discovery solicitations until there are fewer than 100 discovery solicitations outstanding.

.....

3. When the gateway receives a multicast discovery solicitation from an ENode, it responds by sending a unicast discovery advertisement to that ENode.

When the gateway receives a unicast discovery solicitation from an ENode, it also responds with a unicast discovery advertisement to the ENode.

To the ENode, the gateway appears to be an FCF.

The FIP discovery process adds the ENode to the gateway ENode database.

### Proxy FIP Initialization

---

1. If the ENode chooses to log in to the gateway, it responds to the gateway's unicast discovery advertisement by sending a login request in the form of a FIP FLOGI if it is the initial connection to the gateway. If the ENode already has an established session with the gateway and another application or virtual machine wants to connect to the gateway, the ENode sends a FIP FDISC to the gateway.
2. The gateway receives the FIP FLOGI or FIP FDISC from the ENode, converts it into an FC FDISC, and sends it through the least-loaded NP\_Port to the FC switch on behalf of the ENode. The FC FDISC message requests an FCID for the new virtual link.



**NOTE:** The gateway converts both ENode FIP FLOGI and FIP FDISC messages into FC FDISC messages, because the gateway has already performed FC FLOGI with the FC switch, so all subsequent connection requests on the gateway NP\_Port are FDISC requests for virtual (NPIV) connections. FDISC messages request a virtual N\_Port connection over an existing physical N\_Port connection.

3. The FC switch processes the request, accepts it, assigns a unique FCID for the connection, and then sends the response to the gateway. If the FC switch rejects the FDISC request, no virtual link is established.
4. The gateway maps the FC switch response to the ENode VN\_Port, converts the FC acceptance message to a FIP FLOGI or FIP FDISC response, and sends it to the ENode VN\_Port.
5. The ENode VN\_Port accepts the FCID, and the virtual link is established.

If an ENode sends an FDISC, the proxy gateway switch checks whether the ENode has already performed a FLOGI to create the initial connection. If the ENode has not performed a FLOGI, the FDISC request is dropped.

The FC protocol does not recognize multipoint-to-point connections. Although the gateway can aggregate traffic from multiple FCoE servers on one NP\_Port, each virtual link appears to be an individual point-to-point link between an FCoE ENode VN\_Port and the FC switch, not as an aggregated multipoint-to-point link. The gateway is essentially invisible to the FC protocol, so the virtual link looks and acts like a point-to-point link from the FCoE device to the FC switch.

### Proxy FIP Maintenance

The gateway sends and receives periodic FIP keepalive messages to and from ENode VN\_Ports to maintain the connection between the gateway and the ENodes.

### Proxy FIP Logout

As with FIP discovery and FIP FLOGI, the gateway represents the FCoE device in transactions with the FC switch and represents the FC switch in transactions with the FCoE device:

1. An ENode VN\_Port sends a FIP LOGO message to log off and terminate the virtual link connection.
2. The gateway converts the FIP LOGO to an FC LOGO and relays it to the FC switch.
3. The FC switch responds to the LOGO request.
4. The gateway converts the FC LOGO response to a FIP LOGO response and relays it to the VN\_Port, completing the logout and terminating the virtual link.

#### Related Documentation

- [Overview of FIP on page 4960](#)
- [Understanding FIP Functions on page 4984](#)

- [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)
- [Understanding Fibre Channel Virtual Links on page 4995](#)
- [Understanding FCoE on page 4968](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)

## Understanding FIP Parameters on an FCoE-FC Gateway

By default, FIP is enabled, and the default FIP settings are valid on all FCoE interfaces that are part of the gateway FC fabric. You can configure some FIP parameters at a global level or on a specific interface. Some FIP parameters can be configured only at the global level or only at the individual interface level. When you configure a parameter at the interface level, the configuration overrides the global setting for that interface only.

- [FIP Keepalive Advertisement Period on page 4992](#)
- [Addressing Mode on page 4992](#)
- [FC-MAP on page 4993](#)
- [FCoE Trusted Fabric on page 4994](#)
- [Maximum Number of FCoE Sessions Per ENode on page 4994](#)
- [Priority on page 4994](#)

### FIP Keepalive Advertisement Period

---

The FIP keepalive advertisement period (fka-adv-period) is the time interval between messages that verify the connection is still valid and the device at the other end of the virtual link is still reachable. The ENode sends an ENode FIP keepalive advertisement to the gateway with the ENode MAC address as the source address to verify its reachability. The ENode also sends VN\_Port FIP keepalive messages for every VN\_Port on the ENode that is logged in to the gateway, with the VN\_Port MAC address as the source address.

The FIP keepalive advertisement period also determines the time interval between unsolicited multicast discovery advertisements from the gateway to the *ALL-ENode-MACs* multicast address. Unsolicited multicast discovery advertisements serve as keepalive messages from the gateway to the ENodes and also advertise the gateway's presence on the network.

The gateway sends the periodic unsolicited multicast discovery advertisements to the ENodes. On the gateway, you can configure a global FIP keepalive advertisement period for an FC fabric and you can configure a FIP keepalive advertisement period for individual interfaces to override the global setting.

### Addressing Mode

---

For FIP transactions, the ENode identifies itself using the globally unique MAC address assigned to the CNA by the manufacturer. After FIP has established a virtual link between an ENode VN\_Port and the gateway, for FCoE transactions, the VN\_Port identifies itself using a locally unique MAC address. The format of the locally unique MAC address depends on the addressing mode the fabric supports and the addressing mode the ENode is programmed to use.

The addressing mode is not a configurable parameter on the gateway. FC fabrics on the gateway support only the fabric provided MAC address (FPMA) addressing mode for FCoE transactions. The gateway does not support the server provided MAC address (SPMA) addressing mode. ENodes that use SPMA cannot log in to the gateway.

The FC switch assigns a locally unique FPMA to an ENode MAC through the FLOGI or FDISC process:

1. During the FIP discovery process, the ENode compiles a list of compatible FCFs (including the gateway) in the fabric. A compatible addressing mode is one of the criteria an FCF must meet to be added to an ENode's compatible FCFs list.
2. The ENode MAC transmits a FLOGI or FDISC to the FCF that includes the addressing modes the ENode supports.
3. If the FCF supports an addressing mode the ENode uses, the FCF accepts the FLOGI or FDISC and assigns the FPMA in the accept message (FIP FLOGI LS\_ACC or FIP NPIV FDISC LS\_ACC). If the ENode uses an addressing mode that is incompatible with the FCF, the FLOGI or FDISC is rejected.

The FPMA uniquely identifies a single VN\_Port at that ENode MAC in FCoE transactions with the FCF. Each VN\_Port connection receives its own unique FPMA to identify its virtual link connection. When an ENode uses NPIV to create multiple VN\_Ports, each VN\_Port virtual link receives its own unique FPMA to identify its traffic.

An FPMA consists of two concatenated 24-bit values:

1. The upper 24 bits are the FCF's FC-MAP value, which is a MAC address prefix that is unique to the fabric.
2. The lower 24 bits are the locally unique FCID that the FCF (FC switch) assigns to the VN\_Port.

The combination of these values guarantees that each FPMA is unique within a fabric.

### FC-MAP

The FCoE mapped address prefix (FC-MAP) value is a MAC address prefix used by the FCF that is unique within a given fabric. The FCF uses the FC-MAP for FCoE traffic within that fabric. The FCF rejects FCoE traffic that uses an FC-MAP value that does not match the FCF's FC-MAP value. In most cases, the FCF uses the default FC-MAP value (0EFC00), but a pool of 256 values is available (0EFC00 through 0EFCFF).

The gateway learns FC switches in the fabric that match the gateway fabric's FC-MAP value. To learn and communicate with an FC switch, the FC-MAP value for a fabric (or for the fabric's FCoE VLAN) on the gateway must match the FC switch's FC-MAP value. If the FC-MAP values do not match, no connection is established.



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

## FCoE Trusted Fabric

---

By default, all interfaces are untrusted interfaces. You can globally configure all of the ports in a specified gateway FC fabric to be FCoE trusted. This reduces system overhead by eliminating the need for filters. The total number of FCoE sessions (ENode to FCF sessions) the system can support is 2500 sessions. Sessions are defined as the combined number of VN\_Port to VF\_Port sessions and VN\_Port to VN\_Port sessions. (Although VN2VF and VN2VN sessions run in different FCoE VLANs, the session limit is a system limit, not a per-VLAN limit.)



**NOTE:** A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions. There is no limit to the number of end-to-end server-to-storage sessions.



**NOTE:** Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports and terminates the existing sessions. Changing the fabric ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate FIP snooping filters.

## Maximum Number of FCoE Sessions Per ENode

---

You can configure the maximum number of FCoE session logins from each ENode that are permitted on the gateway FC fabric. The number of sessions is the ENode FLOGI session plus the VN\_Port FDISC sessions on that ENode. Regardless of whether the fabric is trusted or untrusted, the maximum number of FCoE sessions per ENode is 2500 sessions. The total number of sessions cannot exceed the gateway fabric's maximum limit of 2500 sessions.

The maximum number of FCoE sessions per ENode is a global configuration for all members of the gateway FC fabric and cannot be configured on a per-interface basis.



**NOTE:** Session does not refer to end-to-end server-to-storage sessions. There is no limit to the number of end-to-end server-to-storage sessions.

## Priority

---

When the FIP discovery process offers an ENode the choice of more than one FCF-MAC on a given FCF to use for login, the ENode chooses the FCF-MAC to which to send a login request based on the FCF-MAC priority. The lower the priority number, the higher the FCF-MAC's priority. The ENode selects the highest-priority (lowest priority number) FCF-MAC for the login request.

An ENode can receive multiple FCF-MAC advertisements from the same FCF in two ways:

- During the FIP discovery process, an FCF can receive an ENode MAC's multicast discovery solicitation on multiple FCF-MACs. Each FCF-MAC replies with a unicast discovery advertisement to the ENode. The ENode determines that the advertisements are from the same FCF, because the value in the Name\_Identifier descriptor is the same in each advertisement.
- During the FIP discovery process, an ENode MAC can receive unsolicited multicast discovery advertisements from multiple FCF-MACs on the same FCF. The ENode determines that the advertisements are from the same FCF, because the value in the Name\_Identifier descriptor is the same in each advertisement.

On the gateway, you can configure the priority value for an entire fabric or for an individual interface. The default value for both the fabric and the individual interfaces is 128 (the highest priority is 0; the lowest priority is 255).

#### Related Documentation

- [Overview of FIP on page 4960](#)
- [Understanding FIP Functions on page 4984](#)
- [Understanding FIP Implementation on page 4988](#)
- [Understanding Fibre Channel Virtual Links on page 4995](#)
- [Understanding FCoE on page 4968](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)

## Understanding Fibre Channel Virtual Links

A virtual link emulates a secure point-to-point connection between the virtual node port (VN\_Port) of a Fibre Channel over Ethernet (FCoE) node (ENode) and the virtual fabric port (VF\_Port) of an FCoE forwarder (FCF). The combination of the FCF media access control (MAC) address and the VN\_Port MAC address uniquely identifies each virtual link. Uniquely identifying each virtual link enables the logical separation of traffic that belongs to each virtual link. A single physical link between an ENode and an FCF can carry multiple virtual links and maintain secure, separate transport of traffic on the different virtual links.

Virtual links are necessary because Fibre Channel protocol does not recognize multipoint-to-point connections. Even when multiple connections are aggregated on one physical port, FCoE Initialization Protocol (FIP) presents each virtual link as an individual point-to-point link between an ENode VN\_Port and an FCF VF\_Port.

#### Related Documentation

- [Overview of FIP on page 4960](#)
- [Understanding FIP Functions on page 4984](#)
- [Understanding FIP Implementation on page 4988](#)
- [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)
- [Understanding FCoE on page 4968](#)

## Understanding Interfaces on an FCoE-FC Gateway

When the QFX Series functions as an FCoE-FC gateway to connect FCoE devices on an Ethernet network to a Fibre Channel (FC) switch in a storage area network (SAN), it handles FCoE traffic from hosts and native FC traffic from the FC switch. To support this architecture, each local FC fabric configured on the gateway (in the **fc-fabrics** configuration hierarchy) must have:

- An Ethernet-network-facing F\_Port interface for the FCoE VLAN to connect to FCoE device VN\_Ports in the form of an FCoE VLAN interface. Multiple VF\_Ports are initiated on the F\_Port interface, one VF\_Port for each ENode that logs in to the FC network.
- One or two blocks of six proxy N\_Port (NP\_Port) interfaces to connect to FC switch fabric ports (F\_Ports).

Each FC fabric is local to the gateway on which you configure it. This means that both the FC switch and the FCoE devices must be connected to the same gateway (QFX3500 switch or QFabric system Node device), and that all of the interfaces configured for the local fabric also must be on that gateway. FC fabric traffic does not flow between different Node devices in a QFabric system.

This topic describes:

- [Native FC Interfaces to the FC Switch on page 4996](#)
- [FIP Login Session Limits on page 4998](#)
- [Trusted and Untrusted Interfaces on page 5001](#)
- [Buffer-to-Buffer Credit Recovery on page 5002](#)
- [FCoE VLAN Interface to FCoE Devices on page 5003](#)
- [Assigning Interfaces to a Fibre Channel Fabric on page 5006](#)
- [Deleting a Fibre Channel Interface on page 5006](#)

---

### Native FC Interfaces to the FC Switch

You must configure either 6 or 12 of the physical interfaces on the gateway as native FC NP\_Port interfaces to connect to FC switch F\_Port interfaces. By default, all of the gateway interfaces are Ethernet interfaces, so you must explicitly configure the interfaces that you want to use as FC interfaces.

You can configure the FC-capable ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to create blocks of native FC interfaces. You cannot individually configure a single port as a native FC interface. Within these port blocks, you cannot mix FC interfaces with Ethernet interfaces. All of the ports in a block must be either native FC interfaces or Ethernet interfaces.

You cannot configure ports xe-0/0/6 through xe-0/0/41 and ports xe-0/1/1 through xe-0/1/15 as native FC ports; they can only be Ethernet ports. Native FC ports do not handle Ethernet traffic (including FCoE traffic); they handle only native FC traffic and must connect to native FC ports.



You can configure:

- Six native FC interfaces by configuring either ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 or ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- Twelve native FC interfaces by configuring ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- No native FC interfaces by leaving ports xe-0/0/0 through xe-0/0/5 and ports xe-0/0/42 through xe-0/0/47 in their default state as Ethernet interfaces.

Each native FC interface can belong to only one local FC fabric configured on the gateway. You can configure up to 12 FC fabrics on a gateway, but each FC fabric must use different native FC interfaces to connect to an FCF. (Although the native FC ports are configured in blocks, each individual port can belong to a different FC fabric.) Native FC interfaces can be configured as loopback interfaces.

- [Port Mode on page 4997](#)
- [NPIV on page 4997](#)
- [Port Speed on page 4998](#)

### **Port Mode**

The gateway presents a proxy N\_Port (NP\_Port) interface to the FC switch. An NP\_Port connects to a single FC switch F\_Port using a point-to-point link (in other architectures an N\_Port can also connect in a point-to-point link to another N\_Port, but that is not a valid configuration on the gateway).

You must explicitly configure each native FC interface connected to an FC switch as an NP\_Port. The gateway NP\_Ports act as a proxy for the FCoE device virtual N\_Ports (VN\_Ports) when the VN\_Ports attempt to connect to the FC switch.

When the FC switch is a trusted switch, configure the fabric as **fcoe-trusted** to reduce overhead caused by the VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping filters that are automatically installed on untrusted ports.

### **NPIV**

FC requires a unique point-to-point link between the FC switch and each host N\_Port. The gateway creates an independent virtual link for each FCoE device session by mapping each FCoE device to a virtualized N\_Port through the gateway's proxy function. This process is called N\_Port ID virtualization (NPIV).

NPIV makes each virtual link look like a dedicated point-to-point link to the FC switch. In this way, multiple FCoE devices, multiple applications, and multiple virtual machines on an FCoE device can connect to an FC switch using one physical port instead of using a physical port for each host connection. The virtual link creates a secure boundary between traffic from different sources that are on a single physical port.

FCoE-FC gateway mode implements NPIV as follows:

1. An NP\_Port on the gateway comes up and logs in to the attached F\_Port on the FC switch. The FC switch sees the gateway port as a physical FC device N\_Port and assigns it a unique FCID. This establishes the physical point-to-point link between the gateway and the FC switch.
2. The gateway receives a FIP discovery message from an FCoE device that seeks to log in to the FC network. To the FCoE device, the gateway presents a virtual F\_Port (VF\_Port) interface and appears to be an FCF.
3. The gateway converts the FCoE device's message into an FC fabric discovery (FDISC) message and sends it through the least-loaded physical NP\_Port to the FC switch. The FDISC message requests an FCID for the new virtual link.
4. The FC switch processes the request, accepts it, assigns a unique FCID for the connection, and sends the response.
5. The gateway maps the FC switch response to the host FCoE device's VN\_Port and sends a FIP acceptance advertisement to the FCoE device.
6. The FCoE device accepts the FCID.

If the FC switch rejects the FDISC, the gateway relays the rejection to the FCoE device VN\_Port.

### **Port Speed**

The gateway supports configuring FC port speeds of 2 Gbps, 4 Gbps, or 8 Gbps. FC ports can also autonegotiate the port speed to 2, 4, or 8 Gbps.

### **FIP Login Session Limits**

---

A FIP login session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. (A session here does not refer to an end-to-end server-to-storage session; there is no limit to the number of end-to-end server-to-storage sessions.) You can limit the maximum number of FIP login sessions on each gateway Node device (QFX3500 switch or QFabric system Node device configured in FCoE-FC gateway mode), on each local gateway FC fabric, and on each individual NP\_Port interface in a local FC fabric:

- Gateway Node devices and Node groups—The total number of FIP login sessions on the gateway Node or Node group (the sum of the sessions on all of the NP\_Port interfaces in all of the local FC fabrics on the gateway Node or Nodes) cannot exceed the limit. When a gateway reaches the maximum session limit, the gateway sends subsequent multicast discovery advertisements (MDAs) with the availability bit set to 0 (zero) to prevent additional ENode login attempts. If the maximum number of sessions is running on the gateway, ENodes cannot use the gateway to log in new sessions to the FC switch. When the number of sessions falls below the maximum, the gateway sets the availability bit in MDAs to 1 so that ENodes can again log in new sessions. When a session slot becomes available, the system accepts the first session request to fill the slot.
- FC fabric—The total number of FIP login sessions on an FC fabric (the sum of the sessions on all of the NP\_Port interfaces that belong to the fabric) cannot exceed the

limit. When a fabric reaches the maximum session limit, the gateway sends MDAs associated with that fabric with the availability bit set to 0 to prevent additional ENode login attempts.



**NOTE:** Other FC fabrics on the same gateway can still accept ENode logins as long as the maximum session limit for those fabrics and the maximum session limit for the gateway (the Node device) have not been met.

- **NP\_Port interfaces**—The total number of FIP login sessions cannot exceed the interface's limit. When an interface reaches the maximum session limit, the gateway removes it from the load-balancing list for that FC fabric to prevent the gateway from attempting to assign new sessions to the interface. Other interfaces in the FC fabric can still accept logins until the FC fabric or gateway reaches its maximum session limit. However, the interface that reached the maximum session limit cannot be assigned new sessions until the number of sessions on the interface falls below the limit.



**BEST PRACTICE:** Configure a maximum session limit for each NP\_Port interface that is less than or equal to the number of FIP sessions the directly connected FC switch port is configured to support. This prevents the gateway from attempting to assign new login sessions to an interface when the connected FC switch port reaches its maximum number of sessions.

- [FCoE Trusted and Untrusted Interface Session Limits on page 4999](#)
- [Configuring Consistent Session Limits on page 4999](#)
- [Decreasing Session Limits on page 5000](#)
- [Increasing Session Limits on page 5001](#)
- [Effect of Deactivating and Then Reactivating the Configuration on Session Limits on page 5001](#)

### ***FCoE Trusted and Untrusted Interface Session Limits***

The maximum number of VN2VF\_Port FCoE login sessions that each gateway can support is 2500 sessions, regardless of whether interfaces are trusted or untrusted. (In software releases earlier than Junos OS Release 12.3, the session limit on untrusted interfaces and untrusted fabrics was 376 sessions.)

### ***Configuring Consistent Session Limits***

The system does not perform commit checks to enforce consistent session limit configuration. For example, the system does not prevent you from configuring a higher limit for ENode sessions than the total session limit for the gateway Node device, or from configuring a higher limit on an interface than on the fabric to which the interface belongs.

To prevent unexpected FIP login rejections, you should configure consistent Node device, fabric, and interface session limits. For example:

- The session limit of an interface should not exceed the session limit of the fabric to which it belongs.
- For interfaces that belong to the same fabric, the sum of the interface session limits should not exceed the fabric session limit.
- The fabric session limit should not exceed the session limit of the gateway Node device.
- For fabrics that belong to the same gateway Node device, the sum of the fabric session limits should not exceed the Node device session limit.

Session limit configuration considerations include:

- The fabric session limit restricts how many sessions can run on the NP\_Port interfaces that belong to that fabric. If the combined session limits of the interfaces exceed the fabric session limit, the total number of sessions on the interfaces is the fabric limit.

For example, if a fabric has three NP\_Port interfaces, and each NP\_Port interface has a limit of 500 sessions (total of 1500 sessions for the three interfaces), but the fabric has a limit of 1000 sessions, the combined number of sessions on the three interfaces is limited to 1000 sessions.

- The gateway Node device session limit restricts how many sessions can run on the fabrics that belong to that gateway. If the combined session limits of the fabrics exceed the gateway Node device session limit, the total number of sessions on the fabrics is the gateway Node device limit.

For example, if a gateway has two fabrics, and each fabric has a limit of 1000 sessions (total of 2000 sessions for the two fabrics), but the gateway has a limit of 1500 sessions, the combined number of sessions on the two fabrics is limited to 1500 sessions.

Hierarchically, the gateway Node device session limit is the maximum limit for all sessions on the gateway, regardless of fabric and interface session limits. In the same way, the fabric session limit supersedes the interface session limit.

When session limits are exceeded, no new logins are accepted until a session slot becomes free.

### ***Decreasing Session Limits***

If you decrease the session limit, the currently logged in sessions are terminated as follows:

- Gateway Node devices and Node groups—Decreasing the session limit terminates all of the sessions on the Node device (all sessions on all interfaces on all fabrics). If the gateway Node device is part of a Node group, all sessions on all members of the Node group are terminated.
- Fabric—Decreasing the session limit terminates all of the sessions on all of the interfaces that belong to the fabric.
- NP\_Port interfaces—Decreasing the session limit terminates all of the sessions on the interface and also terminates all of the sessions on any other interfaces that belong to the same fabric.

After you decrease a session limit, the sessions are terminated even if the new session limit is greater than the number of currently active sessions. For example:

- An interface has 300 active sessions.
- The current session limit is 1000 sessions.
- You decrease the session limit to 500 sessions and commit the new configuration.
- All 300 sessions are logged out, even though the new session limit is greater than the number of sessions running.

After the session limit change takes effect, the ENodes log in again and establish new sessions, up to the new session limits.

### ***Increasing Session Limits***

Increasing the session limits does not disrupt logged in sessions.

### ***Effect of Deactivating and Then Reactivating the Configuration on Session Limits***

If you decrease session limits, all ENodes are logged out. Deactivating and then reactivating the configuration can have the same effect as decreasing the session limit, which results in the ENodes being logged out.

The ENode logouts occur because when you deactivate the configuration, the system reverts to the default session limit of 2500 sessions (the maximum number of sessions). When you reactivate the configuration, the system uses the configured session limit. Unless the configured session limit is equal to the maximum session limit, reactivating the configuration decreases the session limit, which causes the ENodes to be logged out.

For example, if you:

1. Configure and commit a limit of 400 sessions.
2. Allow ENodes to log in and start sessions.
3. Deactivate the configuration.
4. Reactivate the configuration.
5. The ENode sessions are logged out because deactivating the session increased the session limit from 400 to 2500.

Because an increase in the session limit does not affect existing sessions, the running ENode sessions are not affected. However, reactivating the configuration decreased the session limit from 2500 back to 400. The session limit decrease causes the ENode sessions to be logged out.

### ***Trusted and Untrusted Interfaces***

By default, gateway fabric interfaces are untrusted interfaces. If you do not configure a gateway fabric as an FCoE trusted fabric to set all of the gateway fabric interfaces as trusted interfaces, the gateway installs VN2VF\_Port FIP snooping filters on the fabric ports.

If you configure a gateway fabric as an FCoE trusted fabric, the gateway does not install VN2VF\_Port FIP snooping filters on the fabric interfaces. This is usually done when the gateway is connected to an FCoE transit switch that has VN2VF\_Port FIP snooping enabled.

Regardless of whether an interface is trusted or untrusted, the maximum session limit is 2500 sessions.



**NOTE:** The session limit for a Node group is the same as the session limit for an individual Node device, 2500 sessions. Even if more than one Node device in a Node group is acting as an FCoE-FC gateway, the total maximum number of sessions on all Node devices in the Node group is 2500 sessions.

The default maximum login session value for Node devices (on QFabric systems, the maximum applies to each Node device), trusted fabrics, and interfaces in trusted fabrics is 2500 sessions.

### Buffer-to-Buffer Credit Recovery

---

Buffer-to-buffer credits represent the number of receive buffers an interface can use to store FC frames. Buffer-to-buffer credit determines buffer-to-buffer flow control. When an interface transmits a frame, it decrements its buffer-to-buffer credit count by one. When the destination interface forwards the frame and frees a buffer, it sends a receiver ready (R\_RDY) primitive to the transmitting interface. Each R\_RDY primitive the transmitting interface receives increments its buffer-to-buffer credit count by one.

Both interfaces on an FC link track buffer-to-buffer credits. As long as buffer-to-buffer credits are available, the transmitter continues to send frames. If the number of buffer-to-buffer credits reaches zero (0), transmission stops until buffer-to-buffer credits are available, as indicated by the reception of an R\_RDY primitive. Buffer-to-buffer credits can compensate for long cable distances to limit throughput and prevent buffer overflow.

However, if frame corruption or errors transmitting R\_RDY primitives occur, the buffer-to-buffer credit counters on the sending and receiving interfaces do not have the same values. This causes the permanent loss of buffer-to-buffer credits. When credits are lost, the buffer credit count can decrement to zero and indicate that there is no available buffer space even if buffer space is actually available. This can result in unnecessary link idle time.

To recover lost buffer-to-buffer credits, you can configure a buffer-to-buffer credit state change number (BB\_SC\_N). BB\_SC\_N must be configured on both ends of the connection. If only one end of the connection is configured for BB\_SC\_N, the feature is disabled. The two directly connected FC interfaces communicate the BB\_SC\_N value during fabric login (FLOGI).

When you enable BB\_SC\_N on the interfaces on both ends of an FC link, the interfaces exchange buffer-to-buffer state change send (BB\_SCs) and buffer-to-buffer state change receive (BB\_SCr) primitives to track the number of frames sent and the number of R\_RDY primitives received. The state change number determines the number of frames and R\_RDY primitives the interfaces exchange between consecutive BB\_SCn primitives and

between consecutive BB\_SCr primitives. The state change primitives inform each interface of the other interface's frame count and R\_RDY count states.

The state counters should match so that each interface knows and agrees with the other interface's state. If the interface at either end of the link detects a discrepancy, it knows that a frame or an R\_RDY primitive was corrupted or dropped.

For example, if a receiving interface has sent two R\_RDY primitives but the BB\_SCr that the interface receives from the sending interface only counts one R\_RDY primitive received, it reveals that one R\_RDY primitive was not delivered successfully and that one buffer-to-buffer credit was lost. When one of the interfaces on the link detects a discrepancy, the interfaces can take corrective action and recover the lost buffer-to-buffer credits.

Enabling the buffer-to-buffer credit recovery feature does not impact buffer resources and has an insignificant impact on processing resources.

If buffer-to-buffer credit recovery is not used, then when there is no buffer credit on a port, a timeout and recovery mechanism prevents buffer overflow.

### FCoE VLAN Interface to FCoE Devices

Each FC fabric configured on the gateway includes at least one FCoE VLAN interface to connect the FCoE devices on the FCoE VLAN to the FC switch. (Including the FCoE VLAN interface and the native FC interfaces in the FC fabric configuration connects them.) FCoE VLANs can include any Ethernet interface on the switch that is in tagged-access or trunk mode. The best practice is to configure Ethernet interfaces that belong to FCoE VLANs in **tagged-access** port mode.



**NOTE:** The Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.

FCoE VLANs should carry only FCoE traffic. You should not mix FCoE traffic and standard Ethernet traffic on the same VLAN.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

Each FCoE VLAN interface can belong to only one FC fabric configured on the gateway. A gateway FC fabric can have more than one FCoE VLAN, but each FCoE VLAN in the FC fabric must belong only to that FC fabric. You can configure more than one FC fabric on a gateway; each FC fabric must use different FCoE VLAN interfaces to connect to FCoE devices.



**NOTE:** Storm control must be disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped.

- [Port Mode on page 5004](#)
- [Disabling Storm Control on FCoE Interfaces on page 5005](#)
- [NPIV Support on page 5006](#)
- [VN2VF\\_Port FIP Snooping on page 5006](#)

### **Port Mode**

You must explicitly configure the FCoE VLAN interface in F\_Port mode. All members of the FCoE VLAN use the FCoE VLAN interface as the connection to the gateway NP\_Port interfaces and ultimately to the FC switch.

All of the 10-Gigabit Ethernet interfaces that are members of an FCoE VLAN should be configured as **tagged-access** port mode interfaces. However, the system also supports configuring these interfaces in **trunk** port mode.



**BEST PRACTICE:** Use **tagged-access** port mode for Ethernet interfaces that are connected to converged network adapters (CNAs) in FCoE access devices.

Use trunk port mode when an Ethernet interface is an interswitch link (ISL)—that is, when the port is connected to another switch. For example, if a port is connected to a transit switch that is performing VN2VF\_Port FIP snooping, configure the port in trunk mode and as an FCoE trusted port.

The **tagged-access** port mode was not available in Junos OS Release 11.3 and earlier releases. In Release 11.3 and earlier, only trunk port mode was used for Ethernet interfaces that belong to an FCoE VLAN. Because **tagged-access** mode is now available, using trunk mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses trunk mode for ports connected to FCoE CNAs, you can change the port mode to **tagged-access** without disrupting traffic. Although we recommend changing the port mode of these ports from trunk mode to **tagged-access** mode as a best practice, it is not mandatory. New configurations should use **tagged-access** mode for interfaces that connect to FCoE devices.

There are several advantages of configuring Ethernet ports connected to FCoE devices in **tagged-access** mode instead of in **trunk** mode:

- It is standard practice to configure ISL ports as trunk ports.
- It is standard practice not to configure ports that connect to servers as trunk ports.
- When an interface goes down, if that interface is in **trunk** mode, then the FCoE sessions on that interface are terminated only after the gateway stops receiving FIP keepalive



messages from the ENode and exceeds 2.5 times the FIP keepalive timeout advertisement value. If the interface is in **tagged-access** mode and the interface goes down, the gateway sends a FIP message to terminate the sessions on the interface.

- Similarly, if an ENode session moves from one interface to another interface, if the original interface is in **trunk** mode, the session is not removed from the interface until the gateway stops receiving FIP keepalive messages and exceeds 2.5 times the FIP keepalive advertisement timeout value. But if the interface is in **tagged-access** mode, the gateway detects that the session is no longer on the interface, does not refresh the FIP keepalive timer, and thus ages out the session.



**NOTE:** FIP is enabled on the FCoE VLAN, which is a Layer 3 interface. As with other Layer 3 interfaces under Junos OS, when the last member (10-Gigabit Ethernet interface) of the FCoE VLAN is deleted, the FCoE VLAN interface is internally marked as “down.” When the Layer 3 FCoE VLAN interface is marked as “down”, FIP stops running on it. When the last member interface is deleted from an FCoE VLAN and FIP stops running, the result could be an immediate timeout for the VN\_Ports that were connected on that interface, regardless of whether the port mode is **tagged-access** or **trunk**.

### ***Disabling Storm Control on FCoE Interfaces***

Storm control is not supported on the FCoE interfaces of an FCoE-FC gateway VLAN. Enabling storm control on an FCoE-FC gateway VLAN interface may cause FCoE packet loss. Storm control is disabled by default on all interfaces. However, if you enabled storm control globally on all switch interfaces or on any interfaces that are part of the FCoE VLAN interface, you must disable storm control on the Ethernet interfaces of the FCoE VLAN.

If storm control is enabled on only a few interfaces of the FCoE VLAN, you can disable storm control on individual interfaces by including the **set ethernet-switching-options storm-control interface *interface-name*** statement in the configuration, where ***interface-name*** is the name of the interface on which you want to disable storm control.

If storm control is enabled globally on the switch when the switch is acting as an FCoE-FC gateway, it is often easiest to disable storm control on all interfaces, then enable storm control only on Ethernet interfaces that are not part of the FCoE VLAN interface.

If storm control is enabled globally, you can disable storm control in either of two ways:

- Disable storm control on all interfaces, then enable storm control on the interfaces you want to use storm control. (From the default configuration, you cannot disable storm control on individual interfaces because the default configuration enables storm control on **all** interfaces, not on individual interfaces.)

For example, if you want interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22 to use storm control, disable storm control on all interfaces, then enable storm control on those three interfaces:

1. Disable storm control on all interfaces:

```
user@switch# delete ethernet-switching-options storm-control interface all
```

2. Enable storm control on interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22:

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/20
```

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/21
```

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/22
```

- Disable storm control for all unknown unicast traffic on all interfaces by including the following statement in your configuration:

```
user@switch# set ethernet-switching-options storm-control interface all no-unknown-unicast
```

### ***NPIV Support***

The gateway supports FCoE device NPIV. For example, a single physical FCoE device can have multiple virtual machines running on it. Each virtual machine can instantiate a separate virtual connection to the gateway, which results in its own virtual link to the FC switch. In this way, an FCoE device can have multiple separate connections to the FC SAN on a single physical port.

This is similar to the NPIV function the gateway performs with the FC switch to support multiple virtual FCoE device connections on one physical NP\_Port.

The gateway presents multiple VF\_Port interfaces on each FCoE VLAN interface to support the requirement for unique, secure virtual links.

### ***VN2VF\_Port FIP Snooping***

The FCoE-facing ports that belong to an FCoE VLAN on a gateway are enabled for VN2VF\_Port FIP snooping automatically. You can disable VN2VF\_Port FIP snooping on any individual interface by configuring it as a trusted interface.

## ***Assigning Interfaces to a Fibre Channel Fabric***

---

You assign at least one FCoE VLAN interface and at least one native FC interface to each FC fabric you configure on the gateway. All of the interfaces that belong to an FC fabric must reside on the same gateway device. Interfaces on different gateways cannot belong to the same FC fabric, because an FC fabric is local to a single gateway device.

## ***Deleting a Fibre Channel Interface***

---

To delete an FC interface or an FCoE VLAN interface, you must delete the interface from the fabric first and then delete the interface from the switch.

### **Related Documentation**

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding Fibre Channel on page 4961](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Understanding FCoE-FC Gateway Benefits](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Configuring a Fibre Channel Interface on page 5182](#)
- [Disabling VN2VF\\_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)

- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway on page 5189](#)
- [Disabling VN2VF\\_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)
- [Deleting a Fibre Channel Interface on page 5188](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197](#)
- [Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

## Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric

You can balance or rebalance the load on the ports in an FCoE-FC gateway proxy fabric in order to avoid overutilizing or underutilizing the links. Load balancing is distributing sessions across the available native Fibre Channel (FC) interfaces (NP\_Ports) that belong to a local gateway FC fabric to create a relatively equal load on all the fabric links. Load rebalancing is redistributing the existing sessions across the available NP\_Port links on a local gateway FC fabric.



**NOTE:** A session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.

The fabric-facing NP\_Port links of an FCoE-FC gateway use different load-balancing methods than the FCoE-network-facing Ethernet links. You can balance the load on the FCoE-network facing Ethernet port members of the FCoE-FC gateway by configuring the ports in a link aggregation group (LAG). This topic focuses on NP\_Port load balancing.

Balancing the load on FCoE-FC gateway NP\_Port links consists of two steps:

1. Choosing the algorithm used to balance and rebalance the link load
2. Choosing whether to rebalance link loads automatically or only when you explicitly request a rebalance (load-rebalancing method)

You can configure a different load-balancing algorithm and use a different rebalancing method for each local FC fabric on the FCoE-FC gateway. The load-balancing algorithm and automated rebalancing, if configured, apply to all NP\_Port interfaces in the local FC fabric.

This topic describes:

- [Load-Balancing Algorithms on page 5009](#)
- [Load-Rebalancing Methods on page 5013](#)
- [NP\\_Port Interface FIP Session Limit Effect on Load Balancing on page 5014](#)
- [Load-Balancing Triggers and Timing on page 5014](#)
- [Load Rebalancing Behavior When a Link Goes Down on page 5016](#)
- [Interface Load Calculation Algorithm on page 5017](#)
- [Load-Balancing Scenarios on page 5018](#)
- [Load Balancing on the FCoE Interfaces \(Ethernet Links\) on page 5023](#)

## Load-Balancing Algorithms

You can choose one of three load-balancing algorithms to configure the way the switch balances the link loads. The switch uses the configured algorithm to balance the link loads when NP\_Ports are initialized and whenever link loads are rebalanced. Regardless of whether you configure automated load rebalancing or use on-demand load rebalancing, the switch uses the configured algorithm to balance the link load:

- **Simple load balancing**—The switch assigns each ENode FLOGI session and VN\_Port FDISC session to the least-loaded link. The switch can place FDISC sessions on a different link than the parent FLOGI session (an ENode FLOGI session and its subsequent FDISC sessions can be placed on different links). Simple load balancing is the default load-balancing algorithm. Rebalancing the link load disrupts only selected sessions to minimize the impact (the switch uses an algorithm to log out only the sessions that need to be moved to other links to balance the load when those sessions log in again).
- **ENode-based load balancing**—When an ENode logs in to the fabric, the switch places all subsequent VN\_Port FDISC sessions associated with that ENode on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. The switch calculates the link load based on the combined total of FLOGIs and FDISCs on each NP\_Port link. Rebalancing the link load disrupts all sessions (all sessions log out and then log in again).
- **FLOGI-based load balancing**—Similar to ENode-based load balancing; when an ENode logs in to the fabric, the switch places all subsequent VN\_Port FDISC sessions associated with that ENode on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link.

One difference between ENode-based load balancing and FLOGI-based load balancing is that the switch calculates the link load based only on the number of FLOGIs on each NP\_Port link. The algorithm does not count FDISCs. Another difference is that instead of disrupting all sessions on a link load rebalance, the system disrupts only selected sessions to minimize the impact (the switch uses an algorithm to log out only the sessions that need to be moved to other links to balance the load when those sessions log in again).



**NOTE:** Changing the load-balancing algorithm when FCoE sessions are running forces the FCoE sessions to log out and then log in again.

If you do not explicitly configure the load-balancing algorithm, the switch uses simple load balancing by default on the all NP\_Port interfaces that belong to a given local FC fabric.

The following sections describe how each algorithm works, its advantages and disadvantages, and what happens when NP\_Port links come up for the first time, when an NP\_Port link is added to existing links, and when you rebalance the link load:

- [Simple Load Balancing on page 5010](#)
- [ENode-Based Load Balancing on page 5011](#)
- [FLOGI-Based Load Balancing on page 5011](#)
- [Load-Balancing Algorithm Comparison on page 5012](#)

### **Simple Load Balancing**

Simple load balancing provides the most equal load balancing across links because each VN\_Port FDISC session can be assigned to the least-loaded link, regardless of whether the parent ENode FLOGI session is on that link. (The parent ENode is the ENode that originates the logins to the fabric. After the parent ENode logs in, the VN\_Ports on that ENode can log in to the fabric using FDISC.)

The FCoE-FC gateway performs simple load balancing by default on the NP\_Ports that connect the gateway to the FC SAN. When an ENode sends a FLOGI request to the gateway, the gateway checks the NP\_Ports that connect it to the FC SAN and assigns the new session to the least-loaded link.

Every time an ENode sends a FLOGI or an FDISC request, the gateway assigns the new session to the least-loaded NP\_Port link. After the gateway assigns an ENode FLOGI session to an NP\_Port, subsequent FDISC requests by the same ENode can result in sessions being assigned to different NP\_Ports, because the gateway always assigns the new session to the least-loaded interface.



**NOTE:** Because VN\_Port sessions might be placed on a different link than their parent ENode, if the link that contains the ENode goes down, only the ENode session and any of its VN\_Port sessions that are on that link go down. VN\_Port sessions on other links remain active as long as the link is up and the VN\_Port is not logged out.

---

When a new link comes up, the switch logs out enough sessions so that when the sessions log in again, they are placed on the new link and the link loads are balanced. The switch uses an algorithm to log out sessions in the least disruptive manner by first logging out FDISCs whose FLOGI is not on the same link, then the least-loaded FLOGIs (loaded in terms of related FDISC logins).

Similarly, when you rebalance an existing link load, the switch logs out only enough sessions so that when the sessions log in again, they balance the load on the existing links. In this case (rebalance without a new link up), the switch takes into account the dependencies between FLOGIs and FDISCs when selecting sessions to log out.

The simple load-balancing algorithm uses the sum of the FLOGI and FDISC sessions to determine the session load on each link for both initial load balancing and load rebalancing.

### ***ENode-Based Load Balancing***

ENode-based load balancing can result in a less balanced load across the NP\_Port links because the VN\_Port FDISC sessions are assigned to the same link as the parent ENode FLOGI session, regardless of how many FDISC sessions are associated with the ENode. However, ENode-based load balancing has the advantage of keeping all of the sessions associated with a particular ENode on one link, which provides better control and predictability.

When you use the ENode-based load-balancing algorithm, the gateway assigns the ENode to an NP\_Port link when the ENode sends its FLOGI message to the gateway. The gateway places the ENode session on the least-loaded link at that time. The VN\_Port FDISC sessions associated with an ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. Essentially, the ENode sessions are load-balanced, but the VN\_Port sessions are not.

ENode-based load balancing ensures that each ENode and its associated VN\_Port sessions are assigned to the same NP\_Port link. ENode-based load balancing provides more control and predictability and ensures that if the link carrying an ENode goes down, all of the ENodes associated VN\_Port sessions also go down.

The disadvantage of ENode-based load balancing is that if one ENode has a large number of sessions and the other ENodes do not, the link that carries the ENode with the large number of sessions might have a much larger load than the other NP\_Port links in the gateway proxy fabric.

For example, if a gateway fabric has two NP\_Ports connected to the FC fabric, and two ENodes log in to the fabric, one ENode session is placed on each link. If two VN\_Port sessions are initiated on one of the ENodes, those sessions are placed on the same link as the parent ENode. If 1000 VN\_Port sessions are initiated on the other ENode, all of the 1000 VN\_Port sessions are placed on the same link as that ENode. In this case, one link has 3 sessions (1 ENode FLOGI session and 2 VN\_Port FDISC sessions) and the other link has 1001 sessions (1 ENode FLOGI session and 1000 VN\_Port FDISC sessions).

When a new link comes up or when you rebalance an existing load, the switch logs out all sessions (FLOGIs and FDISCs) in the fabric. As the sessions log in again, the switch assigns them to NP\_Ports in a balanced manner, with all FDISCs assigned to the same link as the parent FLOGI. A new link coming up or a rebalance disrupts all of the existing sessions.

The ENode-based load-balancing algorithm uses the sum of the FLOGI and FDISC sessions to determine the session load on each link for both initial load balancing and load rebalancing.

### ***FLOGI-Based Load Balancing***

FLOGI-based load balancing is similar to ENode-based load balancing in most ways:

- It can result in a less balanced load across the NP\_Port links because the VN\_Port FDISC sessions are assigned to the same link as the parent ENode FLOGI session, regardless of how many FDISC sessions are associated with the ENode.
- When an ENode logs in with a FLOGI, the gateway places the session on the least-loaded link, and the FDISC logins associated with the FLOGI are placed on the same link, regardless of link load.
- Provides control and predictability because each ENode and its associated VN\_Port (FDISC) sessions are assigned to the same link, so if the link an ENode is on goes down, all of its associated sessions also go down.
- If one ENode has a large number of sessions and the other ENodes do not, the link that carries the ENode with the large number of sessions might have a much larger load than the other NP\_Port links in the gateway proxy fabric.

FLOGI-based load balancing differs from ENode-based load balancing in two important ways:

1. The switch uses the sum of the FLOGI sessions on a link to determine the link load. The switch does not use FDISC sessions when calculating the number of sessions on a link. (ENode-based load balancing uses the sum of the FLOGI and FDISC sessions to calculate the number of sessions on a link.)
2. When a new link comes up or when you rebalance an existing load, the switch logs out enough FLOGI (and FDISC) sessions so that when the FLOGI sessions log in again, the load is balanced. The switch balances the load based only on the number of FLOGI sessions, not the sum of FLOGI and FDISC sessions. However, the FDISC sessions associated with a FLOGI follow the FLOGI to the new link if the FLOGI session is part of the rebalancing.

The FLOGI-based load-balancing algorithm uses only the FLOGI sessions to determine the session load on each link for both initial load balancing and load rebalancing.

### ***Load-Balancing Algorithm Comparison***

[Table 387 on page 5012](#) compares the three load-balancing algorithms and summarizes their differences, advantages, and disadvantages.

**Table 387: Load-Balancing Algorithm Comparison**

| Load-Balancing Algorithm   | Session Assignment                                                            | Session Disruption on Rebalance                                                                                           | Session Count Method            | Advantages                                                                                                                                                                                               | Disadvantages                                                                               |
|----------------------------|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Simple (default algorithm) | FDISC sessions can be placed on different links than the parent FLOGI session | Minimum number of selected sessions logged out (FDISC sessions can be logged out independent of the parent FLOGI session) | Sum of FLOGI and FDISC sessions | <ul style="list-style-type: none"> <li>• Most equal session distribution across links</li> <li>• Minimum number of sessions logged out when rebalancing</li> <li>• Least disruptive algorithm</li> </ul> | <ul style="list-style-type: none"> <li>• Less session control and predictability</li> </ul> |



Table 387: Load-Balancing Algorithm Comparison (*continued*)

| Load-Balancing Algorithm | Session Assignment                                                            | Session Disruption on Rebalance                                                                                        | Session Count Method                                                   | Advantages                                                                                                                                                                                                                | Disadvantages                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENode-based              | FDISC sessions are always placed on the same link as the parent FLOGI session | All sessions are logged out                                                                                            | Sum of FLOGI and FDISC sessions                                        | <ul style="list-style-type: none"> <li>Better session control and predictability (on link down, all sessions associated with an ENode go down)</li> </ul>                                                                 | <ul style="list-style-type: none"> <li>Most disruptive algorithm; all sessions logged out on rebalance</li> <li>Might result in less balanced link load because FDISCs are placed on the same link as parent FLOGI</li> </ul> |
| FLOGI-based              | FDISC sessions are always placed on the same link as the parent FLOGI session | Minimum number of selected sessions logged out (but FDISC sessions logged out when parent FLOGI session is logged out) | FLOGI sessions only (FDISC sessions not included in the session count) | <ul style="list-style-type: none"> <li>Better session control and predictability (on link down, all sessions associated with an ENode go down)</li> <li>Minimum number of sessions logged out when rebalancing</li> </ul> | <ul style="list-style-type: none"> <li>Might result in less balanced link load because FDISCs are placed on the same link as parent FLOGI</li> </ul>                                                                          |

### Load-Rebalancing Methods

The load-rebalancing method determines the way the system redistributes sessions to balance the load on the NP\_Ports that belong to a local FC fabric on an FCoE-FC gateway.

You can rebalance the existing load on existing NP\_Port links using either of two methods:

- Automated load rebalancing—When a load rebalancing trigger occurs, the switch automatically rebalances the link loads by redistributing the sessions across the active NP\_Port links. There are three possible load rebalancing triggers:

- When you enable automated load rebalancing, the switch checks the load balance on the existing NP\_Port links. If the links are already balanced, the switch does not rebalance the link load. If the links are not balanced, the switch rebalances the link loads using the configured load-balancing algorithm.

Enabling automated load rebalancing causes sessions to be logged out in accordance with the configured load-balancing algorithm if the link load is unbalanced. If the link load is already balanced when you enable automated load rebalancing, the links are not rebalanced. (Disabling automated load rebalancing is not disruptive because the link load is already balanced.)

- When a new NP\_Port link comes up on a local FCoE-FC gateway fabric, the switch rebalances the link load using the configured load-balancing algorithm if automated load balancing is enabled.
- When the port speed is changed (unless the port speed change does not change the actual port speed, for example, changing the port speed from auto to 8 Gbps).

Use automated load rebalancing if you want link loads to be rebalanced automatically when a load-balancing trigger occurs, instead of at times of your choosing. Keep in mind that load rebalancing is a disruptive event (sessions are logged out).

- On-demand load rebalancing—You choose when to rebalance the NP\_Port links by explicitly requesting a load rebalance using an operational command. The system rebalances the link load only when you issue the rebalancing command.

Use on-demand load rebalancing if you only want to rebalance the link load once or if you want to rebalance the link loads at controlled times instead of automatically.

You can also request a load rebalancing *dry run*. A dry run simulates rebalancing and lists the sessions that might be affected if you choose to perform an actual load-rebalancing operation. The link loads are not rebalanced when you request a dry run.

---

### NP\_Port Interface FIP Session Limit Effect on Load Balancing

The maximum number of FIP login sessions configured for each NP\_Port interface affects load balancing. When an interface reaches its maximum number of FIP login sessions, that interface is removed from the list of interfaces used for load balancing. The other interfaces in the gateway fabric continue to accept ENode login sessions until they reach their configured maximum session limit. Only interfaces that have not reached their maximum session limit are included in the load-balancing calculations.



**NOTE:** If all NP\_Port interfaces in a gateway fabric reach their FIP login session limits, the fabric sends subsequent multicast discovery advertisements (MDAs) with the availability bit set to 0 (zero) to prevent additional ENode login attempts. While the maximum number of sessions is running on the gateway fabric, ENodes cannot use that fabric to log in to the FC switch. When the number of sessions falls below the maximum, the gateway sets the availability bit in MDAs to 1 so that ENodes can log in to the fabric again.

---

### Load-Balancing Triggers and Timing

Several events trigger load balancing. Some of the events trigger load balancing only when automated load balancing is enabled. Other events trigger load rebalancing whether or not automated rebalancing is enabled.

This section describes the load-balancing triggers, what happens when the trigger action occurs, and how the switch determines if and when to balance the link load:

- [Load-Balancing Triggers on page 5014](#)
- [Load-Balancing Timer on page 5015](#)

#### **Load-Balancing Triggers**

[Table 388 on page 5015](#) describes the four different events can trigger load balancing or load rebalancing. In every case, link load rebalancing uses the configured load-balancing algorithm to determine the placement of sessions on links.

Table 388: Load-Balancing Triggers and Actions

| Trigger Event                                          | Action                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| New link comes up                                      | <p>Triggers a load-rebalancing operation regardless of whether or not automated load rebalancing is enabled. (The new link has no sessions, so the sessions on other links must be redistributed to balance the load.)</p> <p>The link load is not rebalanced if there are no sessions on the existing links or if there are so few sessions on the existing links that they cannot be redistributed.</p>                                       |
| On-demand load rebalancing request issued from CLI     | <p>The switch checks the NP_Port link load. If the load is not balanced across the links, the switch rebalances the link load. If the load is already balanced, nothing happens.</p> <p><b>NOTE:</b> Requesting a dry run displays sessions that might be disrupted if you rebalance the link load, but does not rebalance the link load.</p>                                                                                                   |
| Automated load balancing configured for the first time | <p>The switch checks the NP_Port link load. If the load is not balanced across the links, the switch rebalances the link load. If the load is already balanced, nothing happens.</p>                                                                                                                                                                                                                                                            |
| NP_Port speed change                                   | <p>If automated rebalancing is enabled, changing the port speed brings the port up and down (flaps the port) and causes the switch to rebalance the link loads. If the port speed change does not change the actual port speed (for example, changing the port speed from <i>auto</i> to 8 Gbps), the link loads are not rebalanced.</p> <p>If automated rebalancing is not enabled, port speed changes do not cause link load rebalancing.</p> |



**NOTE:** When an NP\_Port link goes down, it does not trigger load rebalancing. The loads on the remaining active links are already balanced, and as the sessions logged out from the down link log in again, they are they assigned to links in a balanced manner determined by the configured load-balancing algorithm.

### Load-Balancing Timer

When you trigger load balancing from the CLI, the load-balancing action occurs immediately after you execute the command. However, when a load-balancing trigger occurs that is not a CLI command, the switch does not balance the link loads immediately. Instead, the switch follows an intelligent timer process:

1. The switch checks the current load balance on the NP\_Port links in the local gateway FC fabric. If the load is already balanced, the switch does nothing, and there is no session disruption.
2. If the check shows that the link load is not balanced, the switch starts a 10-second timer. If no other load-balancing triggers occur during the 10-second interval, the switch rebalances the load.

If another load-balancing trigger occurs during the 10-second interval, the timer resets to 10 seconds. The 10-second timer prevents the switch from performing multiple disruptive load-rebalancing actions in a short period of time.



**NOTE:** The switch processes new sessions that log in after the timer starts in the normal manner. The new sessions are considered in the load-balancing evaluation and operation.

3. At a maximum of 30 seconds after the first load-balancing trigger occurs, the switch checks the link load balance again. If the links are already balanced, the switch cancels the load-rebalancing operation. If the links are not balanced, the switch rebalances the link loads.



**NOTE:** If the trigger event that started the load-rebalancing timer is no longer valid when the timer elapses, the switch cancels the rebalancing operation. For example, if a new NP\_Port link comes up and triggers the timer, then goes down before the timer expires, the original link up event is no longer valid, and the switch cancels the rebalancing operation (unless another valid rebalancing trigger occurs in that time frame).

When a link load rebalancing operation is in progress, the switch defers any load-rebalancing triggers that occur until the load-rebalancing operation is complete. The new rebalancing operation begins after the current rebalancing operation finishes if a check shows that rebalancing is required.

If you explicitly request load rebalancing from the CLI using the **request fibre-channel proxy load-rebalance** operational command, the switch rejects the command and displays an error message stating that rebalancing is already in progress.

### Load Rebalancing Behavior When a Link Goes Down

---

If an NP\_Port link goes down, the ENode and VN\_Port sessions on that link are logged out. The ENodes and VN\_Port sessions log in again and are assigned to NP\_Port links based on the link load and the load-balancing algorithm. If a link goes down, the switch does not rebalance the remaining load on the remaining links to avoid disrupting the existing ENode and VN\_Port sessions. (Also, it is not necessary to rebalance the links in that manner because after a link goes down, the sessions on the remaining links are already balanced. As the logged out sessions log back in, the switch places them on the remaining active links in a balanced manner, according to the configured load-balancing algorithm.)



**NOTE:** When you use the simple load-balancing algorithm, an ENode and its associated VN\_Port sessions might be on different links. In that case, if the NP\_Port with the ENode goes down, only the VN\_Ports on the same link are logged out. VN\_Ports on other links remain up and running.

### Interface Load Calculation Algorithm

A weighted round-robin (WRR) algorithm determines the interface load based on:

- The current number of sessions on the interface



**NOTE:** The configured load-balancing algorithm determines how the switch counts the number of sessions. For simple and ENode-based load balancing, the number of sessions is the sum of the FLOGI and FDISC sessions on each link. For FLOGI-based load balancing, the number of sessions is the sum of the FLOGI sessions on each link.

- The interface weight, which is the speed of the Fibre Channel link (2 Gbps, 4 Gbps, or 8 Gbps)

The interface load algorithm is:

$$(\text{number-of-sessions} * \text{max-weight}) / \text{weight}$$

where *max-weight* is an internal constant.

If the load on the FC interfaces is equal, the session is assigned to the interface with the highest link speed (the greatest weight).

For example, if the three FC interfaces have the characteristics shown in [Table 389 on page 5017](#), the loads of the interfaces are not equal:

**Table 389: FC Interface Session-Based Load-Balancing Characteristics for Unequal Loads**

| Interface | Number of Sessions | Weight (Speed) |
|-----------|--------------------|----------------|
| fc-0/0/0  | 4                  | 4 Gbps         |
| fc-0/0/1  | 1                  | 2 Gbps         |
| fc-0/0/2  | 8                  | 8 Gbps         |

In this example, interfaces fc-0/0/0 and fc-0/0/2 have a greater load than fc-0/0/1. For simple load balancing, the gateway assigns the next new FLOGI or FDISC to fc-0/0/1 because it is the least-loaded interface. For both ENode-based and FLOGI-based load balancing, the gateway assigns the next new FLOGI to fc-0/0/1 because it is the least-loaded interface. Then all VN\_Port FDISCs from that ENode follow the ENode FLOGI and are also assigned to fc-0/0/1 regardless of the link load.

For another example, if the three FC interfaces have the characteristics shown in [Table 390 on page 5018](#), the loads of the interfaces are equal:

**Table 390: FC Interface Session-Based Load-Balancing Characteristics for Equal Loads**

| Interface | Number of Sessions | Weight (Speed) |
|-----------|--------------------|----------------|
| fc-0/0/0  | 4                  | 4 Gbps         |
| fc-0/0/1  | 2                  | 2 Gbps         |
| fc-0/0/2  | 8                  | 8 Gbps         |

In this case, all interfaces have the same relative load. For simple load balancing, the gateway assigns the next new FLOGI or FDISC to fc-0/0/2 because although the loads of the three interfaces are equal, fc-0/0/2 has the greatest weight. For both ENode-based and FLOGI-based load balancing, the gateway assigns the next new FLOGI to fc-0/0/2, and all VN\_Port FDISCs from that ENode follow the ENode FLOGI and are also assigned to fc-0/0/2 regardless of the link load.

After the gateway establishes a session between an ENode or a VN\_Port and an FC switch on an NP\_Port, the session remains on that NP\_Port until the ENode or VN\_Port performs a LOGO.

If the physical FC interface link goes down, the FLOGI and FDISC sessions on the down link are logged out. The ENodes and VN\_Ports log in again to start new sessions on other NP\_Ports in the local gateway FC fabric in accordance with the configured load-balancing algorithm (assuming there is more than one NP\_Port connected to the FC fabric).

### Load-Balancing Scenarios

The configured load-balancing algorithm, the sequence in which ENodes log in to the FC network, the current session count (number of sessions per interface) and the interface speed determine the way the session load is balanced across the native FC interfaces (NP\_Ports) in a gateway FC fabric. Whether you are balancing the link load for the first time or rebalancing an existing link load, the way the load is distributed across the active links is the same.



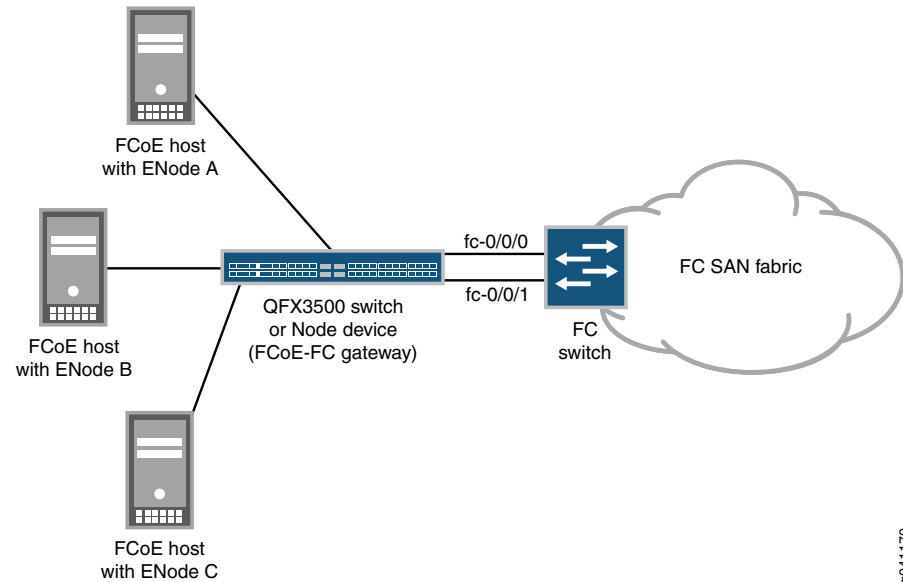
**NOTE:** The way the switch counts the number of sessions on a port depends on the load-balancing algorithm. For simple and ENode-based load balancing, the sum of the FLOGI and FDISC sessions equals the session count. For FLOGI-based load balancing, only the FLOGI sessions are counted in the total session count.

The following scenarios demonstrate how sessions are assigned to links for each load-balancing algorithm:

- [Simple Load-Balancing Algorithm Scenario on page 5019](#)
- [ENode-Based Load-Balancing Algorithm Scenarios on page 5020](#)
- [FLOGI-Based Load-Balancing Algorithm Scenarios on page 5021](#)

All of the scenarios use the topology shown in [Figure 191 on page 5019](#).

**Figure 191: Sample Load-Balancing Topology**



### **Simple Load-Balancing Algorithm Scenario**

Simple load balancing results in the most equal load distribution among the NP\_Ports connected to an FC SAN fabric because VN\_Port FDISC sessions do not need to “follow” the parent ENode FLOGI session on the same link between the gateway and the FC fabric. When a new FLOGI or FDISC session is initiated, it is assigned to the least-loaded link.

The simple load-balancing algorithm example uses the topology shown in [Figure 191 on page 5019](#) and has the following characteristics:

- QFX Series configured as an FCoE-FC gateway
- Two gateway NP\_Ports, **fc-0/0/0** and **fc-0/0/1**, connected to an FC SAN fabric switch at a speed of 8 Gbps
- Three ENodes, ENode\_A, ENode\_B, and ENode\_C connected to the gateway
- NP\_Ports **fc-0/0/0** and **fc-0/0/1**, and ENode\_A, ENode\_B, and ENode\_C, belong to the same local FC fabric on the gateway

When the NP\_Ports initialize, they send FLOGI messages to the FC switch and log in to the FC SAN fabric. The gateway then advertises the fabric to the ENodes on the Ethernet side of the network. At this point, the load on both of the NP\_Ports is equal. Now the ENodes and VN\_Ports start to log in to the fabric:

1. ENode\_A sends a FLOGI to log in to the fabric. Because the loads on the two NP\_Ports are equal, the session for ENode\_A is randomly placed on one of the links. In this example, the ENode\_A FLOGI session is placed on port **fc-0/0/0**.
2. ENode\_B logs in. Because the load is less on port **fc-0/0/1**, the ENode\_B FLOGI session is placed on port **fc-0/0/1**.

3. ENode\_C logs in. Because the link loads are equal, the ENode\_C login session is randomly placed on one of the links. In this example, the ENode\_C login session is placed on port **fc-0/0/0**.
4. A VN\_Port on ENode\_A sends an FDISC to log in to the fabric. Because port **fc-0/0/1** currently is the least-loaded link, the VN\_Port session is placed on port **fc-0/0/1**, even though its parent ENode session is on port **fc-0/0/0**.
5. As each new VN\_Port session comes up, it is placed on the least-loaded link, regardless of the link on which its parent ENode session is placed.

### ***ENode-Based Load-Balancing Algorithm Scenarios***

ENode-based load balancing ensures that VN\_Port FDISC sessions are placed on the same link as their parent ENode FLOGI sessions, regardless of the link load. ENode-based load balancing can result in a less-balanced load among the NP\_Port links, but it provides the control and predictability of keeping ENodes and their VN\_Port sessions on the same link.

The examples in this section use the topology shown in [Figure 191 on page 5019](#).

- QFX Series configured as an FCoE-FC gateway
- Two gateway NP\_Ports, **fc-0/0/0** and **fc-0/0/1**, connected to an FC SAN fabric switch at a speed of 8 Gbps
- Three ENodes connected to the gateway:
  - ENode\_A, which has 2 VN\_Port FDISC sessions
  - ENode\_B, which has 20 VN\_Port FDISC sessions
  - ENode\_C, which has 100 VN\_Port FDISC sessions
- NP\_Ports **fc-0/0/0** and **fc-0/0/1**, and ENode\_A, ENode\_B, and ENode\_C, belong to the same local FC fabric on the gateway

When the NP\_Ports initialize, they send FLOGI messages to the FC switch and log in to the FC SAN fabric. The gateway then advertises the fabric to the ENodes on the Ethernet side of the network. At this point, the load on both of the NP\_Ports is equal. Now the ENodes and VN\_Ports start to log in to the fabric. As the following two scenarios show, how these sessions are placed on the links depends on the sequence in which they log in to the fabric.

#### **Scenario 1:**

1. ENode\_A sends a FLOGI to log in to the fabric. Because the loads on the two NP\_Ports are equal, the session for ENode\_A is randomly placed on one of the links. In this example, the ENode\_A FLOGI session is placed on port **fc-0/0/0**.
2. ENode\_B logs in. Because the load is less on port **fc-0/0/1**, the ENode\_B FLOGI session is placed on port **fc-0/0/1**.
3. The two VN\_Ports on ENode\_A log in to the fabric. Their sessions are placed on port **fc-0/0/0**, following ENode\_A on the link. Now port **fc-0/0/0** has a greater load (one FLOGI session plus two FDISC sessions) than port **fc-0/0/1** (one FLOGI session).



4. The 20 VN\_Ports on ENode\_B log in to the fabric. Their sessions are placed on port **fc-0/0/1**, following ENode\_B on the link. Now port **fc-0/0/0** has a lesser load (one FLOGI, two FDISC) than port **fc-0/0/1**.
5. ENode\_C logs in. Because the load is less on port **fc-0/0/0**, the ENode\_C FLOGI session is placed on port **fc-0/0/0**.
6. The 100 VN\_Ports on ENode\_C log in to the fabric. Their sessions follow the ENode\_C session onto port **fc-0/0/0**.
7. If more VN\_Ports come up, their FDISC sessions are placed on the same link as the corresponding parent ENode session.

#### Scenario 2:

1. ENode\_A sends a FLOGI to log in to the fabric. Because the loads on the two NP\_Ports are equal, the session for ENode\_A is randomly placed on one of the links. In this example, the ENode\_A FLOGI session is placed on port **fc-0/0/0**.
2. ENode\_B logs in. Because the load is less on port **fc-0/0/1**, the ENode\_B FLOGI session is placed on port **fc-0/0/1**.
3. The two VN\_Ports on ENode\_A log in to the fabric. Their sessions are placed on port **fc-0/0/0**, following ENode\_A on the link. Now port **fc-0/0/0** has a greater load (one FLOGI session plus two FDISC sessions) than port **fc-0/0/1** (one FLOGI session).
4. In this step, the login sequence in Scenario 2 differs from the login sequence in Scenario 1, resulting in a different placement of sessions on the links, and therefore a different load on the links. ENode\_C logs in before the ENode\_B VN\_Ports log in, which changes the session count on the links compared to the first scenario. Because the load in this scenario is less on port **fc-0/0/1**, the ENode\_C FLOGI session is placed on port **fc-0/0/1** (instead of port **fc-0/0/0** as in the first scenario).
5. The 20 VN\_Ports on ENode\_B log in to the fabric. Their sessions are placed on port **fc-0/0/1**, following ENode\_B on the link. Now port **fc-0/0/0** carries one FLOGI and two FDISC sessions, and port **fc-0/0/1** carries two FLOGI and 20 FDISC sessions.
6. The 100 VN\_Ports on ENode\_C log in to the fabric. Their sessions follow the ENode\_C session onto port **fc-0/0/1**. Now port **fc-0/0/1** carries 2 FLOGI and 120 FDISC sessions, whereas port **fc-0/0/0** carries one FLOGI and two FDISC sessions.
7. If more VN\_Ports come up, their FDISC sessions are placed on the same link as the corresponding parent ENode session.

Because of the sequence of ENode logins in Scenario 2, port **fc-0/0/1** carries a greater load than port **fc-0/0/0**. If the simple load-balancing algorithm had been used, the FLOGI and FDISC sessions would be allocated to the two links evenly. However, because the FDISC sessions are placed on the same link as their parent FLOGI sessions, this example demonstrates how using the ENode-based load-balancing algorithm can lead to scenarios in which the link loads are not equal.

#### *FLOGI-Based Load-Balancing Algorithm Scenarios*

FLOGI-based load balancing is similar in many ways to ENode-based load balancing. An important difference that affects how the switch places sessions on links is that for

FLOGI-based load balancing, only the FLOGI sessions are counted when the link load is calculated. FDISC sessions are not counted to determine the link load. Because ENode-based load balancing uses the sum of the FLOGI and FDISC sessions to determine the link load, an interface with exactly the same combination of FLOGI and FDISC sessions can have a different session count depending on the algorithm used. A different session count can change the interface to which the switch assigns the next session.

As with ENode-based load balancing, FLOGI-based load balancing ensures that VN\_Port FDISC sessions are placed on the same link as their parent ENode FLOGI sessions, regardless of the link load. FLOGI-based load balancing can result in a less-balanced load among the NP\_Port links, but it provides the control and predictability of keeping ENodes and their VN\_Port sessions on the same link.

The examples in this section use the topology shown in [Figure 191 on page 5019](#).

- QFX Series configured as an FCoE-FC gateway
- Two gateway NP\_Ports, **fc-0/0/0** and **fc-0/0/1**, connected to an FC SAN fabric switch at a speed of 8 Gbps
- Three ENodes connected to the gateway:
  - ENode\_A, which has 2 VN\_Port FDISC sessions
  - ENode\_B, which has 20 VN\_Port FDISC sessions
  - ENode\_C, which has 100 VN\_Port FDISC sessions
- NP\_Ports **fc-0/0/0** and **fc-0/0/1**, and ENode\_A, ENode\_B, and ENode\_C, belong to the same local FC fabric on the gateway

When the NP\_Ports initialize, they send FLOGI messages to the FC switch and log in to the FC SAN fabric. The gateway then advertises the fabric to the ENodes on the Ethernet side of the network. At this point, the load on both of the NP\_Ports is equal. Now the ENodes and VN\_Ports start to log in to the fabric.

Because FLOGI-based load balancing does not count FDISC sessions when calculating the link load, how the sessions are placed on the link depends only on the number of FLOGI sessions per interface, not on the number of FLOGI sessions plus FDISC sessions. This means that an ENode with a FLOGI session and many FDISC sessions is counted as having the same load as an ENode with a FLOGI session and no FDISC sessions.

#### Scenario 1:

1. ENode\_A sends a FLOGI to log in to the fabric. Because the loads on the two NP\_Ports are equal, the session for ENode\_A is randomly placed on one of the links. In this example, the ENode\_A FLOGI session is placed on port **fc-0/0/0**.
2. ENode\_B logs in. Because the load is less on port **fc-0/0/1**, the ENode\_B FLOGI session is placed on port **fc-0/0/1**.
3. The two VN\_Ports on ENode\_A log in to the fabric. Their sessions are placed on port **fc-0/0/0**, following ENode\_A on the link. However, unlike simple load balancing or ENode-based load balancing, the session count of the two ports is still equal (one session each) because the FDISC sessions are not used in the session count.

4. The 20 VN\_Ports on ENode\_B log in to the fabric. Their sessions are placed on port **fc-0/0/1**, following ENode\_B on the link. Again, unlike simple load balancing or ENode-based load balancing, the session count of the two ports is still equal (one session each) because the FDISC sessions are not used in the session count.
5. ENode\_C logs in. Because the link loads are counted as equal, the ENode\_C login session is randomly placed on one of the links. In this example, the ENode\_C login session is placed on port **fc-0/0/0**.
6. The 100 VN\_Ports on ENode\_C log in to the fabric. Their sessions follow the ENode\_C session onto port **fc-0/0/0**.
7. If more VN\_Ports come up, their FDISC sessions are placed on the same link as the corresponding parent ENode session.

If a fourth ENode, ENode\_D, sends a FLOGI to log in to the fabric, it is placed on port **fc-0/0/1** because port **fc-0/0/0** has a session count of two (two FLOGIs from ENode\_A and ENode\_C, FDISCs not counted) and port **fc-0/0/1** has a session count of one (one FLOGI from ENode\_B, FDISCs not counted), so port **fc-0/0/1** is the least-loaded port.

With FLOGI-based load balancing, it is possible for ENodes with many FDISC sessions to be placed on the same link, whereas ENodes with few FDISC sessions are placed on different links because only FLOGIs are used in the session count.

### Load Balancing on the FCoE Interfaces (Ethernet Links)

You can achieve load balancing on the 10-Gigabit Ethernet interfaces that belong to an FCoE VLAN by configuring them in one or more link aggregation groups (LAGs). In addition to load balancing, LAGs offer the advantages of protecting against port failover, increasing available link bandwidth, and preventing spanning-tree algorithms from blocking physical links and wasting bandwidth.

#### Related Documentation

- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Defining the Proxy Load-Balancing Algorithm on page 5190](#)
- [Simulating On-Demand Fibre Channel Link Load Rebalancing \(Dry Run Test\) on page 5191](#)
- [Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175](#)
- [show fibre-channel proxy fabric-state on page 5364](#)
- [request fibre-channel proxy load-rebalance on page 5286](#)
- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

## Understanding OxID Hash Control for FCoE Traffic Load Balancing

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links in an Ethernet link aggregation group (LAG). The originator of an exchange between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) uses the OxID field as an identifier for that exchange. The originator also uses the OxID field to track the progress of the series of sequences that comprise the exchange.

When FCoE traffic traverses a LAG, it can take multiple different links between the source and destination endpoints. The idea is to distribute the FCoE traffic across the LAG links, thus balancing the link load. The switch creates a hash value from some of the packet header fields, and uses the hash value to assign each packet to one of the LAG links. The switch always uses five packet header fields to compute the hash value:

- Source ID (SID)
- Destination ID (DID)
- Fabric ID (FID)
- Source Port ID (SPID)
- Source Module ID (SMID)

In addition, the OxID field is included by default in the FCoE load-balancing hash computation. However, if you do not want to use the OxID field in the FCoE load-balancing hash computation, you can remove it from the computation by using the **set forwarding-options hash-key family fcoe oxid disable** command.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, on the QFX3500 switch, you can assign different IEEE priorities to different lossless FCoE flows as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) to further separate the traffic flows.

### Related Documentation

- [Enabling and Disabling CoS OxID Hash Control on page 5192](#)

## Understanding VN\_Port to VF\_Port FIP Snooping on an FCoE Transit Switch

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping is a security mechanism that is designed to prevent unauthorized access and data transmission to a Fibre Channel (FC) network. It works by filtering traffic to permit only servers that have logged in to an FC network to access that network.

You explicitly enable VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (FC-BB-5) on FCoE VLANs when the QFX Series is an FCoE transit switch connecting FCoE devices on the Ethernet network to FC switches or gateways at the FC storage area network (SAN) edge. The transit switch applies FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VF\_Port FIP snooping. An FCoE transit switch is a data center bridging (DCB) switch with FIP snooping capability.

Through the FIP process, FCoE devices that have a converged network adapter (CNA) present an FCoE Node (ENode) that can log in to the FC network. The login process establishes a dedicated virtual link between a virtual N\_Port (VN\_Port) on the ENode and a virtual F\_Port (VF\_Port) on the FC switch to emulate a point-to-point connection. The emulated connection is called a virtual link.

Virtual links pass transparently through the switch. The ENode VN\_Port and the FC switch VF\_Port do not detect the transit switch, and virtual links appear to be direct point-to-point links.

The QFabric system applies VN2VF\_Port FIP snooping firewall filters at the directly attached edge ports associated with the FCoE VLANs on which you enable VN2VF\_Port FIP snooping. FIP snooping provides security for virtual links by creating firewall filters based on information gathered (snooped) about FC devices during FIP transactions.

The QFX3500 switch also supports VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping (FC-BB-6) to allow FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch, as described in [“Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch” on page 5031](#).



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping (FC-BB-5) or VN2VN\_Port FIP snooping (FC-BB-6), but not both. The same QFX3500 switch can have multiple FCoE VLANs configured, some for VN2VF\_Port FIP snooping traffic and others for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port snooping VLANs, VN2VF\_Port FIP snooping traffic is dropped.

When you enable VN2VF\_Port FIP snooping, the system snoops VN\_Port to VF\_Port packets and enforces security only on VN2VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port packets and enforces security only on VN2VN\_Port virtual links.

- [FC Network Security on page 5026](#)
- [VN2VF\\_Port FIP Snooping Functions on page 5026](#)

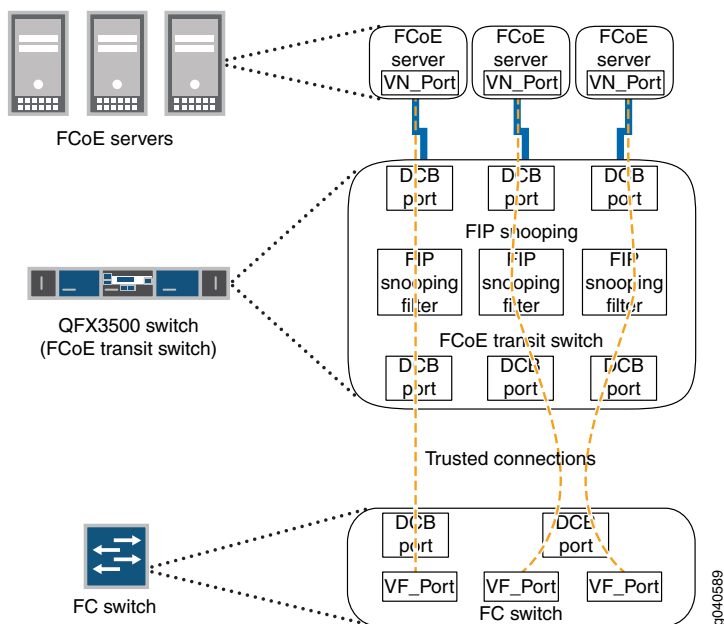
- [FIP Snooping Firewall Filters on page 5027](#)
- [Session Scalability on page 5027](#)
- [VN2VF\\_Port FIP Snooping Implementation on page 5027](#)
- [T11 VN2VF\\_Port FIP Snooping Specification on page 5030](#)

## FC Network Security

In traditional FC networks, the FC switch is usually a trusted entity, and server ENodes connect directly to its VF\_Ports. After an ENode gains access to the network through the fabric login (FLOGI) process, the FC switch enforces zoning configurations, ensures that the ENode uses valid addresses, monitors the connection, and performs other security functions to prevent unauthorized access.

However, FCoE exposes FC frames to Ethernet networks, which do not have the same level of security as native FC networks. VN2VF\_Port FIP snooping firewall filters emulate the native FC network security functions by preventing unauthorized access to the FC switch through the transit switch and by ensuring the security of the virtual link between each ENode and the FC switch, as shown in [Figure 192 on page 5026](#). VN2VF\_Port FIP snooping also prevents man-in-the-middle attacks.

**Figure 192: FCoE Transit Switch Performs VN2VF\_Port FIP Snooping**



The transit switch performs VN2VF\_Port FIP snooping at the ports connected to the FCoE devices. At the SAN edge, the FC switch must be able to convert the FCoE traffic to native FC traffic.

## VN2VF\_Port FIP Snooping Functions

When VN2VF\_Port FIP snooping is enabled, the QFX3500 transit switch sets and applies filters to block all FCoE traffic by default. The transit switch monitors FIP logins, solicitations, and advertisements that pass through it and gathers information about the

ENode address and the address of the port on the FC switch. The transit switch uses the information to construct firewall filters that permit access only to logged-in ENodes. All other traffic on the VLAN is denied.

For example, when an ENode on an FCoE VLAN performs a successful login to an FC switch port, the transit switch snoops the FIP information and constructs a firewall filter that provides access for the ENode to that port on the FC switch.

The firewall filters enable FCoE frames to pass through the transit switch only on a virtual link established between an FCoE device ENode VN\_Port and the FC switch VF\_Port to which it has logged in. The firewall filters ensure that ENodes can only connect to the FC switches they have successfully logged in to and that only valid FCoE traffic along valid paths is transmitted. VN2VF\_Port FIP snooping maintains the filters by tracking FCoE sessions (ENode to FCF sessions).

### FIP Snooping Firewall Filters

The effect of the firewall filters is to protect the FCoE ports. VN2VF\_Port FIP snooping performs the following actions and checks to ensure that FCoE traffic is valid:

- Denies ENodes that use the FC switch media access control (MAC) address as the source address.
- Enables ENodes to transmit FIP and FCoE frames to the FC switch address.
- Ensures that the FCoE source address the FC switch assigns or accepts is only used for FCoE traffic.
- Ensures that FCoE frames are only addressed to the accepting FC switch.

### Session Scalability

The system supports up to 2500 total FIP snooping sessions on an interface, a gateway FC fabric, a QFX3500 switch, a QFabric Node device, or a QFabric Node group. For example, you can:

- Place all 2500 sessions on one gateway FC fabric interface.
- Split the 2500 sessions among the interfaces on one gateway FC fabric.
- Split the 2500 sessions among the interfaces on multiple gateway FC fabrics on a switch.
- Split the 2500 sessions among the interfaces on multiple gateway FC fabrics on multiple Node devices in a QFabric Node group.

Regardless of how you allocate the sessions among interfaces and local FC fabrics on a QFX3500 switch or on a QFabric system Node device or Node group, the combined FIP session limit is a maximum of 2500 sessions.

### VN2VF\_Port FIP Snooping Implementation

You enable VN2VF\_Port FIP snooping on a per-VLAN basis on VLANs that carry FCoE traffic. The switch snoops FIP frames at the ports associated with FCoE VLANs enabled

for VN2VF\_Port FIP snooping. The switch then installs the resulting firewall filters on the ports to ensure that all VN2VF\_Port FIP snooping occurs on the switch network edge.

VN2VF\_Port FIP snooping FCoE VLANs must meet the following criteria:

- An FCoE VLAN should be dedicated to FCoE traffic only.
- An FCoE VLAN cannot support both VN2VF\_Port FIP snooping and VN2VN\_Port FIP snooping simultaneously. You must configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic.



**NOTE:** Changing an FCoE VLAN from VN2VF\_Port FIP snooping mode to VN2VN\_Port snooping mode terminates the existing virtual links on the VLAN. The transit switch removes the existing FIP snooping filters, creates the new FIP snooping filters, and applies them to the FIP snooping ports. If you downgrade the software to Junos OS Release 12.1 or earlier, VLANs configured for VN2VN\_Port FIP snooping revert to VN2VF\_Port FIP snooping VLANs.

- All access ports associated with an FCoE VLAN (ports connected to a converged network adapter [CNA] in an FCoE device) should be configured in **tagged-access** port mode. Access ports associated with an FCoE VLAN should not be configured as access ports or trunk ports.
- All ports connected to an FC switch (or FCoE forwarder) must be configured in **trunk** port mode. Ports connected to an FC switch must be configured as trusted ports.
- FIP traffic uses the native VLAN (FIP VLAN discovery and notification frames are exchanged as untagged packets).
- All FCoE VLAN traffic must be tagged and cannot belong to the native VLAN.
- FCoE VLAN traffic cannot be untagged or priority-tagged.

When you enable VN2VF\_Port FIP snooping, the switch inspects FIP frames.

The VN2VF\_Port FIP snooping implementation includes these considerations:

- [ENode-Facing Interfaces on page 5028](#)
- [Network-Facing Interfaces on page 5029](#)
- [FC-MAP on page 5029](#)

### ***ENode-Facing Interfaces***

We recommend that you enable VN2VF\_Port FIP snooping on all FCoE VLANs that connect VN\_Ports to VF\_Ports to ensure secure connections between server ENodes and FC switches (or enable VN2VN\_Port FIP snooping on FCoE VLANs that connect VN\_Ports to other VN\_Ports). The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) should be configured in **tagged-access** port mode. After VN2VF\_Port FIP snooping is enabled on an FCoE VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login with an FC switch.



The **tagged-access** port mode was not available in Junos OS Release 11.3 and prior releases. In Release 11.3 and earlier, **trunk** port mode was used for Ethernet interfaces that connected to FCoE access devices. Because **tagged-access** mode is now available, using **trunk** mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses **trunk** mode for ports connected to FCoE CNAs, you can change the port mode to **tagged-access** without disrupting traffic. Although we recommend changing the port mode of these ports from **trunk** to **tagged-access** as a best practice, it is not mandatory. New configurations should use **tagged-access** mode for interfaces that connect to FCoE devices.



**NOTE:** Do not configure ENode-facing interfaces as FCoE trusted interfaces when VN2VF\_Port FIP snooping is enabled on those interfaces. If you enable VN2VF\_Port FIP snooping on an FCoE VLAN and you configure ENode-facing interfaces that are members of the FIP snooping VLAN as **fcoe-trusted**, then FCoE devices might not be able to log in to the FC network.



**NOTE:** Changing ports from untrusted to trusted removes any existing VN2VF\_Port FIP snooping filters from the ports and terminates the existing sessions. Changing the fabric ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate VN2VF\_Port FIP snooping filters.

### **Network-Facing Interfaces**

When the switch acts as an FCoE transit switch, you must configure any interface that is connected to a switch as an FCoE trusted interface in **trunk** port mode and as a 10-Gigabit Ethernet interface.

Switch-facing Ethernet interfaces have the following requirements and behaviors:

- You must explicitly configure switch-facing trunk ports on an FCoE transit switch as FCoE trusted interfaces.
- After you configure an FC switch-facing trunk port as a trusted interface, the FCoE transit switch always processes FC switch frames because they come from a source on a trusted interface.
- All ports in an FCoE VLAN must be configured as tagged access or trunk ports.

### **FC-MAP**

When the switch acts as an FCoE transit switch and you enable VN2VF\_Port FIP snooping on an FCoE VLAN, you can optionally specify a 24-bit FCoE mapped address prefix (FC-MAP) value. On a given VLAN, the transit switch learns only those FC switches that have a matching FC-MAP value. If the transit switch FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, the transit switch does not discover the FC switch

on that VLAN, and the ENodes on that VLAN cannot access the FC switch. An FCoE VLAN can have one and only one FC-MAP value.

The FC-MAP value is a MAC address prefix unique to an FC switch in the FC SAN fabric that the FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). The FC switch combines the FC-MAP value with a unique 24-bit FCID value for the ENode VN\_Port during the login process. This creates a 48-bit identifier that is unique to the fabric. The FC switch assigns this 48-bit value to the ENode VN\_Port as its MAC address and unique identifier for the session. Each VN\_Port session the ENode establishes with the FC switch receives a unique FCID from the FC switch, so an FCoE device can host multiple virtual links (one for each VN\_Port) to an FC switch, each with a 48-bit MAC address that is unique to the fabric.

The VN2VF\_Port FIP snooping filter compares the configured FC-MAP value with the FC-MAP value in the header of frames coming from the ENode VN\_Port. If the values do not match, the transit switch denies access.



**NOTE:** Changing the FC-MAP value causes all logins to be dropped and forces ENodes to log in again.

---



**NOTE:** Do not configure static MAC addresses with the FC-MAP value as a prefix (the first 24 bits of the MAC address). If you configure a static MAC address that uses the FC-MAP value as a prefix, the system deletes the static MAC address automatically after you enable FIP snooping. The static MAC address configuration is not restored even if you disable FIP snooping later. (The system considers a static MAC address with the FC-MAP value as the prefix to be a misconfiguration.) Do not use a MAC address with the FC-MAP value as the prefix for any traffic other than the FIP snooping traffic when the QFX Series is acting as a transit switch.

---

### T11 VN2VF\_Port FIP Snooping Specification

For more details about VN2VF\_Port FIP snooping, see <http://www.t11.org/ftp/t11/pub/fc/bb-5/08-264v3.pdf> for the Technical Committee T11 organization document *Increasing FCoE Robustness using FIP Snooping*.

#### **Related Documentation**

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding FCoE Transit Switch Functionality on page 4972](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Overview of FIP on page 4960](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Configuring VN2VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)

- [Disabling VN2VF\\_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)
- [Understanding Fibre Channel Terminology on page 5074](#)

## Understanding VN\_Port to VN\_Port FIP Snooping on an FCoE Transit Switch

VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping (FC-BB-6) on an FCoE transit switch is conceptually similar to VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (FC-BB-5) on an FCoE transit switch. VN2VN\_Port FIP snooping provides security in the form of filters. The filters help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network.

You enable VN2VN\_Port FIP snooping on the FCoE VLAN that transports the VN2VN traffic. The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

A key benefit of VN2VN\_Port FIP snooping is that it enables FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch. The transit switch does not differentiate between initiators and targets because the transit switch sees both VN\_Ports as FIP virtual link end points. Direct VN2VN\_Port communication requires secure access (FIP snooping filters) because ENodes are not trusted entities.

This topic describes:

- [VN2VN\\_Port FIP Snooping and FIP Snooping Virtual Links on page 5031](#)
- [VN2VN\\_Port Communication Modes on page 5032](#)
- [Network Security on page 5033](#)
- [VN2VN\\_Port FIP Snooping Functions on page 5033](#)
- [Scalability on page 5033](#)
- [VN2VN\\_Port FIP Snooping Implementation on page 5033](#)
- [ENode-Facing Interfaces on page 5034](#)
- [Network-Facing Interfaces \(Connecting to Another Transit Switch\) on page 5035](#)
- [Beacon Period \(VN2VN\\_Port FIP Snooping Link Maintenance\) on page 5035](#)
- [QFabric System Differences in VN2VN\\_Port FIP Snooping Traffic Handling on page 5036](#)

### VN2VN\_Port FIP Snooping and FIP Snooping Virtual Links

FIP snooping under the T11 FC-BB-5 specification requires that an FC switch or an FCF be in the path between two VN\_Ports when they communicate. Introduced in the T11 FC-BB-6 specification (see <http://www.t11.org/ftp/t11/pub/fc/bb-6/10-019v3.pdf>), VN2VN\_Port FIP snooping allows the FCoE transit switch to connect two VN\_Ports to each other directly, without going through an FC switch or an FCF, provided that the ENodes have logged in to the FC network.

In VN2VF\_Port FIP snooping, when an ENode logs in to the FC network, the FCoE transit switch snoops the FIP communication between the ENode and the FC switch. In

VN2VN\_Port FIP snooping mode, the transit switch creates filters on the switch access ports to control VN\_Port access to other VN\_Ports on the Ethernet network. The VN2VN\_Port FIP snooping filters allow the switch to establish a dedicated virtual link that emulates a point-to-point connection between two VN\_Ports, through the switch.

Virtual links pass transparently through the transit switch. The VN\_Ports do not detect the transit switch, and virtual links appear to be direct point-to-point links.

You explicitly enable VN2VN\_Port FIP snooping on FCoE VLANs when the QFX Series switch or QFabric system is an FCoE transit switch connecting FCoE devices on the Ethernet network to each other and to FC switches or gateways at the FC storage area network (SAN) edge. An FCoE transit switch is a data center bridging (DCB) switch with FIP snooping capability.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN\_Port to VF\_Port traffic is dropped.

When you enable FIP snooping, the system snoops VN2VF\_Port packets and enforces security only on VN\_Port to VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port FIP packets and enforces security only on VN\_Port to VN\_Port virtual links.

---

The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN\_Port FIP snooping. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

### VN2VN\_Port Communication Modes

The transit switch supports two VN2VN\_Port communication modes:

- Point-to-point mode
- Multipoint mode

In point-to-point mode, two ENodes are connected to the network and form a single VN\_Port to VN\_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.

In multipoint mode, multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN\_Ports. This is analogous to loop mode in traditional FC networks.

The VN2VN\_Port communication mode is not configured; it is determined by the number of ENodes connected to the network.

---

## Network Security

---

In traditional FC networks, the FC switch is usually a trusted entity and the server ENodes are untrusted entities. The ENodes connect directly to the FC switch VF\_Ports. After an ENode gains access to the network through the fabric login (FLOGI) process, the FC switch enforces zoning configurations, ensures that the ENode uses valid addresses, monitors the connection, and performs other security functions to prevent unauthorized access.

However, FCoE exposes FC frames to Ethernet networks, which do not have the same level of security as native FC networks. VN2VN\_Port FIP snooping filters emulate the native FC network security functions by preventing unauthorized access and by ensuring the security of the virtual link between ENode VN\_Ports. The transit switch performs VN2VN\_Port FIP snooping at the ports connected to the FCoE VN\_Port devices.

---

### VN2VN\_Port FIP Snooping Functions

---

When you enable VN2VN\_Port FIP snooping, the transit switch sets and applies filters to block all FCoE traffic on the VLAN by default. The transit switch monitors FIP logins, solicitations, and advertisements that pass through it and gathers information about the ENode address. The transit switch uses the information to construct filters that permit access only to logged-in ENodes. All other traffic on the VLAN is denied.

The filters enable FCoE frames to pass through the transit switch only on a virtual link established between two VN\_Ports. The filters ensure that ENodes can only connect to other ENodes if they have successfully logged in to each other, and that only valid FCoE traffic along valid paths is transmitted. VN2VN\_Port FIP snooping maintains the filters by tracking VN\_Port to VN\_Port sessions.

---

### Scalability

---

Because ENodes are untrusted and the system needs to apply filters to untrusted FIP snooping interfaces, the total number of combined VN2VN\_Port FIP snooping sessions per switch is 376 sessions (ENode to ENode sessions) on untrusted interfaces. On interfaces that are configured as trusted interfaces, no FIP snooping filters are applied.



**NOTE:** The total number of sessions the system can support is the combined number of VN2VF\_Port sessions and VN2VN\_Port sessions. If VN2VF\_Port sessions are active, the total number of available VN2VN\_Port sessions is reduced.

---

---

### VN2VN\_Port FIP Snooping Implementation

---

You enable VN2VN\_Port FIP snooping on a per-VLAN basis on VLANs that carry FCoE traffic. The switch snoops FIP frames at the ports associated with FCoE VLANs enabled for VN2VN\_Port FIP snooping. The switch then installs the resulting filters on the ENode-facing ports to ensure that all FIP snooping occurs on the switch network edge.

VN2VN\_Port FIP snooping FCoE VLANs must meet the following criteria:

- An FCoE VLAN should be dedicated to FCoE traffic only.
- An FCoE VLAN cannot support both VN2VF\_Port FIP snooping (FC-BB-5) and VN2VN\_Port FIP snooping (FC-BB-6) simultaneously. You must configure separate FCoE VLANs for FIP snooping traffic and for VN2VN\_Port FIP snooping traffic.



**NOTE:** Changing an FCoE VLAN from VN2VF\_Port FIP snooping mode to VN2VN\_Port FIP snooping mode terminates the existing virtual links on the VLAN. The transit switch removes the existing FIP snooping filters, creates the new FIP snooping filters, and applies them to the FIP snooping ports. If you downgrade the software to Junos OS Release 12.1 or earlier, VLANs configured for VN2VN\_Port FIP snooping revert to VN2VF\_Port FIP snooping VLANs.

- As a best practice, all access ports associated with an FCoE VLAN (ports connected to a converged network adapter [CNA] in an FCoE device) should be configured in **tagged-access** port mode, but access and trunk port modes are also supported.
- Access ports should be configured as untrusted ports.
- All ports connected to another transit switch must be configured in **trunk** port mode.
- FIP traffic uses the native VLAN.
- You can enable VN2VN\_Port FIP snooping on a native VLAN.

---

### ENode-Facing Interfaces

We recommend that you either enable VN2VN\_Port FIP snooping on all FCoE VLANs to ensure secure connections between VN\_Ports, or enable VN2VF\_Port FIP snooping on FCoE VLANs that connect ENodes to an FC switch. The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) should be configured in **tagged-access** port mode, unless your CNA does not support tagged VN2VN traffic. After you enable VN2VN\_Port FIP snooping on an FCoE VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login (FIP FLOGI) with another ENode.

The **tagged-access** port mode was not available in Junos OS Release 11.3 and prior releases. In Release 11.3 and earlier, **trunk** port mode was used for Ethernet interfaces that connected to FCoE access devices. Because **tagged-access** mode is now available, using **trunk** mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses **trunk** mode for ports connected to FCoE CNAs, you can change the port mode to **tagged-access** without disrupting traffic. Although we recommend changing the port mode of these ports from **trunk** to **tagged-access** as a best practice, it is not mandatory. New configurations should use **tagged-access** mode for interfaces that connect to FCoE devices.



**NOTE:** Do not configure ENode-facing interfaces as FCoE trusted interfaces when VN2VN\_Port FIP snooping is enabled on those interfaces. If you enable VN2VN\_Port FIP snooping on an FCoE VLAN and you configure ENode-facing interfaces that are members of the VN2VN\_Port FIP snooping VLAN as FCoE trusted, then FCoE devices might not be able to log in to each other.



**NOTE:** Changing ports from untrusted to trusted removes any existing VN2VN\_Port FIP snooping filters from the ports and terminates the existing sessions. Changing the ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate VN2VN\_Port FIP snooping filters.

### Network-Facing Interfaces (Connecting to Another Transit Switch)

Configure any interface that is connected to another transit switch (not to an ENode) as an FCoE trusted interface, in **trunk** port mode, and as a 10-Gigabit Ethernet interface.

Network-facing Ethernet interfaces have the following requirements and behaviors:

- You must explicitly configure network-facing trunk ports on an FCoE transit switch as FCoE trusted interfaces.
- After you configure a network-facing trunk port as a trusted interface, the FCoE transit switch always processes frames from the connected switch because they come from a source on a trusted interface.
- As a best practice, configure ports in an FCoE VLAN as tagged access ports, but access and trunk port modes are also supported to accommodate whatever types of VN2VN traffic your CNA supports.

### Beacon Period (VN2VN\_Port FIP Snooping Link Maintenance)

The transit switch needs to maintain the virtual links between VN\_Ports, and needs to know when sessions begin and end, and when to install and remove the FIP snooping filters. FIP snooping uses a FIP keepalive advertisement to accomplish this task. VN2VN\_Port FIP snooping does not exchange FIP keepalive timer information. Instead, you configure a *beacon period*, which performs the same function as a keepalive timer.

The beacon period is the time interval between messages which verify that the connection is still valid and that the device at the other end of the virtual link is still reachable. You set the beacon period value for each FCoE VLAN that you configure to do VN2VN\_Port FIP snooping.



**NOTE:** Explicitly set the beacon period when you configure VN2VN\_Port FIP snooping. VN\_Ports do not automatically send beacons.

ENodes transmit periodic multicast N\_Port\_ID beacons to the ALL-VN2VN-ENode-MACs address. The transmission period varies by a random delay of between 0 ms and 100 ms to avoid synchronized bursts of multicast traffic on the network.

If the transit switch does not receive a beacon message from an ENode within 2.5 times the configured beacon period, the transit switch considers the virtual link to be down and terminates the virtual link to that ENode.

### QFabric System Differences in VN2VN\_Port FIP Snooping Traffic Handling


Configuring VN2VN\_Port FIP snooping on a QFabric system is the same as configuring VN2VN\_Port FIP snooping on the QFX Series. However, there are internal differences in the way a QFabric system handles VN2VN\_Port FIP snooping traffic compared to the way a QFX Series switch handles VN2VN\_Port FIP snooping traffic. The internal differences are transparent. Whether you configure VN2VN\_Port FIP snooping on a QFabric system or on the QFX Series, the proper FIP snooping filters and forwarding information are installed on each device.

On the QFX Series, the VN2VN\_Port FIP snooping traffic does not cross a fabric (Interconnect device). VN2VN\_Port traffic enters and exits ports on a single switch, so the ingress port and the egress port have access to the same *local* forwarding and FIP snooping databases.

However, on a QFabric system, VN2VN\_Port FIP snooping traffic might enter on the ingress port of one Node device, traverse the Interconnect device fabric, and exit on the egress port of a different Node device. In this case, the QFabric system must ensure that the FIP snooping database and forwarding information for the VN2VN\_Port traffic is installed correctly on both of the Node devices so that traffic is correctly filtered and forwarded.

For example, [Figure 193 on page 5036](#) shows that VN2VN\_Port traffic from FCoE host ENode E1 enters the QFabric system at Node device ND1, traverses the Interconnect device fabric, and then exits from Node device ND2 before arriving at FCoE host ENode E2. Similarly, VN2VN\_Port traffic from FCoE host ENode E2 enters the QFabric system at Node device ND2, traverses the Interconnect device fabric, and then exits from Node device ND1 before arriving at FCoE host ENode E1.

Figure 193: VN2VN\_Port Traffic Across a QFabric Interconnect Device





When the QFabric system receives a FLOGI ACC from either ENode E1 or ENode E2, the QFabric system creates and installs the correct VN2VN\_Port FIP snooping filters on both Node devices, and updates the forwarding tables accordingly.

In addition, the QFabric system must also ensure that the VN2VN\_Port FIP snooping session statistics are correctly counted. Even though a session is running on each of the two Node devices, the QFabric system counts the complete VN2VN\_Port connection as one session because the two Node devices belong to the same session. This ensures that VN2VN\_Port sessions that traverse the Interconnect device fabric are counted as one unique session, not as two separate sessions.

**Related  
Documentation**

- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding FCoE Transit Switch Functionality on page 4972](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Overview of FIP on page 4960](#)
- [Understanding Fibre Channel Terminology on page 5074](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5155](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5167](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5202](#)

## Understanding FIP Snooping, FBF, and MVR Filter Scalability

The VLAN filter processor (VFP) ternary content addressable memory (TCAM) stores the VLAN filter configuration for three filter types:

- Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping—FIP snooping filters prevent an FCoE device from gaining unauthorized access to a Fibre Channel (FC) storage device or to another FCoE device. VN2VF\_Port FIP snooping filters prevent an FCoE device from gaining unauthorized access to devices on an FC network. VN2VN\_Port FIP snooping filters prevent an FCoE device from gaining unauthorized access to another FCoE device directly through the QFX Series or QFabric system, without traversing the FC network.

The VFP TCAM stores the VN2VF\_Port and VN2VN\_Port FIP snooping filters that the switch automatically creates when you enable FIP snooping on a VLAN that carries FCoE traffic. See [“Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch” on page 5025](#) and [“Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch” on page 5031](#) for more information.

- Filter-based forwarding (FBF)—FBF enables you to use firewall filters to direct packets to virtual routing instances. The switch then forwards the matching packets based on the configuration of the routing instances. The VFP TCAM stores the terms you configure for FBF filters. See [“Understanding Filter-Based Forwarding” on page 4637](#) for more information.
- Multicast VLAN registration (MVR)—MVR enables you to configure a multicast source VLAN (MVLAN) that is shared across a Layer 2 network. An MVLAN distributes IPTV multicast streams across different VLANs without having to create a separate multicast stream for each VLAN, and without compromising the security and separation of traffic in the different VLANs. The VFP TCAM stores the MVR rules you configure for MVLANS. See [“Understanding Multicast VLAN Registration” on page 4241](#) for more information.

FIP snooping filters, FBF filters, and MVR rules share the VFP TCAM memory space. In most use cases, the VFP TCAM memory is sufficient to store filter terms and information for all three applications.

- [VFP TCAM Architecture and Allocation on page 5038](#)
- [VFP TCAM Entry Consumption on page 5039](#)
- [Rejected Filter Configurations \(No Available VFP TCAM Space\) on page 5042](#)
- [VFP TCAM Allocation and Consumption \(Scaling\) Examples on page 5043](#)
- [Filter Configuration Recommendations on page 5045](#)

---

### VFP TCAM Architecture and Allocation

When packets arrive at an ingress interface, the VFP TCAM is the first TCAM in the packet pipeline. The VFP TCAM stores a total of 1024 entries. The 1024 entries are partitioned into four equal *slices* of 256 entries.

The VFP TCAM allocates entries to three filter types (FIP snooping filters, FBF filter terms, and MVR rules) in 256-entry slices. The VFP TCAM dynamically allocates the minimum number of memory slices required to store the filters for a particular filter type, as needed.

The TCAM does not allocate partial slices to a filter type, and slices cannot be shared among filter types. At any given time, each slice contains entries for one and only one filter type.

For example, if you configure one MVR rule, the system allocates a whole slice to MVR rules, even if the MVR rule consumes only one TCAM entry. The remaining 256 entries in the slice allocated to MVR rules can store subsequently configured MVR rules, but not FIP snooping or FBF filters. Similarly, if FIP snooping filters consume 50 entries of a 256-entry slice, the remaining 206 entries in the FIP snooping slice are available only to store more FIP snooping filters, not to store FBF filter terms or MVR rules.

The VFP TCAM allocates slices to a filter type only if there is at least one configured filter or rule for that filter type. If no filters exist for a filter type, then the VFP TCAM does not allocate a slice to that filter type.



**NOTE:** The VFP TCAM rejects partial filters. For example, if an FBF filter contains six terms, but there is only space in the TCAM for four of those terms, the whole filter is not committed.

Each filter type can use from zero slices to all four slices of VFP TCAM space. However, if one filter type uses three slices, then only one slice remains, so only one other filter type can use the remaining slice. In that situation, if you configure filters for all three filter types, the last filter type that you configure receives no TCAM space for its filter entries. Filters that receive no TCAM entry space are not implemented.

### VFP TCAM Entry Consumption

FIP snooping filters, FBF filters, and MVR rules consume VFP TCAM entry space in different ways:

- [FIP Snooping Filter VFP TCAM Consumption on page 5039](#)
- [FBF Filter VFP TCAM Consumption on page 5040](#)
- [MVR Filter VFP TCAM Consumption on page 5041](#)
- [VFP TCAM Consumption Summary Table on page 5041](#)

#### *FIP Snooping Filter VFP TCAM Consumption*

VN2VF\_Port FIP snooping filters consume VFP TCAM entry space differently than VN2VN\_Port FIP snooping filters:

- [VN2VF\\_Port FIP Snooping Filter VFP TCAM Consumption on page 5040](#)
- [VN2VN\\_Port FIP Snooping Filter VFP TCAM Consumption on page 5040](#)



**NOTE:** One FCoE VLAN cannot support both VN2VF\_Port traffic and VN2VN\_Port traffic. Configure separate FCoE VLANs for VN2VF\_Port traffic and for VN2VN\_Port traffic.

### ***VN2VF\_Port FIP Snooping Filter VFP TCAM Consumption***

The switch uses an algorithm that allows one 256-entry slice of the VFP TCAM to store the maximum possible number of VN2VF\_Port FIP snooping filters (2500 filters). VN2VF\_Port FIP snooping filters never consume more than one slice of the VFP TCAM.

Regardless of whether there is one VN2VF\_Port FIP snooping session or there are 2500 VN2VF\_Port FIP snooping sessions, VN2VF\_Port FIP snooping filters consume one slice of the VFP TCAM. (If there are no VN2VF\_Port or VN2VN\_Port FIP snooping sessions, the TCAM does not allocate a slice for FIP snooping filters.)

### ***VN2VN\_Port FIP Snooping Filter VFP TCAM Consumption***

VN2VN\_Port FIP snooping filters consume one VFP TCAM entry for each VN2VN\_Port session. The maximum number of VN2VN\_Port FIP snooping sessions is 376 sessions per switch. (If you configure an interface that carries VN2VN\_Port FIP snooping traffic as a trusted interface, the switch does not apply filters on the trusted interface.)

Because the switch can have up to 376 VN2VN\_Port sessions running simultaneously, with each session consuming one entry, VN2VN\_Port FIP snooping filters consume VFP TCAM space as follows:

- 1–256 filters consume one slice
- 257–376 filters consume two slices

### ***FBF Filter VFP TCAM Consumption***

Each FBF filter term is double-wide, so each FBF filter term consumes two entries in the VFP TCAM. One 256-entry slice can contain up to 128 FBF filter terms. FBF filters consume VFP TCAM space as follows:

- 1–128 entries consume one slice
- 129–256 entries consume two slices
- 257–384 entries consume three slices
- 385–512 entries consume four slices



**NOTE:** In practice, FBF filters can consume only three slices of the VFP TCAM because FBF filters are also stored simultaneously in the ingress filter processor (IFP) TCAM, and the IFP TCAM can store only 384 FBF filter terms (768 entries, or 3 TCAM slices).

---

For example, if you configure FBF filters that contain 200 terms, then the FBF filters require 400 VFP TCAM entries and consume 2 slices.

FBF filter entries are simultaneously stored in the VFP TCAM and the IFP TCAM. The IFP TCAM can only contain up to 768 entries—256 fewer entries (1 slice) than the VFP TCAM. As with the VFP TCAM, FBF filters consume two IFP TCAM entries per filter term. In addition to FBF filter terms, the IFP TCAM stores filter entries for firewall filters.



**CAUTION:** There must be enough space in the VFP TCAM *and* the IFP TCAM for the FBF filter entries. If both TCAMs do not have enough space for the FBF filters, the switch rejects the portion of the configuration that it cannot store and sends a syslog message to notify you.

For example, if you configure FBF filters that have 400 terms, even though the VFP TCAM has enough space to store the resulting 800 entries, the switch rejects a portion of the configuration because the IFP TCAM can store a maximum of only 768 entries. If the IFP TCAM stores no other filter entries, the switch rejects 32 FBF filter entries.

In another example, if you configure firewall filters that have a total of 200 terms, which consume 200 entries in the IFP TCAM, and you then configure FBF filters that have a total of 300 terms, the switch rejects a portion of the configuration because the FBF filters require 600 entries. Combined with the 200 entries required for the firewall filters, the total number of 800 entries exceeds the maximum of 768 entries that the IFP TCAM can store. In this case, the switch accepts the first 768 entries and rejects the rest of the filter entries. The switch installs the filter entries in the order that they are committed; the rejected entries are the last entries the switch attempts to commit after the TCAM space is exhausted.

The IFP TCAM limit of 768 entries means that the true maximum number of FBF filter terms is 384 terms, even though the VFP TCAM can store up to 512 FBF terms.

#### ***MVR Filter VFP TCAM Consumption***

Each MVR rule consumes one entry in the VFP TCAM, so MVR rules consume VFP TCAM space as follows:

- 1–256 rules consume one slice
- 257–512 rules consume two slices
- 513–758 rules consume three slices
- 759–1024 rules consume four slices

#### ***VFP TCAM Consumption Summary Table***

Table 391 on page 5042 summarizes VFP TCAM consumption.



**NOTE:** FBF filters are simultaneously stored in the VFP TCAM and in the IFP TCAM. Due to the IFP TCAM limit of 768 entries (384 FBF filters), which is 256 entries fewer than the VFP TCAM, the effective VFP TCAM consumption limit for FBF filters is lower than the total amount of VFP TCAM entry space, even when no other filters consume VFP TCAM space.

Table 391: VFP TCAM Entry Consumption Summary

| Filter Type                     | VFP TCAM Entry Consumption         | Maximum VFP TCAM Slices Consumed             | Other Limitations                        |
|---------------------------------|------------------------------------|----------------------------------------------|------------------------------------------|
| VN2VF_Port FIP snooping filters | Never consumes more than one slice | One slice (regardless of number of sessions) | 2500 session maximum                     |
| VN2VN_Port FIP snooping filters | One entry per session              | Two                                          | 376 session maximum                      |
| FBF filters                     | Two entries per filter             | Three (due to IFP TCAM limitation)           | 384 filters (due to IFP TCAM limitation) |
| MVR rules                       | One entry per rule                 | Four                                         | 1024 rule maximum                        |

#### Rejected Filter Configurations (No Available VFP TCAM Space)

If there is not enough space available in the VFP TCAM to store the FIP snooping filters, the configured FBF filters, and the MVR rules, the switch rejects only the portion of the configuration that it cannot store. Any portion of the filter configuration that the TCAM can store, is stored. In most cases, even if the switch rejects part of the configuration, part of the configuration is also stored.

If the switch rejects any portion of a configuration, the switch sends a syslog message to notify you of the failure. The switch does not generate a commit error, and the rejected portion of the configuration remains on the switch, even though the rejected configuration does not function. (The accepted portions of the configuration function as expected.) The syslog message shows you the filter configuration that the switch rejected.

We strongly recommend that you always delete rejected filter configurations from the switch. It is important to delete rejected filter configurations because:

- Even though the rejected configuration remains on the switch, it does not function.
- After a reboot, there is no guarantee that the same filters will be rejected. The previously rejected filters might be accepted, and other filters that had previously been accepted might be rejected. Therefore, the functioning filter configuration could be changed inadvertently and unexpectedly.
- Even if a VFP TCAM slice becomes available, the switch does not automatically allocate the available slice to the rejected configuration. To use the available slice, you must delete and reconfigure the rejected configuration.

For example, you configure FBF filters and MVR rules on a switch, and that switch also transports FCoE traffic with VN2VF\_Port FIP snooping (never consumes more than one slice) enabled on FCoE access interfaces. After you commit the configuration, you check the syslog. You find that the VN2VF\_Port FIP snooping and FBF filters consume all four slices of the VFP TCAM, and the MVR configuration was rejected. Instead of deleting the MVR configuration, you leave it on the switch. Subsequently, all VN2VF\_Port FIP snooping sessions end, the FIP snooping filters time out and are removed from the VFP TCAM, so the slice that was allocated to VN2VF\_Port FIP snooping filters becomes free. However, the MVR rules do *not* automatically receive the free slice.

To force the switch to allocate the free slice to the MVR rules, you should delete the MVR rules from the configuration and then reconfigure the MVR rules. When you commit the new configuration, check the syslog messages to ensure that the MVR rule configuration was accepted.

In this example, you could also choose to free a VFP TCAM slice for MVR rule storage by deleting some of the FBF filters. To do this, you delete both the unneeded FBF filters and the MVR rule configuration. Then you reconfigure the MVR rules, and check the syslog to ensure that the configuration was successful.

### VFP TCAM Allocation and Consumption (Scaling) Examples

The following examples illustrate how FIP snooping entries, FBF filter entries, and MVR rule entries consume VFP TCAM slices:

- [Example 1: Three Filter Types Consume Three Slices on page 5043](#)
- [Example 2: Three Filter Types Consume Four Slices on page 5043](#)
- [Example 3: Two Filter Types Consume Four Slices on page 5044](#)
- [Example 4: Three Filter Types Oversubscribe the VFP TCAM on page 5044](#)

#### **Example 1: Three Filter Types Consume Three Slices**

Filters and rules are configured in the following sequence:

- 100 VN2VN\_Port FIP snooping filters (1 slice)
- 2 MVR rules (1 slice, 2 entries)
- 60 FBF filter terms (1 slice, 120 entries)

One slice remains free. The slice allocated to VN2VN\_Port FIP snooping filters can store 156 more filters before another slice is required. The slice allocated to MVR rules can store 254 more rules before another slice is required. The slice allocated to FBF filters can store 68 more filter terms (136 entries) before another slice is required. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.

#### **Example 2: Three Filter Types Consume Four Slices**

Filters and rules are configured in the following sequence:

- 2000 VN2VF\_Port FIP snooping filters (always 1 slice)
- 18 MVR rules (1 slice, 18 entries)
- 150 FBF filter terms (2 slices, 300 entries)

All four slices are allocated to filter types. The slice allocated to MVR rules can store 238 more rules before it is full. The slice allocated to FBF filters can store 106 more filter terms (212 entries) before it is full. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.



**NOTE:** If you configure more MVR rules or FBF filters than entry space remaining in the slices, the switch rejects those rules and filters because no slice is available. The switch installs filters in the order that they were configured, so if filters are rejected, the filters configured last are rejected.

---

### ***Example 3: Two Filter Types Consume Four Slices***

Filters and rules are configured in the following sequence:

- 50 VN2VF\_Port FIP snooping filters (always 1 slice)
- 300 FBF filter terms (3 slices, 600 entries)

All four slices are allocated to filter types. No slices are available for MVR rules. The third slice allocated to FBF filters can store 84 more filter terms (168 entries) before it consumes all of its entry space. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.



**NOTE:** If you configure MVR rules or if you configure more than 84 more FBF filters, the switch rejects those rules and filters because no slice is available for the MVR rules, and the FBF filter slice has entry space for only 84 more filter terms.

---

### ***Example 4: Three Filter Types Oversubscribe the VFP TCAM***

Filters and rules are configured in the following sequence:

- 1750 VN2VF\_Port FIP snooping filters (always 1 slice)
- 10 MVR rules (1 slice, 10 entries)
- 275 FBF filter terms (2 slices, 512 accepted entries, 38 rejected entries)

All four slices are allocated to filter types. The slice allocated to MVR rules can store 246 more rules before it is full, but the number of FBF filter terms exceeds the amount of available VFP TCAM storage space. (The 275 FBF filter terms consume 550 VFP TCAM entries. However, there are only two available slices, for a total of 512 available entry spaces, so only 256 FBF filter terms can be stored, leaving 19 rejected FBF filter terms.)

The switch accepts the VN2VF\_Port FIP snooping filters, the MVR rules, and 256 FBF filter terms. The switch retains the excess FBF filters in the configuration, but does not install those filters in the VFP TCAM. In this case, you delete the rejected FBF filter terms from the configuration. Alternatively, you could delete the MVR rules from the configuration to free a slice of the TCAM, and then delete and reconfigure the rejected FBF filters so that the system allocates the freed slice to the FBF filters.





**NOTE:** The sequence of configuration makes a difference; if there is not enough VFP TCAM space for a given filter type, the switch installs the filters that fit in the order they are configured. For example, if you configure the FBF filters before you configure the MVR rules, the VFP TCAM allocates one slice to FIP snooping filters, three slices to FBF filters (assuming the IFP TCAM has available space), and no slices to MVR rules, because all four slices are allocated before the switch attempts to install the MVR rules in the VFP TCAM.

### Filter Configuration Recommendations

To utilize the VFP TCAM space most efficiently:

- [Configure and Maintain the Fewest Number of Filters Needed on page 5045](#)
- [Always Delete Rejected Filter Configurations on page 5046](#)

#### ***Configure and Maintain the Fewest Number of Filters Needed***

To conserve VFP TCAM entry space, and because FBF filter storage also depends on the availability of IFP TCAM space, we recommend that you configure as few FBF filters and MVR rules as is practical to serve your network needs. The more filters you configure, the greater the possibility of exceeding TCAM storage capacity.

Several factors determine VFP TCAM consumption:

- **Type of filters configured**—Different filter types consume different amounts of VFP TCAM space. VN2VF\_Port FIP snooping filters never consume more than one slice. MVR rules and VN2VN\_Port FIP snooping filters consume entries in a slice at a rate of one entry per MVR rule or VN2VN\_Port session. FBF filter terms consume entries in a slice at a rate of two entries per FBF filter term.
- **Number of filters configured**—Although the number of filters does not affect the number of slices allocated to the VN2VF\_Port FIP snooping filter type (it is always one slice for one or more VN2VF\_Port FIP snooping filters and no slice for no FIP snooping filters), the number of VN2VN\_Port FIP snooping filters, MVR rules, and FBF filter terms that you configure determine how many VFP TCAM slices are required for each filter type.

For example, if you configure 257 MVR rules, the MVR rule entries consume 2 slices. One slice stores 256 MVR rules (entries), and one slice stores 1 MVR rule (entry). In this case, if you can eliminate one MVR rule, you can free a slice to allocate to other filter types.

- **Sequence of filter configuration**—If you configure too many filters for the VFP TCAM to store, the last filters you configure are not stored in the TCAM.

Always check the syslog after you configure FBF filters or MVR rules to ensure that the configuration was not rejected. If you enable FIP snooping on access ports, check the syslog to ensure that the configuration was not rejected due to lack of VFP TCAM space.

If you check the syslog and a filter configuration has been rejected, delete the filters that were rejected from the configuration.



**TIP:** If you no longer need an FBF filter or an MVR rule, delete it from the configuration to conserve VFP TCAM space. Enable VN2VF\_Port or VN2VN\_Port FIP snooping on access ports only if the switch port is directly connected to FCoE devices. (FIP snooping should be performed at the access edge. FIP snooping should not be performed on traffic that has already been snooped and filtered at the access edge. If another switch between the QFX Series or QFabric system and the FCoE devices already performs FIP snooping, do not enable FIP snooping on the QFX Series or QFabric system.)

---

### ***Always Delete Rejected Filter Configurations***

The switch does not return a commit error if it rejects any portion of a configuration. Instead, the switch sends a syslog message to report the rejected portion of the configuration. The rejected portion of the configuration remains on the switch, but does not function.

After you configure FBF filters or MVR rules, or enable FIP snooping, check the syslog messages to ensure that the switch accepted the configuration. If the switch rejected any portion of the configuration, delete that portion of the configuration. (You do not need to delete the portion of the configuration that was accepted, unless you want to reconfigure those filters or rules.)



**CAUTION:** If you do not delete rejected filter configurations, and if you reboot the system, you cannot predict which filters the system installs after the reboot. For example, a switch with the following configuration has more configured filters than the VFP TCAM can support:

- VN2VF\_Port FIP snooping sessions (always consumes one slice)
- 20 MVR rules (consume one slice)
- 300 FBF filters (attempt to consume three slices, but because only two slices are available, 256 filters consume two slices, and the remaining 44 filters are rejected)

If you do not delete the 44 rejected FBF filters, then if the switch reboots, the 44 FBF filters that were rejected might be accepted, and 44 different FBF filters might be rejected. This unpredictable behavior is the reason that you should check the syslog messages after you configure filters, and if any filters were rejected, you should always delete the rejected filters from the configuration.

---

### **Related Documentation**

- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)
- [Understanding Filter-Based Forwarding on page 4637](#)
- [Understanding Multicast VLAN Registration on page 4241](#)

- [Configuring VN2VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5155](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5167](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 4697](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 4321](#)

## Understanding MC-LAGs on an FCoE Transit Switch

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX3500 switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree protocol (STP).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX3500 switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because MC-LAGs do not carry forwarding class and IEEE 802.1p priority information.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as passthrough transit switch ports.

QFX3500 Node devices support MC-LAGs. QFabric Node devices do not support MC-LAGs.


This topic describes:

- [Supported Topology on page 5047](#)
- [FIP Snooping and FCoE Trusted Ports on page 5049](#)
- [CoS and Data Center Bridging \(DCB\) on page 5050](#)

### Supported Topology

QFX3500 switches that are not directly connected to FCoE hosts and that act as passthrough transit switches support MC-LAGs for FCoE traffic in an *inverted-U* network topology, as shown in [Figure 194 on page 5048](#):

Figure 194: Supported Topology for an MC-LAG on an FCoE Transit Switch



The following rules and guidelines apply to MC-LAGs when used for FCoE traffic. The rules and guidelines help ensure the proper handling and lossless transport characteristics required for FCoE traffic:

- The two QFX3500 switches that form the MC-LAG (Switches S1 and S2) cannot use ports that are part of an FCoE-FC gateway fabric. The MC-LAG switch ports must be passthrough transit switch ports (used as part of an intermediate transit switch that is not directly connected to FCoE hosts).
- MC-LAG Switches S1 and S2 cannot be not directly connected to the FCoE hosts.
- The two QFX3500 switches that serve as access devices for FCoE hosts (FCoE Transit Switches TS1 and TS2) use standard LAGs to connect to MC-LAG Switches S1 and S2. FCoE Transit Switches TS1 and TS2 can be standalone QFX3500 switches or they can be Node devices in a QFabric.
- Transit Switches TS1 and TS2 must use transit switch ports for the FCoE hosts and for the standard LAGs to MC-LAG Switches S1 and S2.
- Enable FIP snooping on the FCoE VLAN on Transit Switches TS1 and TS2. You can configure either VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping or VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, depending on whether the FCoE hosts need to access targets in the FC SAN (VN2VF\_Port FIP snooping) or targets in the Ethernet network (VN2VN\_Port FIP snooping).

FIP snooping should be performed at the access edge and is not supported on MC-LAG switches. Do not enable FIP snooping on MC-LAG Switches S1 and S2. (Do not enable FIP snooping on the MC-LAG ports that connect Switches S1 and S2 to Switches TS1 and TS2 or on the LAG ports that connect Switch S1 to S2.)

- The class of service (CoS) configuration must be consistent on the MC-LAG switches. Because the MC-LAG carries no forwarding class or priority information, each MC-LAG

switch needs to have the same CoS configuration to support lossless transport. (On each MC-LAG switch, the name, egress queue, and CoS provisioning of each forwarding class must be the same, and the PFC configuration must be the same.)

### ***Transit Switches (Server Access)***

The role of FCoE Transit Switches TS1 and TS2 is to connect FCoE hosts in a multi-homed fashion to the MC-LAG switches. In essence, Transit Switches TS1 and TS2 act as access switches for the FCoE hosts. (FCoE hosts are directly connected to Transit switches TS1 and TS2.)

The transit switch configuration depends on whether you want to do VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, and whether the QFX3500 transit switches also have ports configured as part of an FCoE-FC gateway virtual fabric. Ports that the QFX3500 switch uses in an FCoE-FC gateway virtual fabric cannot be included in the transit switch LAG connection to the MC-LAG switches. (Ports cannot belong to both a transit switch and an FCoE-FC gateway; you must use different ports for each mode of operation.)

### ***MC-LAG Switches (FCoE Aggregation)***

The role of MC-LAG Switches S1 and S2 is to provide redundant, load-balanced connections between FCoE transit switches. In essence, MC-LAG Switches S1 and S2 act as aggregation switches. FCoE hosts are not directly connected to the MC-LAG switches.

The MC-LAG switch configuration is the same regardless of which type of FIP snooping that FCoE Transit Switches TS1 and TS2 perform.

### **FIP Snooping and FCoE Trusted Ports**

To maintain secure access, enable VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping at the transit switch access ports connected directly to the FCoE hosts. FIP snooping should be performed at the access edge of the network to prevent unauthorized access. For example, in [Figure 194 on page 5048](#), you enable FIP snooping on the FCoE VLANs on Transit Switches TS1 and TS2 that include the access ports connected to the FCoE hosts.

Do not enable FIP snooping on the switches used to create the MC-LAG. For example, in [Figure 194 on page 5048](#), you would not enable FIP snooping on the FCoE VLANs on Switches S1 and S2.

Configure links between switches as FCoE trusted ports to reduce FIP snooping overhead and ensure that the system performs FIP snooping only at the access edge. In the sample topology, configure the Transit Switch TS1 and TS2 LAG ports connected to the MC-LAG switches as FCoE trusted ports, configure the Switch S1 and S2 MC-LAG ports connected to Switches TS1 and TS2 as FCoE trusted ports, and configure the ports in the LAG that connects Switches S1 to S2 as FCoE trusted ports.

## CoS and Data Center Bridging (DCB)

---

The MC-LAG links do not carry forwarding class or priority information. The following CoS properties must have the same configuration on each MC-LAG switch or on each MC-LAG interface to support lossless transport:

- FCoE forwarding class name—For example, the forwarding class for FCoE traffic could use the default **fcoe** forwarding class on both MC-LAG switches.
- FCoE output queue—For example, the **fcoe** forwarding class could be mapped to queue 3 on both MC-LAG switches (queue 3 is the default mapping for the **fcoe** forwarding class).
- Classifier—The forwarding class for FCoE traffic must be mapped to the same IEEE 802.1p code point on each member interface of the MC-LAG on both MC-LAG switches. For example, the FCoE forwarding class **fcoe** could be mapped to IEEE 802.1p code point 011 (code point 011 is the default mapping for the **fcoe** forwarding class).
- Priority-based flow control (PFC)—PFC must be enabled on the FCoE code point on each MC-LAG switch and applied to each MC-LAG interface using a congestion notification profile.

You must also configure enhanced transmission selection (ETS) on the MC-LAG interfaces to provide sufficient scheduling resources (bandwidth, priority) for lossless transport. The ETS configuration can be different on each MC-LAG switch, as long as enough resources are scheduled to support lossless transport for the expected FCoE traffic.

LLDP and DCBX must be enabled on each MC-LAG member interface (LLDP and DCBX are enabled by default on all interfaces).



**NOTE:** As with all other FCoE configurations, FCoE traffic requires a dedicated VLAN that carries only FCoE traffic, and IGMP snooping must be disabled on the FCoE VLAN.

---

### Related Documentation

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5121](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)

## Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.

This topic describes:

- [DCBX Basics on page 5051](#)
- [DCBX Modes and Support on page 5052](#)
- [DCBX Attribute Types on page 5055](#)
- [DCBX Application Protocol TLV Exchange on page 5056](#)
- [DCBX and PFC on page 5057](#)
- [DCBX and ETS on page 5057](#)

### DCBX Basics

---

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

---

### DCBX Modes and Support

---

This section describes DCBX support on the QFX Series:

- [DCBX Modes \(Versions\) on page 5052](#)
- [Autonegotiation on page 5054](#)
- [CNA Support for DCBX Modes on page 5055](#)
- [Interface Support for DCBX on page 5055](#)

#### **DCBX Modes (Versions)**

The QFX Series supports the two most common DCBX modes:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



**NOTE:** The QFX Series does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The QFX Series drops LLDP frames that contain pre-CEE DCBX TLVs.

[Table 392 on page 5053](#) summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:



**Table 392: Summary of Differences Between IEEE DCBX and DCBX Version 1.01**

| Characteristic                                                                   | IEEE DCBX                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | DCBX Version 1.01                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OUI                                                                              | 0x0080c2                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 0x001b21                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Frame Format                                                                     | Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs                                                                                                                                                                                                                                                                       | Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.                                                                                                                                                                                                                                                                                                         |
| Symmetric/asymmetric configuration with peer                                     | Asymmetric or symmetric                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Symmetric only                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Differences in the <b>show dcbx interface interface-name</b> operational command | <ul style="list-style-type: none"> <li>• Synchronization information is not shown because symmetric configuration is not required.</li> <li>• Operational state information is not shown because the operational states do not have to be symmetric.</li> <li>• TLV type is shown because unique TLVs are sent for each DCBX attribute.</li> <li>• ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs.</li> </ul> | <ul style="list-style-type: none"> <li>• Synchronization information is shown because symmetric configuration is required.</li> <li>• Operational state information is shown because the operational states do have to be symmetric.</li> <li>• TLV type is not shown because one TLV is used for all attribute information.</li> <li>• Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).</li> </ul> |

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—<http://www.ieee802.org/1/files/public/docs2008/az-wadelaar-dcbx-capability-exchange-discovery-protocol-T108-v101.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>



**NOTE:** As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.

- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



**NOTE:** On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

### ***Autonegotiation***

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on QFX3500 switches and QFabric systems:

- QFX3500 switch—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



**NOTE:** If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

### ***CNA Support for DCBX Modes***

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on QFX Series interfaces depends on the DCBX features that the CNAs in your network support.

### ***Interface Support for DCBX***

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

### ***DCBX Attribute Types***

DCBX has three attribute types:

- **Informational**—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- **Asymmetric**—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- **Symmetric**—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 5055](#)
- [Symmetric Attributes on page 5056](#)

### ***Asymmetric Attributes***

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing,” the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface

sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

### ***Symmetric Attributes***

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

### **DCBX Application Protocol TLV Exchange**

---

DCBX advertises the switch’s capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 5056](#)
- [FCoE Application Protocol TLV Exchange on page 5056](#)
- [Disabling Application Protocol TLV Exchange on page 5057](#)

### ***Application Protocol TLV Exchange***

For all applications, DCBX advertises the application’s state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application’s TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

### ***FCoE Application Protocol TLV Exchange***

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



**NOTE:** If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

---

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.



**NOTE:** If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

### ***Disabling Application Protocol TLV Exchange***

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

## **DCBX and PFC**

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

## **DCBX and ETS**

This section describes:

- [Default DCBX ETS Advertisement on page 5058](#)
- [ETS Advertisement and Peer Configuration on page 5058](#)
- [ETS Recommendation TLV on page 5058](#)

### ***Default DCBX ETS Advertisement***

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

### ***ETS Advertisement and Peer Configuration***

DCBX does not control the switch's ETS (hierarchical scheduling) operational state. If the connected peer is configured as "willing," DCBX pushes the switch's ETS configuration to the switch's peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

### ***ETS Recommendation TLV***

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is "willing," it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



**NOTE:** You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

#### Related Documentation

- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding FCoE on page 4968](#)
- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)

## Understanding DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

---

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 5060](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case (see [“Application Maps” on page 5061](#)) and only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic.

This topic describes:

- [Applications on page 5060](#)
- [Application Maps on page 5061](#)
- [Classifying and Prioritizing Application Traffic on page 5062](#)
- [Enabling Interfaces to Exchange Application Protocol Information on page 5063](#)
- [Disabling DCBX Application Protocol Exchange on page 5063](#)

---

### Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise, except FCoE if FCoE is the only application that you want the interface to advertise.





**NOTE:** If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

### Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

---

### Classifying and Prioritizing Application Traffic

---

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

### Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



**NOTE:** You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

### Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

#### Related Documentation

- [Understanding DCBX on page 5051](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)

- [Example: Configuring Unicast Classifiers on page 5658](#)

## Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

The QFX Series supports two methods of flow control:

- IEEE 802.3X Ethernet PAUSE
- IEEE 802.1Qbb priority-based flow control (PFC)

Ethernet PAUSE and PFC are link-level flow control mechanisms.

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority is mapped to a 3-bit IEEE 802.1p CoS code point flag in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on specified types of traffic (for example, Fibre Channel over Ethernet traffic).



**NOTE:** Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error. Ethernet PAUSE and PFC are mutually exclusive configurations on an interface.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

- [Ethernet PAUSE on page 5064](#)
- [PFC on page 5068](#)
- [Lossless Transport Support Summary on page 5072](#)

### Ethernet PAUSE

---

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions

on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

*Symmetric flow control* means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

*Asymmetric flow control* allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, the PFC configuration overrides Ethernet PAUSE flow control.) The QFX Series supports both symmetric and asymmetric flow control.

- [Symmetric Flow Control on page 5065](#)
- [Asymmetric Flow Control on page 5066](#)

### ***Symmetric Flow Control***

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

### ***Asymmetric Flow Control***

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 393 on page 5066](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

**Table 393: Asymmetric Ethernet PAUSE Flow Control Configuration**

| Receive (Rx) Buffer | Transmit (Tx) Buffer | Configured Flow Control State                                                                                                                                                                                         |
|---------------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| On                  | Off                  | Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).         |
| Off                 | On                   | Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.) |
| On                  | On                   | Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.                                                               |
| Off                 | Off                  | Ethernet PAUSE flow control is disabled.                                                                                                                                                                              |

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, the QFX Series supports autonegotiation of Ethernet PAUSE with the connected peer. (The QFX Series does not support autonegotiation on 10-Gigabit Ethernet interfaces.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected

peer using a combination of the Ethernet PAUSE and ASM\_DIR bits, as described in [Table 394 on page 5067](#):

**Table 394: Flow Control State Advertised to the Connected Peer (Autonegotiation)**

| Rx Buffer State | Tx Buffer State | PAUSE Bit | ASM_DIR Bit | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------|-----------------|-----------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Off             | Off             | 0         | 0           | The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.                                                                                                                                                                                                                                                                                                                                                                                       |
| On              | On              | 1         | 0           | The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).                                                                                                                                                                                                                                                                                                                                      |
| On              | Off             | 0         | 1           | The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).                                                                                                                                                                                                                                                                                                                                     |
| Off             | On              | 1         | 1           | The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.) |

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control behavior (resolution) between them, as shown in [Table 395 on page 5068](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series and on the connected peer (also known as the link partner). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



**NOTE:** In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

**Table 395: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces**

| Local Interface (QFX Series) |             | Peer Interface |             | Local Resolution                                                  | Peer Resolution                                                   |
|------------------------------|-------------|----------------|-------------|-------------------------------------------------------------------|-------------------------------------------------------------------|
| PAUSE Bit                    | ASM_DIR Bit | PAUSE Bit      | ASM_DIR Bit |                                                                   |                                                                   |
| 0                            | 0           | Don't care     | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                            | 1           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                            | 1           | 1              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0                            | 1           | 1              | 1           | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive | Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive |
| 1                            | 0           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1                            | 0           | 1              | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |
| 1                            | 1           | 0              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1                            | 1           | 0              | 1           | Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive |
| 1                            | 1           | Don't care     | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |



**NOTE:** For your convenience, [Table 395 on page 5068](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

## PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits



a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic with the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as Fibre Channel over Ethernet (FCoE), LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of link-level flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)
- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 396 on page 5070](#):

**Table 396: Default PFC Priority to Queue and Forwarding Class Mapping**

| IEEE 802.1p Priority (Code Point) | Queue | Forwarding Class |
|-----------------------------------|-------|------------------|
| 0 (000)                           | 0     | best-effort      |
| 1 (001)                           | 1     | best-effort      |
| 2 (010)                           | 2     | best-effort      |
| 3 (011)                           | 3     | fcoe             |
| 4 (100)                           | 4     | no-loss          |
| 5 (101)                           | 5     | best-effort      |
| 6 (110)                           | 6     | network-control  |
| 7 (111)                           | 7     | network-control  |

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default forwarding class configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) and [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#).

You enable PFC on a priority by:

1. Specifying the IEEE 802.1p code point to pause in the input stanza of a CNP
2. Applying the CNP to the ingress interfaces on which you want to pause the traffic



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
  1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
  2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
  3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

Although unicast traffic and multdestination (multicast, broadcast, and destination lookup fail) traffic must use different classifiers, you can map a unicast queue (queue 0 through 7) and a multdestination queue (queue 8, 9, 10, or 11) to the same PFC priority so that both unicast and multicast traffic use that priority. Do not map multdestination traffic to lossless priorities. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



**NOTE:** You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone QFX3500 switches and QFX3600 switches, you can create two CNPs with an explicitly configured output stanza.

When a QFX3500 switch or a QFX3600 switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 5506 describes configuring output flow control.

---

### Lossless Transport Support Summary

---

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



**NOTE:** Junos OS Release 12.2 introduced changes to the way the QFX Series handles lossless forwarding classes (including the default fcoe and no-loss forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the fcoe and no-loss forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the fcoe or the no-loss forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the fcoe or the no-loss forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit fcoe and no-loss forwarding class configuration before you upgrade to Junos OS Release 12.2.

See [“Overview of CoS Changes Introduced in Junos OS Release 12.2” on page 5428](#) for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the fcoe and no-loss forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new no-loss packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) provides detailed information about the explicit configuration of lossless priorities and about the default configuration of lossless priorities, including the input and output stanzas of the CNP.



**NOTE:** PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

#### Related Documentation

- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)

- [Understanding DCB Features and Requirements on page 4965](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)

## Understanding Fibre Channel Terminology

To understand the Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) capabilities of the QFX Series, you should become familiar with the terms defined in [Table 397 on page 5074](#).

**Table 397: Fibre Channel Terms**

| Term                            | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| addressing mode                 | <p>Format for the locally unique MAC address the FC switch assigns to FCoE devices for FCoE transactions after FIP establishes a connection between an FCoE device and the FC switch. The two addressing modes are <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>. The QFX Series supports only FPMA.</p> <p>During FLOGI or FDISC, the ENode advertises the addressing modes it supports. If the FC switch supports an addressing mode that the ENode uses, the virtual link can be established, and the devices can communicate.</p> <p>See also <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>.</p> |
| ALL-ENode-MACs                  | <p>Well-known multicast MAC address to which all FCoE ENodes listen. FCFs send multicast <i>FIP discovery advertisement</i> messages and <i>FIP keepalive</i> messages to the ALL-ENode-MACs address so that ENodes can discover and maintain connections to FCFs. The hexadecimal format of the address is <b>01:10:18:01:00:01</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>                                                                                                                                                                                                                                                                                               |
| ALL-FCF-MACs                    | <p>Well-known multicast MAC address to which all FCFs listen. ENodes send multicast <i>FIP discovery solicitation</i> messages to the ALL-FCF-MACs address to find out which FCFs can accept a login. The hexadecimal format of the address is <b>01:10:18:01:00:02</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>                                                                                                                                                                                                                                                                                                                                                            |
| congestion notification         | See <i>quantized congestion notification (QCN)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| converged network adapter (CNA) | <p>Physical adapter that combines the functions of a Fibre Channel <i>host bus adapter (HBA)</i> to process FCoE frames and a <i>lossless Ethernet network interface card (NIC)</i> to process non-FCoE Ethernet frames. CNAs have one or more Ethernet ports. CNAs encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.</p> <p>See also <i>host bus adapter (HBA)</i>.</p>                                                                                                                                                                                                                        |

Table 397: Fibre Channel Terms (*continued*)

| Term                                                     | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data center bridging (DCB)                               | <p>Set of IEEE specifications that enhance Ethernet to allow it to support converged Ethernet (LAN) and Fibre Channel (SAN) traffic on one Ethernet network. DCB features include <i>priority-based flow control (PFC)</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, <i>quantized congestion notification (QCN)</i>, and full-duplex 10-Gigabit Ethernet ports.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>Ethernet PAUSE</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, and <i>quantized congestion notification (QCN)</i>.</p>                      |
| expansion port (E_Port)                                  | An expansion port in an FC switch/FCF that connects the FC switch/FCF to the E_Port of another FC switch/FCF to form an <i>Interswitch Link (ISL)</i> in a common FC fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Data Center Bridging Capability Exchange protocol (DCBX) | <p>Discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB)</p> <p>See also <i>data center bridging (DCB)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| enhanced transmission selection (ETS)                    | <p>Mechanism that provides finer granularity of bandwidth management within a link.</p> <p>See also <i>data center bridging (DCB)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| ENode                                                    | See <i>FCoE Node (ENode)</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| ENode MAC                                                | <p><i>Lossless Ethernet MAC</i> paired with an <i>FCoE controller</i> in an ENode.</p> <p>See also <i>FCoE node (ENode)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ENode MAC address                                        | Globally unique address assigned to the CNA by the manufacturer and used to identify the node for FIP transactions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Ethernet PAUSE                                           | <p>As defined in IEEE 802.3X, a flow control mechanism that temporarily stops the transmission of Ethernet frames on a link for a specified period. A receiving element sends an Ethernet PAUSE frame when a sender transmits data faster than the receiver can accept it. Ethernet PAUSE affects the entire link, not just an individual flow. An Ethernet PAUSE frame temporarily stops all traffic transmission on the link and allows the receiver's input buffer to empty sufficiently to restart traffic on the link. Ethernet PAUSE messages are sent to the previous hop and do not automatically propagate to the source of the congestion.</p> <p>See also <i>priority-based flow control (PFC)</i>.</p> |
| fabric                                                   | Interconnection of network nodes using one or more network switches that function as a network single logical entity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 397: Fibre Channel Terms (*continued*)

| Term                               | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fabric discovery (FDISC)           | <p>Subsequent logins from the same ENode for different users, applications, or virtual machines after an ENode performs an initial FLOGI to log in to a switch.</p> <p>FC and FIP FDISC messages serve the same function in FC and FCoE networks, respectively. N_Ports send FC FDISC messages to the FC switch and VN_Ports send FIP FDISC messages to the FCF.</p> <p>After an N_Port acquires its initial N_Port ID through the FC FLOGI process, it can acquire additional N_Port IDs by sending an FC FDISC with a new worldwide port name and a source ID of 0x000000. The new port name and blank source ID tell the FC switch to assign a new N_Port ID to the N_Port. The different N_Port IDs allow multiple virtual machines or users on the N_Port to have separate, secure virtual links on the same physical N_Port. These additional ports are also referred to as VN_Ports.</p> <p>FIP FDISC works the same way, except the VN_Port logs in using a FIP FLOGI message.</p> <p>See also <i>fabric login (FLOGI)</i> and <i>N_Port ID</i>.</p> |
| fabric login (FLOGI)               | <p>Creation of a logical connection to the FC switch and establishment of a node's operating environment.</p> <p>For FC devices, an N_Port logs in to the FC network by sending an FC FLOGI message to the F_Port of an FC switch.</p> <p>For FCoE devices, a VN_Port logs in to the FC network by sending a FIP FLOGI message to the VF_Port of an FC switch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| fabric port (F_Port)               | <p>FC port on an FC switch or an FCF that connects point-to-point to an FC node port (N_Port) on an FC host (server or storage device). An F_Port provides access to fabric services for FC devices.</p> <p>F_Ports are intermediate ports in a connection between FC device end-point N_Ports. For example, a connection between an FC host server and an FC storage device through an FC switch looks like this: FC server N_Port to FC switch ingress F_Port to FC switch egress F_Port to FC storage device N_Port.</p> <p>See also <i>node port (N_Port)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| fabric-provided MAC address (FPMA) | <p>MAC address that an FCF assigns to a single ENode MAC through the FLOGI or FDISC process that is unique to the local fabric. The FPMA uniquely identifies a single VN_Port at that ENode MAC in FCoE transactions with the FCF.</p> <p>Because an ENode can have more than one ENode MAC, an FCF can assign multiple FPMAs to an ENode, one FPMA per ENode MAC.</p> <p>An FPMA is a 48-bit value that consists of two 24-bit values, the N_Port ID and the FC-MAP value. The N_Port ID uniquely identifies the VN_Port and the FC-MAP value identifies the FCF.</p> <p>See also <i>FCoE node (ENode)</i>, <i>N_Port ID</i>, and <i>FCoE mapped address prefix (FC-MAP)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                           |
| FCF-MAC                            | Lossless Ethernet MAC paired with an FCoE controller in an FCF. The FCF-MAC enables the FCF to handle FCoE traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



Table 397: Fibre Channel Terms (*continued*)

| Term                                | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE controller                     | <p>Instantiates and terminates VN_Port and VF_Port instances on an ENode. An ENode can have more than one FCoE controller. Each FCoE controller is paired with a lossless Ethernet MAC on the ENode.</p> <p>See also <i>lossless Ethernet MAC</i>.</p>                                                                                                                                                                                                                                                                          |
| FC forwarder (FCF)                  | Alternative term and acronym to refer to an FC switch that has all physical Fibre Channel ports and the necessary set of services as defined in the T11 Organization <i>Fibre Channel Switched Fabric</i> (FC-SW) standards.                                                                                                                                                                                                                                                                                                    |
| FCoE forwarder (FCF)                | Defined by the <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification available at <a href="http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf">http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf</a> as a device that has the necessary set of services as defined in FC-SW and the FCoE capabilities to act as an FCoE-based FC switch.                                                                                                                                                                        |
| FCoE Initialization Protocol (FIP)  | <p>Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP enables FCoE devices and FC switches to discover one another. Through FIP, FCoE nodes can log in to an FC switch, access the SAN FC fabric, and communicate with target FC devices. FIP messages also maintain the connection between the FCoE initiator and the FCF.</p> <p>FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic.</p>                             |
| FCoE link endpoint (LEP)            | Virtual FC interface mapped onto a physical Ethernet interface to handle FC frame encapsulation and de-encapsulation and transmission and reception of FC frames encapsulated in Ethernet through a single virtual link.                                                                                                                                                                                                                                                                                                        |
| FCoE mapped address prefix (FC-MAP) | <p>24-bit value that identifies the FC switch and is half of the 48-bit FPMA MAC address. The FC-MAP value can be configured on the FC switch and has a default value of 0EFC00h. The FC-MAP value was originally called the Fibre Channel Organizationally Unique Identifier (FC-OUI).</p> <p>See also <i>fabric-provided MAC address (FPMA)</i>.</p>                                                                                                                                                                          |
| FCoE node (ENode)                   | <p>Fibre Channel node that has one or more lossless Ethernet MACs, each paired with an <i>FCoE Controller</i> in order to transmit FCoE frames. An ENode combines FCoE termination functions and the FC stack on a CNA. ENodes present virtual FC interfaces to FC switches or FCFs in the form of VN_Ports, which can establish FCoE virtual links with FC switch/FCF VF_Ports. ENodes perform FCoE related functions in a <i>converged network adapter (CNA)</i>.</p> <p>See also <i>converged network adapter (CNA)</i>.</p> |
| FCoE-FC gateway                     | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FC ports.                                                                                                                                                                                                                                                                                                                                                                                                         |
| FCoE-FCoE gateway                   | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FCoE ports.                                                                                                                                                                                                                                                                                                                                                                                                       |
| FC-FC gateway                       | A form of N_Port virtualizer in which the node-facing ports are FC ports and the FC switch-facing ports are FC ports.                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 397: Fibre Channel Terms (*continued*)

| Term                                                      | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE transit switch (also known as a FIP snooping bridge) | <p>Switch with a minimum set of features designed to support FCoE Layer 2 forwarding and FCoE security. The switch can also have optional additional features.</p> <p>Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the FCoE application TLV</li> <li>• FIP snooping (minimum support is FIP automated filter programming at the ENode edge)</li> </ul> <p>Additional FIP snooping capabilities can include learning the virtual FC connection paths (VN2VF, VN2VN, or VE2VE) and monitoring the FIP keepalive mechanisms. Other optional capabilities can also enhance FCoE within the standards. FIP snooping is typically configurable on a per-VLAN basis.</p> <p>A transit switch has an FC stack even though it is not an FC switch or an FCF.</p> |
| FCoE VLAN                                                 | VLAN dedicated to carrying only FCoE traffic. FCoE traffic must travel in a VLAN. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE traffic must travel in a different VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Fibre Channel                                             | High-speed network technology used for storage area networks (SANs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Fibre Channel fabric                                      | <p>Network of Fibre Channel devices that allows communication among devices, device name lookup, security, and redundancy.</p> <p>Also a local fabric on a QFX3500 switch with FCoE interfaces connected to FCoE devices on the Ethernet network and native FC interfaces connected to an FC switch in a SAN.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Fibre Channel ID (FCID)                                   | <p>24-bit value the FC switch assigns to the N_Port or VN_Port as a unique identifier within the local FC network. The FCID consists of an 8-bit domain value, an 8-bit area value, and an 8-bit port value. The FCID is sometimes called an N_Port ID.</p> <p>See also <i>N_Port ID</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Fibre Channel over Ethernet (FCoE)                        | <p>Standard for transporting FC frames over Ethernet networks. FCoE encapsulates Fibre Channel frames in Ethernet so that the same high-speed Ethernet physical infrastructure can transport both data and storage traffic while preserving the lossless CoS that FC requires. FCoE has its own EtherType (0x8906) to differentiate it from other Ethernet traffic.</p> <p>FCoE runs on a DCB network. FCoE servers connect to a switch that supports both FCoE and native FC protocols. This allows FCoE servers on the Ethernet network to access FC storage devices in the SAN fabric on one converged network.</p> <p>See also <i>data center bridging (DCB)</i>.</p>                                                                                                                                                                                                                                                                             |

Table 397: Fibre Channel Terms (*continued*)

| Term                        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fibre Channel services      | Functions required for establishing FC network connectivity among devices and for managing devices on the FC network, such as login servers, domain managers, name servers, and zone servers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| FC stack                    | <p>FC or FCoE protocol capability implemented on a device to support the FC or FCoE functionality. Having an FC stack does not imply consuming a domain ID.</p> <p>Each FC or FCoE enabled server or storage device has an FC stack. Similarly, an FC or FCoE switch, an FCF, an FCoE-FC gateway, and an FCoE transit switch have FC stacks.</p>                                                                                                                                                                                                                                                                                                                                                                           |
| Fibre Channel switch        | Network switch that implements the Fibre Channel protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| FIP discovery advertisement | <p>Multicast or unicast message that the FC switch (or FCF) transmits to ENodes to advertise the switch's presence on the network so that ENodes can discover the switch and request to log in to the FC fabric.</p> <p>The FC switch periodically sends multicast FIP discovery advertisements to the ALL-ENode-MACs address, a well-known address to which all ENodes listen. The multicast messages advertise the FC switch to all ENodes on the VLAN and serve as keepalive messages to maintain connectivity between the FC switch and ENodes.</p> <p>When an ENode sends a FIP discovery solicitation message to the FC switch, the FC switch responds with a unicast FIP discovery advertisement to that ENode.</p> |
| FIP discovery solicitation  | <p>Multicast or unicast message that an ENode transmits to FC switches (or FCFs) to find compatible switches in the network.</p> <p>When an ENode initializes, it sends a multicast FIP discovery solicitation to the ALL-FCF-MACs address, a well-known address to which all FC switches and FCFs listen. Compatible switches reply with a unicast FIP discovery advertisement.</p> <p>The ENode compiles a list of compatible switches, selects a switch, and logs in to that switch.</p>                                                                                                                                                                                                                                |
| FIP keepalive               | Periodic multicast FIP discovery advertisement sent from the FC switch or FCF to all ENodes to maintain connectivity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 397: Fibre Channel Terms (*continued*)

| Term                   | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FIP snooping           | <p>For VN_Port to VF_Port (VN2VF) paths (Technical Committee T11 BB-FC-5 specification), FIP snooping is a security feature enabled for FCoE VLANs on an Ethernet switch that connects ENodes to FC switches or FCFs. FIP snooping inspects data in FIP frames and uses that data to create firewall filters. The filters permit only traffic from sources that perform a successful FLOGI to the FC switch. All other traffic on the VLAN is denied. FIP snooping filters are installed on the ports in the FCoE VLAN.</p> <p>For VN_Port to VN_Port (VN2VN) paths (Technical Committee T11 BB-FC-6 specification), the FIP snooping security feature filters access between VN_Ports in a similar manner to VN2VF_Port FIP snooping.</p> <p>FIP snooping can also apply similarly to VE_Port to VE_Port (VE2VE) paths.</p> <p>FIP snooping can also snoop to provide additional visibility of FCoE Layer 2 operation.</p> <p>See also <i>FCoE node (ENode)</i>.</p> |
| FIP snooping bridge    | See <i>FCoE transit switch</i> and <i>FIP snooping</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| host bus adapter (HBA) | Physical mechanism that connects a host system to other FC network and storage devices. HBAs have a unique worldwide node name (WWNN) for the HBA node, which all of the ports on the HBA share, and each port on an HBA has a unique worldwide port name (WWPN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| initiator              | System component that originates an I/O command over an I/O bus or network. An FCoE server sending a request to an FC storage device is an example of an initiator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| iSCSI transit switch   | <p>Layer 2 Ethernet switch with a minimum set of best-practice Ethernet features to support iSCSI, along with optional enhancements. Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• IEEE 802.3X asymmetric and symmetric flow control on ports not running in DCB mode</li> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the iSCSI application TLV</li> </ul> <p>Other capabilities such as Internet storage name service (iSNS) are optional.</p>                                                                                                                                                                                                                                                                                                                                                               |
| interswitch link (ISL) | Link between the <i>E_Ports</i> of two FC switches in a common FC fabric. When two FCoE-based FC switches are connected together, there is a virtual ISL through Layer 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| logout (LOGO)          | <p>For FC devices, an N_Port logs out from the FC network by sending an FC LOGO message to the F_Port of an FC switch. The switch can also send a LOGO message to an N_Port to terminate its connection.</p> <p>For FCoE devices, a VN_Port logs out from the FC network by sending a FIP LOGO message to the VF_Port of an FC switch. The switch can also send a LOGO message to a VN_Port to terminate its connection.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 397: Fibre Channel Terms (*continued*)

| Term                            | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lossless Ethernet MAC           | <p>Full-duplex Ethernet MAC that implements Ethernet extensions to avoid Ethernet frame loss due to congestion and supports at least 2.5-KB jumbo frames. Each lossless Ethernet MAC combines with an FCoE Controller to perform FCoE termination functions on an ENode.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>quantized congestion notification (QCN)</i>, <i>FCoE controller</i>, and <i>FCoE node (ENode)</i>.</p>                                                                                                                                                                                                                                                                                                                 |
| lossless Ethernet network       | Ethernet network composed of only full-duplex links and lossless Ethernet MACs and with CoS and flow control to prevent dropping of frames.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| lossless transport              | In DCB networks, the ability to switch FCoE frames over an Ethernet network without dropping any frames. Lossless transport uses mechanisms such as priority-based flow control and quantized congestion notification to control traffic flows and avoid congestion.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| N_Port ID                       | See <i>Fibre Channel ID (FCID)</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| N_Port ID virtualizer           | <p>Presents itself as an FC or FCoE switch to external devices, but connects to an actual FC or FCoE switch in the other direction to provide the FC-SW services.</p> <p>An N_Port ID virtualizer logs in to the actual FC or FCoE switch in the same way as a normal node device and uses the NPIV mechanism to proxy incoming FLOGIs to FDISCs on the actual FC or FCoE switch.</p> <p>An N_Port ID virtualizer has an FC stack even though it is not an FC switch or an FCF.</p> <p>The acronym <i>NPV</i> is commonly used for N_Port ID virtualizer even though the acronym is not defined in the standards.</p>                                                                                                                                        |
| N_Port ID Virtualization (NPIV) | <p>NPIV enables a physical N_Port to acquire multiple N_Port IDs. Each N_Port ID maps to a different application (such as a virtual machine) or to a different user. This allows you to associate one F_Port with many N_Port IDs and create multiple discrete, secure virtual links over one physical point-to-point connection.</p> <p>NPIV increases resource and bandwidth utilization and allows the implementation of access control, zoning, and port security on a per-application or per-user basis.</p> <p>After an N_Port performs a FLOGI and receives its first N_Port ID, it can request more N_Port IDs by sending FDISC messages.</p> <p>See also <i>fabric login (FLOGI)</i>, <i>fabric discovery (FDISC)</i>, and <i>virtual link</i>.</p> |

Table 397: Fibre Channel Terms (*continued*)

| Term                              | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node port (N_Port)                | <p>N_Ports can be in two modes:</p> <ul style="list-style-type: none"> <li>Fabric N_Port—Node port that is an FC host or storage device end port in a point-to-point link between the device and the F_Port of an FC switch. The point-to-point link can be virtual or physical.</li> <li>Point-to-point N_Port—Node port that connects to another N_Port. The QFX3500 switch does not support this configuration.</li> </ul> <p>N_Ports handle creation, detection, and flow of messages to and from the connected devices.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| node worldwide name (NWWN)        | WWN that is unique worldwide and is assigned to an FC node. An NWWN is valid for multiple ports that are on that node (this identifies the ports as network interfaces of a particular node).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| port mode                         | <p>Role that the port plays in the FC fabric (endpoint device, FC switch connection to endpoint devices, interswitch link).</p> <p>See also <i>node port (N_Port)</i>, <i>virtual node port (VN_Port)</i>, <i>proxy node port (NP_Port)</i>, <i>fabric port (F_Port)</i>, and <i>virtual fabric port (VF_Port)</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| port worldwide name (PWWN)        | WWN that is unique worldwide and is assigned to an FC port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| priority-based flow control (PFC) | <p>Link-level flow control mechanism defined by IEEE 802.1Qbb that allows independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.</p> <p>PFC is an enhancement of the Ethernet PAUSE mechanism, but PFC controls classes of flows, whereas Ethernet PAUSE indiscriminately pauses all of the traffic on a link. With PFC, a receiving device can signal a transmitting device to pause transmission based on traffic class.</p> <p>PFC provides application-specific bandwidth reservations so you can ensure that time-critical protocols and applications such as FCoE receive the priority necessary to prevent frame loss. PFC allows the same physical link to carry FCoE traffic and provide lossless service while also carrying loss-tolerant Ethernet traffic.</p> <p>See also <i>Ethernet PAUSE</i>.</p> |
| proxy gateway mode                | Connects FCoE initiators to FC switches in a converged Ethernet and Fibre Channel network and acts as an intermediary for these devices. The FCoE-FC gateway represents and acts for the FCoE initiators in transactions from the FCoE initiators destined for an FC switch, including converting FIP and FCoE frames to FC frames. The gateway represents and acts for an FC switch in transactions from the FC switch destined for an FCoE initiator, including converting FC frames to FIP frames and encapsulating FC frames in Ethernet.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| proxy node port (NP_Port)         | N_Port on the QFX Series that performs proxy functions when it is configured as an FCoE-FC gateway. The NP_Port acts as a proxy for the FCoE device VN_Ports in transactions with the FC switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 397: Fibre Channel Terms (*continued*)

| Term                                    | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| quantized congestion notification (QCN) | Mechanism defined by IEEE 802.1Qau that manages network congestion within a Layer 2 domain. When a queue reaches a configured threshold, QCN throttles traffic at the source of the congestion by transmitting messages that propagate back to the source and temporarily stop the source from transmitting. When the queue crosses the threshold that indicates the congestion has dissipated, QCN sends a message to allow the source to resume transmitting frames. |
| session                                 | Fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.                                                                                                                                                                                                                                                                                                                          |
| server-provided MAC address (SPMA)      | <p>MAC address that an ENode assigns to one of its ENode MACs and is not assigned to any other ENode MAC in the same FCoE VLAN. An SPMA can be associated with more than one VN_Port at that ENode MAC.</p> <p>The QFX Series does not support SPMA.</p> <p>See also <i>ENode MAC</i> and <i>fabric-provided MAC address (FPMA)</i>.</p>                                                                                                                               |
| storage area network (SAN)              | Network whose primary purpose is the transfer of data between computer systems and storage devices. This term is most commonly used in the context of any network that supports block storage, usually iSCSI, FC, and FCoE networks.                                                                                                                                                                                                                                   |
| target                                  | System component that receives an I/O command. An FC storage device that receives a request from a server is an example of a target.                                                                                                                                                                                                                                                                                                                                   |
| VE_Port                                 | Virtual ports created to form a connection (an <i>interswitch link</i> ) between two FCoE-based FC switches as part of a common FC fabric.                                                                                                                                                                                                                                                                                                                             |
| VE2VE (VE_Port to VE_Port)              | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of FCFs to connect to each other as a single FCoE FC SAN.                                                                                                                                                                                                                                                                                                                            |
| VN2VF (VN_Port to VF_Port)              | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of an ENode to connect to an FCF or to an FCoE-enabled FC SAN.                                                                                                                                                                                                                                                                                                                       |
| VN2VN (VN_Port to VN_Port)              | The <i>Fibre Channel Backbone - 6 (FC-BB-6)</i> specification capability of an ENode to connect directly over Layer 2 to another ENode without the need of any FC-related services. This capability is most often used in small-scale FCoE SANs.                                                                                                                                                                                                                       |
| virtual fabric port (VF_Port)           | <p>Data-forwarding component that emulates an F_Port. A VF_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VN_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>See also <i>fabric port (F_Port)</i>.</p>                                                                                                                                                   |
| virtual link                            | <p>Logical link connecting two FCoE Link End Points (LEPs) over a lossless Ethernet network, for example, the link between a VF_Port and a VN_Port. The MAC addresses of the two LEPs identifies a virtual link.</p> <p>See also <i>FCoE link end point (LEP)</i> and <i>lossless Ethernet network</i>.</p>                                                                                                                                                            |

Table 397: Fibre Channel Terms (*continued*)

| Term                        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| virtual node port (VN_Port) | <p>Data-forwarding component that emulates an N_Port. With FCoE, a VN_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VF_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>VN_Port is also used for the virtual N_Ports created in both FC and FCoE when additional NPIV-based logins occur over a previously created N_Port-to-VN_Port or N_Port-to-VF_Port connection.</p> <p>See also <i>node port</i> (N_Port).</p> |
| well-known address (WKA)    | Address identifier used to access a service provided by an FC fabric. The service can be distributed in many elements throughout a fabric, or it can be centralized in one element. A WKA is always accessible, regardless of zoning. An example of a WKA is the <i>ALL-FCF-MACs</i> address to which all FCFs listen.                                                                                                                                                                                                             |
| worldwide name (WWN)        | 64-bit identifier that is similar to a MAC address except that it is not used for forwarding. It uniquely identifies an FC device. The WWN is derived from the IEEE organizationally unique identifier (OUI) and vendor-supplied information. A WWN is unique worldwide.                                                                                                                                                                                                                                                           |
| worldwide node name (WWNN)  | See <i>node worldwide name</i> (NWWN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| worldwide port name (WWPN)  | See <i>port worldwide name</i> (PWWN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

- Related Documentation**
- [Overview of Fibre Channel on the QFX Series on page 4956](#)
  - [Understanding QFabric System Terminology on page 1169](#)

## QFabric Specific

- [Understanding Fibre Channel Fabrics on the QFabric System on page 5084](#)
- [Understanding CoS Fabric Forwarding Class Sets on page 5086](#)

### Understanding Fibre Channel Fabrics on the QFabric System

A Fibre Channel (FC) fabric on a QFabric system is a construct that you configure on a QFX3500 Node device when the Node device is in FCoE-FC gateway mode. The FC fabric on a QFabric Node device is not the same as an FC fabric on a storage area network (SAN). The FC fabric on a QFabric Node device is local to that particular node device. We call the FC fabric on a QFabric Node device a *local FC fabric* to differentiate it from an FC fabric on the SAN.



**NOTE:** The QFX3600 Node device does not support FC or FCoE features.

A local FC fabric does not span Node devices and does not span the fabric Interconnect device. Local FC fabrics are entirely contained on a single Node device. A local FC fabric



creates associations that connect FCoE devices that have converged network adapters (CNAs) on the Ethernet network to an FC switch or FCoE forwarder (FCF) on the FC network. A local FC fabric consists of:

- A unique fabric name.
- A unique fabric ID.
- One or more FCoE VLAN interfaces that include one or more 10-Gigabit Ethernet interfaces connected to FCoE devices. The FCoE VLANs transport traffic between the FCoE servers and the FCoE-FC gateway. Each FCoE VLAN must carry only FCoE traffic. You cannot mix FCoE traffic and standard Ethernet traffic on the same VLAN.

The 10-Gigabit Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic because FIP VLAN discovery and notification frames are exchanged as untagged packets.

Each FCoE VLAN interface can present multiple VF\_Port interfaces to the FCoE network.

- One or more native FC interfaces. The native FC interfaces transport traffic between the gateway and the FC switch or FCF.



**TIP:** If the network does not use a dual-rail architecture for redundancy, configure more than one native FC interface for each local FC fabric to create redundant connections between the FCoE devices and the FC network. If one physical link goes down, any sessions it carried can log in again and connect to the FC network on a different interface.

All of the FC and FCoE traffic that belongs to a local FC fabric on a Node device must enter and exit that Node device. This means that the FC switch or FCF and the FCoE devices in the Ethernet network must be connected to the same Node device. The interfaces that connect to the FC switch and the interfaces that connect to the FCoE devices must be included in the local FC fabric. You cannot configure a local FC fabric that spans more than one Node device.

Traffic flows from FC and FCoE devices that are not in the same local FC fabric remain separate and cannot communicate with each other through the FCoE-FC gateway.



**NOTE:** The QFabric system enforces commit checks to ensure that local FC fabrics and FCoE VLANs on FCoE-FC gateways do not span more than one Node device.

#### Related Documentation

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

## Understanding CoS Fabric Forwarding Class Sets

Fabric forwarding class sets (fabric fc-sets) are similar to the fc-sets (priority groups) you configure on Node devices. The major differences are:

1. Fabric fc-sets group traffic for transport across the QFX3008-I or QFX3600-I Interconnect device (the fabric). Node device fc-sets group traffic on a Node device for transport across that Node device.
2. Fabric fc-sets are global. They apply to the entire fabric. Node device fc-sets apply only to the Node device on which they are configured.
3. You can configure class of service (CoS) scheduling for Node device fc-sets, but you cannot configure CoS for fabric fc-sets.
4. Fabric fc-sets map to Interconnect device fabric output queues statically—you cannot configure the mapping of fabric fc-sets to fabric output queues. All traffic in a fabric fc-set maps to the same output queue.

Node device fc-sets include forwarding classes that map to Node device output queues, and you can configure the mapping of forwarding classes to output queues (or you can use the default mapping). Because output queues are mapped to forwarding classes, different classes of traffic in a Node device fc-set can be mapped to different output queues.

Node device fc-sets consist of forwarding classes containing traffic that requires similar CoS treatment. (Forwarding classes are default forwarding classes or user-defined forwarding classes.) You can configure CoS for each fc-set to determine how the traffic of its forwarding classes is scheduled on a Node device.

When traffic exits a Node device interface and enters an Interconnect device fabric interface, the Interconnect device uses the same forwarding classes to group traffic. The forwarding classes are mapped to global fabric fc-sets for transport across the fabric. Like fc-sets on a Node device, fabric fc-sets also contain traffic that requires similar CoS treatment. Unlike fc-sets on a Node device, you cannot configure CoS on fabric fc-sets.

Fabric fc-sets reside on the Interconnect device and are global to the QFabric system. Fabric fc-sets apply to all traffic that traverses the fabric. The mapping of forwarding classes to fabric fc-sets is global and applies to all forwarding classes with traffic that traverses the fabric from all connected Node devices. You can change the mapping of forwarding classes to fabric fc-sets. All mapping changes you make are global. For example, if you change the fabric fc-set to forwarding class mapping of the default best-effort forwarding class, then every Node device's best-effort forwarding class traffic that traverses the fabric is mapped to that fabric fc-set.

This topic describes:

- [Default Fabric Forwarding Class Sets on page 5087](#)
- [Fabric Forwarding Class Set Configuration and Implementation on page 5090](#)
- [Fabric Forwarding Class Set Scheduling \(CoS\) on page 5092](#)
- [Support for Flow Control and Lossless Transport Across the Fabric on page 5094](#)

- [Viewing Fabric Forwarding Class Set Information on page 5096](#)
- [Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences on page 5098](#)

### Default Fabric Forwarding Class Sets

Interconnect devices have 12 default fabric fc-sets, including five visible default fabric fc-sets, four for unicast traffic and one for multdestination (multicast, broadcast, and destination lookup failure) traffic.

There are also seven hidden default fabric fc-sets. There are three hidden default fabric fc-sets for multdestination traffic that you can use if you want to map different multdestination forwarding classes to different multdestination fabric fc-sets. There are four hidden default fabric fc-sets for lossless traffic that you can use to map different lossless forwarding classes (priorities) to different lossless fabric fc-sets.

[Table 117 on page 1239](#) shows the default fabric fc-sets:

**Table 398: Default Fabric Forwarding Class Sets**

| Fabric Forwarding Class Set Name      | Characteristics                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>fabric_fcset_be</code>          | Transports best-effort unicast traffic across the fabric.                                                                                                                                                                                                                                                                                                                                                                             |
| <code>fabric_fcset_strict_high</code> | Transports unicast traffic that has been configured with <b>strict-high</b> priority and in the <b>network-control</b> forwarding class across the fabric. This fabric fc-set receives as much bandwidth across the fabric as it needs to service the traffic in the group up to the entire fabric interface bandwidth. For this reason, exercise caution when mapping traffic to this fabric fc-set to avoid starving other traffic. |
| <code>fabric_fcset_noloss1</code>     | Transports unicast traffic in the default <b>fcoe</b> forwarding class across the fabric.                                                                                                                                                                                                                                                                                                                                             |
| <code>fabric_fcset_noloss2</code>     | Transports unicast traffic in the default <b>no-loss</b> forwarding class across the fabric.                                                                                                                                                                                                                                                                                                                                          |
| <code>fabric_fcset_noloss3</code>     | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.                                                                                                                                                                                                                |
| <code>fabric_fcset_noloss4</code>     | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.                                                                                                                                                                                                                |
| <code>fabric_fcset_noloss5</code>     | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.                                                                                                                                                                                                                |

Table 398: Default Fabric Forwarding Class Sets (*continued*)

| Fabric Forwarding Class Set Name | Characteristics                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>fabric_fcset_noloss6</b>      | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.        |
| <b>fabric_fcset_multicast1</b>   | Transports multdestination traffic in the <b>mcast</b> forwarding class across the fabric. This fabric fc-set is valid only for multdestination forwarding classes.                                                           |
| <b>fabric_fcset_multicast2</b>   | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes. |
| <b>fabric_fcset_multicast3</b>   | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes. |
| <b>fabric_fcset_multicast4</b>   | (Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes. |

The five default forwarding classes (**best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast**) are mapped to the fabric fc-sets by default as shown in [Table 118 on page 1240](#).

Table 399: Default Forwarding Class to Fabric Forwarding Class Set Mapping

| Forwarding Class                                                          | Fabric Forwarding Class Set | Fabric Output Queue | Maximum MTU Supported for Lossless Operation |
|---------------------------------------------------------------------------|-----------------------------|---------------------|----------------------------------------------|
| <b>best-effort</b>                                                        | fabric_fcset_be             | 0                   | NA                                           |
| <b>network-control</b>                                                    | fabric_fcset_strict_high    | 7                   | NA                                           |
| <b>fcoe</b>                                                               | fabric_fcset_noloss1        | 1                   | 9K                                           |
| <b>no-loss</b>                                                            | fabric_fcset_noloss2        | 2                   | 9K                                           |
| <b>mcast</b>                                                              | fabric_fcset_multicast1     | 8                   | NA                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_noloss3        | 3                   | 9k                                           |

**Table 399: Default Forwarding Class to Fabric Forwarding Class Set Mapping (*continued*)**

| Forwarding Class                                                          | Fabric Forwarding Class Set | Fabric Output Queue | Maximum MTU Supported for Lossless Operation |
|---------------------------------------------------------------------------|-----------------------------|---------------------|----------------------------------------------|
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_noloss4        | 4                   | 9k                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_noloss5        | 5                   | 9k                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_noloss6        | 6                   | 9k                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_multicast2     | 9                   | NA                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_multicast3     | 10                  | NA                                           |
| No forwarding classes are mapped by default to this hidden fabric fc-set. | fabric_fcset_multicast4     | 11                  | NA                                           |

The maximum fiber cable length between the QFabric system Node device and the QFabric system Interconnect device is 150 meters.



**TIP:** If you explicitly configure lossless forwarding classes, we recommend that you map each user-configured lossless forwarding class to an unused fabric fc-set (fabric\_fcset\_noloss3 through fabric\_fcset\_noloss6) on a one-to-one basis: one lossless forwarding class mapped to one lossless fabric fc-set.

The reason for one-to-one mapping is to avoid fate sharing of lossless flows. Because each fabric fc-set is mapped statically to an output queue, when you map more than one forwarding class to a fabric fc-set, all of the traffic in all of the forwarding classes that belong to the fabric fc-set uses the same output queue. If that output queue becomes congested due to congestion caused by one of the flows, the other flows are also affected. (They share fate because the flow that congests the output queue affects flows that are not experiencing congestion.)

However, it is important to understand that fabric\_fcset\_noloss1 and fabric\_fcset\_noloss2 have a scheduling weight of 35, while the other fabric fc-sets have a scheduling weight of 1. The scheduling weights mean that fabric\_fcset\_noloss1 and fabric\_fcset\_noloss2 receive most of the bandwidth available to lossless fabric fc-sets if the amount of traffic on fabric\_fcset\_noloss1 and fabric\_fcset\_noloss2 requires the bandwidth.

If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will consume most of that bandwidth, then you should place all lossless traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2`. If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will **<emphasis>not</emphasis>** consume most of that bandwidth, then you can map lossless forwarding classes in a one-to-one manner to lossless fabric fc-sets to avoid fate sharing.

---

If you want to map different multidestination forwarding classes to different multidestination fabric fc-sets, use one or more of the hidden multidestination fabric fc-sets.



**NOTE:** The global mapping of forwarding classes to fabric fc-sets is independent of the mapping of forwarding classes to Node device fc-sets. Global mapping of forwarding classes to fabric fc-sets occurs only on the Interconnect device. The Node device mapping of forwarding classes to fc-sets does not affect the global mapping of forwarding classes to fabric fc-sets on the Interconnect device, and vice versa.

---

When you define new forwarding classes on a Node device, you explicitly map those forwarding classes to Node device fc-sets. However, new (user-created) forwarding classes are mapped by default to fabric fc-sets. (You can override the default mapping if you want to configure the forwarding class to fabric fc-set mapping explicitly, as described in the next section.)

By default:

- All best-effort traffic forwarding classes that you create are mapped to the **`fabric_fcset_be`** fabric fc-set.
- All lossless traffic forwarding classes that you create are mapped to the **`fabric_fcset_noloss1`** or **`fabric_fcset_noloss2`** fabric fc-set.
- All multidestination traffic forwarding classes that you create are mapped to the **`fabric_fcset_multicast1`** fabric fc-set.
- All **strict-high** priority traffic and **network-control** forwarding classes that you create are mapped to the **`fabric_fcset_strict_high`** fabric fc-set.

---

### Fabric Forwarding Class Set Configuration and Implementation

You can map forwarding classes to fabric fc-sets, but no other attributes of fabric fc-sets are user-configurable, including CoS. This section describes:

- [Mapping Forwarding Classes to Fabric Forwarding Class Sets on page 5091](#)
- [Fabric Forwarding Class Set Implementation on page 5091](#)

### Mapping Forwarding Classes to Fabric Forwarding Class Sets

If you do not want to use the default mapping of forwarding classes to fabric fc-sets, you can map forwarding classes to fabric fc-sets the same way as you map forwarding classes to Node device fc-sets. To do this, use exactly the same statement that you use to map forwarding classes to fc-sets, but instead of specifying a Node device fc-set name, specify a fabric fc-set name.



**NOTE:** The global mapping of forwarding classes to fabric fc-sets does not affect the mapping of forwarding classes to Node device fc-sets. The global forwarding class mapping to fabric fc-sets pertains to the traffic only when it enters, traverses, and exits the fabric. The forwarding class mapping to fc-sets on a Node device is valid within that Node device.

Mapping forwarding classes to fabric fc-sets does not affect the scheduling configuration of the forwarding classes or fc-sets on Node devices. Fabric fc-set scheduling (which is not user-configurable) pertains to traffic only when it enters, traverses, and exits the Interconnect device fabric.

If you change the mapping of a forwarding class to a fabric fc-set, the new mapping is global and applies to all traffic in that forwarding class, regardless of which Node device forwards the traffic to the Interconnect device.

- To assign one or more forwarding classes to a fabric fc-set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric-forwarding-class-set-name class
forwarding-class-name
```

For example, to map a user-defined forwarding class named **best-effort-2** to the fabric fc-set **fabric\_fcset\_be**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric_fcset_be class best-effort-2
```



**NOTE:** Because fabric fc-set configuration is global, in this example all forwarding classes with the name **best-effort-2** on all of the Node devices attached to the fabric use the **fabric\_fcset\_be** fabric fc-set to transport traffic across the fabric.

### Fabric Forwarding Class Set Implementation

The following rules apply to fabric fc-sets:

- You cannot create new fabric fc-sets. Only the twelve default fabric fc-sets are available.
- You cannot delete a default fabric fc-set.
- You cannot attach a fabric fc-set to a Node device interface. Fabric fc-sets are used only on the Interconnect device fabric, not on Node devices.

- You can map only multdestination forwarding classes to multdestination fabric fc-sets.
- You cannot map multdestination forwarding classes to unicast fabric fc-sets.
- You cannot map unicast forwarding classes to multdestination fabric fc-sets.
- You cannot configure CoS for fabric fc-sets. (However, default CoS scheduling properties are applied to traffic on the fabric, and the fabric interfaces use link layer flow control (LLFC) for flow control.)

### **Fabric Forwarding Class Set Scheduling (CoS)**

---

Although fabric fc-set CoS is not user-configurable, CoS is applied to traffic on the fabric. (In addition, fabric interfaces use LLFC to ensure lossless transport for lossless traffic flows.) This section describes how the fabric applies CoS scheduling to traffic:

- [Class Groups for Fabric Forwarding Class Sets on page 5092](#)
- [Class Group Scheduling on page 5092](#)
- [QFabric System CoS on page 5094](#)

#### ***Class Groups for Fabric Forwarding Class Sets***

To transport traffic across the fabric, the QFabric system organizes the fabric fc-sets into three classes called *class groups*. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric\_fcset\_strict\_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes and in any forwarding classes you create on a Node device that consist of **strict-high** priority or **network-control** forwarding class traffic.
- **Unicast**—All traffic in the fabric fc-sets **fabric\_fcset\_be**, **fabric\_fcset\_noloss1**, and **fabric\_fcset\_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes and in any forwarding classes you create on a Node device that consist of best-effort or lossless traffic. If you use any of the hidden no loss fabric fc-sets (**fabric\_fcset\_noloss3**, **fabric\_fcset\_noloss4**, **fabric\_fcset\_noloss5**, or **fabric\_fcset\_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric\_fcset\_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric\_fcset\_multicast2**, **fabric\_fcset\_multicast3**, or **fabric\_fcset\_multicast4**), that traffic is also classified as part of this class group.

#### ***Class Group Scheduling***

You cannot configure CoS for class groups or for fabric fc-sets (that is, you cannot attach a traffic control profile to a fabric fc-set—you attach traffic control profiles to Node device fc-sets to apply scheduling to the traffic that belongs to the Node device fc-set). By default, the fabric uses weighted round-robin (WRR) scheduling in which each class group receives a portion of the total available fabric bandwidth based on its type of traffic, as shown in [Table 119 on page 1245](#):



Table 400: Class Group Scheduling Properties and Membership

| Class Group          | Fabric fc-sets                                                                                                                                                                                                                                                                                                                                                                                              | Forwarding Classes (Default Mapping)                                                                                                                                                                                    | Class Group Scheduling Properties (Weight)                                                                                                                                                                                                                                                                                                                              |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Strict-high priority | <b>fabric_fcset_strict_high</b>                                                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>All <b>strict-high</b> priority forwarding classes</li> <li><b>network-control</b></li> </ul>                                                                                    | Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic. |
| Unicast              | <ul style="list-style-type: none"> <li><b>fabric_fcset_be</b></li> <li><b>fabric_fcset_noloss1</b></li> <li><b>fabric_fcset_noloss2</b></li> </ul> <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> <li><b>fabric_fcset_noloss3</b></li> <li><b>fabric_fcset_noloss4</b></li> <li><b>fabric_fcset_noloss5</b></li> <li><b>fabric_fcset_noloss6</b></li> </ul> | <ul style="list-style-type: none"> <li><b>best-effort</b></li> <li><b>fcoe</b></li> <li><b>no-loss</b></li> </ul> <p><b>NOTE:</b> No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p> | Traffic in the unicast class group receives an 80% weight in the WRR calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)                                                                |
| Multidestination     | <b>fabric_fcset_multicast1</b> <p>Includes the hidden multidestination fabric fc-sets if used:</p> <ul style="list-style-type: none"> <li><b>fabric_fcset_multicast2</b></li> <li><b>fabric_fcset_multicast3</b></li> <li><b>fabric_fcset_multicast4</b></li> </ul>                                                                                                                                         | <ul style="list-style-type: none"> <li><b>mcast</b></li> </ul> <p><b>NOTE:</b> No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>                                            | Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)                                      |

The fabric fc-sets within each class group are weighted equally and receive bandwidth using round-robin scheduling. For example:

- If the unicast class group has three member fabric fc-sets, **fabric\_fcset\_be**, **fabric\_fcset\_noloss1**, and **fabric\_fcset\_noloss2**, then each of the three fabric fc-sets receives one-third of the bandwidth available to the unicast class group.
- If the multidestination class group has one member fc-set, **fabric\_fcset\_multicast1**, then that fc-set receives all of the bandwidth available to the multidestination class group.
- If the multidestination class group has two member fc-sets, **fabric\_fcset\_multicast1** and **fabric\_fcset\_multicast2**, then each of the two fabric fc-sets receives one-half of the bandwidth available to the multidestination class group.

### ***QFabric System CoS***

When traffic enters and exits the same Node device, CoS works the same as it works on a standalone QFX3500 switch.

However, when traffic enters a Node device, crosses the Interconnect device, and then exits a different Node device, CoS is applied differently:

1. Traffic entering the ingress Node device receives the CoS configured at the Node ingress (packet classification, congestion notification profile for PFC).
2. When traffic goes from the ingress Node device to the Interconnect device, the fabric fc-set CoS is applied as described in the discussion of fabric forwarding class set scheduling.
3. When traffic goes from the Interconnect device to the egress Node device, the egress Node device applies CoS at the egress port (egress queue scheduling, WRED, IEEE 802.1p or DSCP code-point rewrite).

### **Support for Flow Control and Lossless Transport Across the Fabric**

---

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the fabric interfaces use LLFC to support lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class on the Node devices.
2. The lossless forwarding class must be mapped to a lossless fabric fc-set on the Interconnect device (**fabric\_fcset\_noloss1**, **fabric\_fcset\_noloss2**, **fabric\_fcset\_noloss3**, **fabric\_fcset\_noloss4**, **fabric\_fcset\_noloss5**, or **fabric\_fcset\_noloss6**).

When traffic meets the two configuration constraints, the fabric propagates the back pressure from the egress Node device across the fabric to the ingress Node device during periods of congestion. However, to achieve end-to-end lossless transport across the switch, you must also configure a congestion notification profile to enable PFC on the Node device ingress ports.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric\_fcset\_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric\_fcset\_be** fabric fc-set, then the congestion control mechanism is packet drop.

- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric\_fcset\_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric\_fcset\_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric\_fcset\_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric\_fcset\_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



**NOTE:** Lossless transport across the fabric also must meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is a 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- **MTU size**—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- **Total number of Node device fabric ports connected to the Interconnect device**—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multideestination traffic.

Table 120 on page 1248 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 120 on page 1248, the QFabric system supports six lossless priorities.



**NOTE:** The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

**Table 401: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported**

| MTU in Bytes | Fiber Cable Length in Meters (Approximate) | Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device |
|--------------|--------------------------------------------|-------------------------------------------------------------------------------|
| 9216 (9K)    | 100                                        | 5                                                                             |
| 9216 (9K)    | 150                                        | 5                                                                             |



**NOTE:** The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

### Viewing Fabric Forwarding Class Set Information

You can display information about fabric fc-sets using the same CLI command you use to display information about Node device fc-sets:

```
user@switch> show class-of-service forwarding-class-set
Forwarding class set: fabric_fcset_be, Type: fabric-type, Forwarding class set
index: 1
 Forwarding class Index
 best-effort 0

Forwarding class set: fabric_fcset_mcast1, Type: fabric-type, Forwarding class
set index: 5
 Forwarding class Index
 mcast 8
```

Forwarding class set: fabric\_fcset\_mcast2, Type: fabric-type, Forwarding class set index: 6

Forwarding class set: fabric\_fcset\_mcast3, Type: fabric-type, Forwarding class set index: 7

Forwarding class set: fabric\_fcset\_mcast4, Type: fabric-type, Forwarding class set index: 8

Forwarding class set: fabric\_fcset\_noloss1, Type: fabric-type, Forwarding class set index: 2

|                  |       |
|------------------|-------|
| Forwarding class | Index |
| fcoe             | 1     |

Forwarding class set: fabric\_fcset\_noloss2, Type: fabric-type, Forwarding class set index: 3

|                  |       |
|------------------|-------|
| Forwarding class | Index |
| no-loss          | 2     |

Forwarding class set: fabric\_fcset\_noloss3, Type: fabric-type, Forwarding class set index: 9

Forwarding class set: fabric\_fcset\_noloss4, Type: fabric-type, Forwarding class set index: 10

Forwarding class set: fabric\_fcset\_noloss5, Type: fabric-type, Forwarding class set index: 11

Forwarding class set: fabric\_fcset\_noloss6, Type: fabric-type, Forwarding class set index: 12

Forwarding class set: fabric\_fcset\_strict\_high, Type: fabric-type, Forwarding class set index: 4

|                  |       |
|------------------|-------|
| Forwarding class | Index |
| network-control  | 3     |

[Table 121 on page 1249](#) describes the meaning of the **show class-of-service forwarding-class-set** output fields when you display fabric fc-set information.

**Table 402: show class-of-service forwarding-class-set Command Output Fields**

| Field Name                 | Field Description                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class set       | Name of the fabric forwarding class set.                                                                                                          |
| Type                       | Type of forwarding class set: <ul style="list-style-type: none"> <li>Fabric-type—Fabric fc-set</li> <li>Normal-type—Node device fc-set</li> </ul> |
| Forwarding class set index | Index of this forwarding class set.                                                                                                               |
| Forwarding class           | Name of a forwarding class.                                                                                                                       |
| Index                      | Index of the forwarding class.                                                                                                                    |

## Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences

Table 122 on page 1250 summarizes the differences between fabric fc-sets and Node device fc-sets:

**Table 403: Summary of Differences Between Fabric fc-sets and Node Device fc-sets**

| Characteristic                                 | Fabric fc-set                                                                                                                             | Node device fc-set                                                                                    |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Location                                       | QFX3008-I or QFX3600-I Interconnect device (the fabric).                                                                                  | QFabric system Node device.                                                                           |
| Global or Node-device specific                 | Global, valid for the entire fabric.                                                                                                      | Local to the Node device on which the fc-set is configured.                                           |
| Ability to create (define) a new fc-set        | No. Use the 12 default fabric fc-sets provided.                                                                                           | Yes.                                                                                                  |
| Ability to configure CoS                       | Default CoS settings only. CoS is not user-configurable.                                                                                  | User-configurable using traffic control profiles.                                                     |
| Ability to map forwarding classes to an fc-set | Yes. Mapping is global and applies to all forwarding classes across Interconnect device fabric (traffic from all connected Node devices). | Yes. Mapping is local to a Node device and applies only to the forwarding classes on the Node device. |

### Related Documentation

- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [show class-of-service forwarding-class-set on page 5938](#)

## CHAPTER 60

# Configuration

- [Configuration Examples on page 5099](#)
- [Configuration Tasks on page 5178](#)
- [Configuration Statements on page 5210](#)

### Configuration Examples

---

- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5121](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5155](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5167](#)
- [Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175](#)

### Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

To transmit Fibre Channel (FC) traffic between FCoE devices and a storage area network (SAN) FC switch, you configure a local FC fabric on the gateway. The gateway FC fabric includes FCoE and native FC interfaces, and a VLAN to carry FCoE traffic from FCoE-capable devices. The gateway FC fabric creates the path between the FCoE devices and the SAN.

This example describes how to configure the interfaces, VLAN, and FC fabric to connect FCoE devices to the FC switch and route traffic between the VLAN and FC interfaces:

- [Requirements on page 5100](#)
- [Overview on page 5100](#)
- [Configuration on page 5103](#)
- [Verification on page 5110](#)

## Requirements

---

This example uses the following hardware and software components:

- A configured and provisioned Juniper Networks QFX3500 Switch to act as an FCoE-FC gateway
- FCoE-capable devices in an Ethernet network equipped with converged network adapters (CNAs)
- An FC switch to transmit and receive native FC traffic
- FC storage devices in the SAN
- Junos OS Release 11.1 or later for the QFX Series

## Overview

---

No interfaces are configured for FC network connectivity by default. You need to configure the FC fabric and its interfaces explicitly. Each FC fabric consists of a combination of at least one FCoE VLAN interface between the FCoE-FC gateway and the FCoE devices, and one or more native FC interfaces between the FCoE-FC gateway and the FC switch.

An FCoE VLAN interface connects the FCoE-FC gateway to FCoE devices. FCoE traffic between the devices and the FCoE-FC gateway requires a dedicated VLAN used only for FCoE traffic. You cannot mix standard Ethernet traffic and FCoE traffic on the FCoE VLAN.



**NOTE:** IGMP snooping is not supported on FCoE VLANs. Disable IGMP snooping on FCoE VLANs.

Storm control is not supported on Ethernet interfaces that belong to the FCoE VLAN. Ensure that storm control is disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped.

When FCoE frames enter the FCoE-FC gateway, the gateway:

1. Strips the Ethernet encapsulation from the FCoE frames.
2. Sends the resulting native FC frames to the FC switch through the gateway's native FC interfaces.

Each FC interface and FCoE VLAN interface can belong to only one FC fabric. Different FC fabrics must use different native FC interfaces and different FCoE VLAN interfaces. Multiple FC fabrics on the FCoE-FC gateway can connect to the same FC switch, but they must use different FC interfaces and different FCoE VLAN interfaces.

The Ethernet interfaces that belong to the FCoE VLAN should be configured in tagged-access port mode and must include the native VLAN because FIP VLAN discovery and notification frames are exchanged as untagged packets. These Ethernet interfaces require a maximum transmission unit (MTU) size of at least 2180 bytes to accommodate the FC payload and FCoE encapsulation. (Sometimes the MTU is rounded up to 2500 bytes. If larger frames are expected on the interface, set the MTU size accordingly.)



This example shows a simple configuration to illustrate the basic steps for creating:

- The FCoE-device-facing VLAN and its 10-Gigabit Ethernet interfaces
- The VLAN interface
- The FC-switch-facing native FC interfaces
- One FC fabric on the FCoE-FC gateway

Configuring these elements results in traffic being routed between the VLAN and FC interfaces, thus connecting the FCoE devices to the FC switch through the FCoE-FC gateway.

A VLAN called **blue** transports FCoE traffic between FCoE devices and the FCoE-FC gateway using an FCoE VLAN interface called **vlan.100**. The FCoE-FC gateway's **vlan.100** interface presents an F\_Port interface to the FCoE devices on the VLAN. For each FCoE device ENode that logs in to the FCoE-FC gateway, the gateway instantiates a virtual F\_Port (VF\_Port) interface. This creates a virtual link between the ENode VN\_Port and the FCoE-FC gateway. The FCoE-FC gateway's native FC interfaces transport FC traffic between the gateway and the FC switch.

Configuring both the FCoE VLAN interface and the native FC interfaces as part of a gateway fabric associates them in the switch and makes the connection between the FCoE servers and the FC switch.

**Topology**

The topology for this example consists of one QFX3500 switch with FC-capable ports to connect to the FC switch and with Ethernet ports in tagged-access mode to connect to the FCoE devices. [Table 404 on page 5101](#) and [Figure 195 on page 5102](#) show the configuration components of this example.


**Table 404: Components of the Fibre Channel Interface Configuration Topology**

| Property                       | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware                | QFX3500 switch in gateway mode                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| FCoE VLAN name and tag ID      | <b>blue</b> , tag <b>100</b><br><br>IGMP snooping disabled on the FCoE VLAN.                                                                                                                                                                                                                                                                                                                                                                                                           |
| Interfaces in VLAN <b>blue</b> | Interfaces: <b>xe-0/0/6</b> , <b>xe-0/0/7</b> , <b>xe-0/0/8</b> , <b>xe-0/0/9</b> , <b>xe-0/0/10</b> , <b>xe-0/0/11</b><br>Port mode: <b>tagged-access</b><br>MTU: <b>2180</b><br>Native VLAN: 1<br><br><b>NOTE:</b> You can bundle two or more of the VLAN interfaces in a link aggregation group (LAG) if you wish.<br><br><b>NOTE:</b> FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features. |

Table 404: Components of the Fibre Channel Interface Configuration Topology (*continued*)

| Property                             | Settings                                                                                                                             |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| FCoE VLAN interface                  | vlan.100<br>Port mode: <b>f-port</b>                                                                                                 |
| Native Fibre Channel interfaces      | Interfaces: <b>fc-0/0/0, fc-0/0/1, fc-0/0/2, fc-0/0/3, fc-0/0/4, fc-0/0/5</b><br>Port mode: <b>np-port</b><br>Speed: <b>4 Gbps</b>   |
| Fibre Channel fabric <b>fcproxy1</b> | Fabric type: <b>proxy</b><br>Fabric ID: <b>1</b><br>FC interfaces: <b>fc-0/0/0, fc-0/0/1, fc-0/0/2, fc-0/0/3, fc-0/0/4, fc-0/0/5</b> |

Figure 195: Fibre Channel Interface Configuration Topology



This configuration example creates a VLAN for FCoE traffic and routes its traffic to an FCoE VLAN interface that is part of the FC fabric. It also creates the FC interfaces needed to connect to the FC switch.

To set up FC interfaces and FCoE VLAN interfaces:

- Configure a VLAN to use as a dedicated FCoE VLAN:
  - Configure the interfaces the FCoE VLAN uses as Ethernet switching interfaces in tagged-access port mode.
  - If storm control is enabled, disable it on the interfaces.
  - Configure the interfaces the FCoE VLAN uses with the native VLAN.

- Configure the FCoE VLAN to use the desired Ethernet interfaces.
- Disable IGMP snooping on the FCoE VLAN. (IGMP snooping is enabled by default on all VLANs, but is not supported on FCoE VLANs).
- Configure the FCoE VLAN interface.
- Define the interface for the FCoE VLAN (associate the VLAN with the FCoE VLAN interface).
- Configure the physical FC interfaces (either one or two 6-port blocks) that connect to the FC switch.
- Configure the logical FC interfaces that connect to the FC switch.
- Configure the FCoE-FC gateway fabric:
  - Configure the fabric ID.
  - Configure the fabric as a proxy fabric.
  - Add the FCoE VLAN interface and the native FC interfaces to the fabric.

To keep the example simple, the configuration steps show six Ethernet interfaces in the FCoE VLAN and six native FC interfaces in the FC fabric. Use the same configuration procedure to add more interfaces to the FCoE VLAN or to the FC fabric.

### Configuration

#### CLI Quick Configuration

To quickly configure FCoE and native FC interfaces on an FCoE-FC gateway and route traffic between the FCoE VLAN and FC interfaces, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans blue vlan-id 100
set vlans native vlan-id 1
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/7 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/8 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode tagged-access vlan members blue
set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/7 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/8 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/9 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/11 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/6 mtu 2180
set interfaces xe-0/0/7 mtu 2180
set interfaces xe-0/0/8 mtu 2180
set interfaces xe-0/0/9 mtu 2180
set interfaces xe-0/0/10 mtu 2180
```

```

set interfaces xe-0/0/11 mtu 2180
set vlans blue interface xe-0/0/6.0
set vlans blue interface xe-0/0/7.0
set vlans blue interface xe-0/0/8.0
set vlans blue interface xe-0/0/9.0
set vlans blue interface xe-0/0/10.0
set vlans blue interface xe-0/0/11.0
set protocols igmp-snooping vlan blue disable
set interfaces vlan unit 100 family fibre-channel port-mode f-port
set vlans blue l3-interface vlan.100
set chassis fpc 0 pic 0 fibre-channel port-range 0 5
set interfaces fc-0/0/0 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/1 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/2 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/3 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/4 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/5 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/0 fibrechannel-options speed 4g
set interfaces fc-0/0/1 fibrechannel-options speed 4g
set interfaces fc-0/0/2 fibrechannel-options speed 4g
set interfaces fc-0/0/3 fibrechannel-options speed 4g
set interfaces fc-0/0/4 fibrechannel-options speed 4g
set interfaces fc-0/0/5 fibrechannel-options speed 4g
set fc-fabrics fcproxy1 fabric-id 1
set fc-fabrics fcproxy1 fabric-type proxy
set fc-fabrics fcproxy1 interface vlan.100
set fc-fabrics fcproxy1 interface fc-0/0/0.0
set fc-fabrics fcproxy1 interface fc-0/0/1.0
set fc-fabrics fcproxy1 interface fc-0/0/2.0
set fc-fabrics fcproxy1 interface fc-0/0/3.0
set fc-fabrics fcproxy1 interface fc-0/0/4.0
set fc-fabrics fcproxy1 interface fc-0/0/5.0

```

**Step-by-Step Procedure** Configure FCoE and FC interfaces in an FCoE-FC gateway FC fabric and set up traffic routing between the FCoE VLAN and FC interfaces:

1. Configure the VLAN for FCoE traffic:
 

```

[edit vlans]
user@switch# set blue vlan-id 100

```
2. Configure the native VLAN:
 

```

[edit vlans]
user@switch# set native vlan-id 1

```
3. Configure the Ethernet interfaces for the FCoE VLAN in tagged-access mode and as members of the FCoE VLAN (VLAN blue):
 

```

[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue
user@switch# set xe-0/0/7 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue
user@switch# set xe-0/0/8 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue
user@switch# set xe-0/0/9 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue
user@switch# set xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue

```

```
user@switch# set xe-0/0/11 unit 0 family ethernet-switching port-mode tagged-access
vlan members blue
```

4. Configure the native VLAN on the Ethernet interfaces in the FCoE VLAN:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
user@switch# set xe-0/0/7 unit 0 family ethernet-switching native-vlan-id 1
user@switch# set xe-0/0/8 unit 0 family ethernet-switching native-vlan-id 1
user@switch# set xe-0/0/9 unit 0 family ethernet-switching native-vlan-id 1
user@switch# set xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
user@switch# set xe-0/0/11 unit 0 family ethernet-switching native-vlan-id 1
```

5. Set the MTU to 2180 for each Ethernet interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 mtu 2180
user@switch# set xe-0/0/7 mtu 2180
user@switch# set xe-0/0/8 mtu 2180
user@switch# set xe-0/0/9 mtu 2180
user@switch# set xe-0/0/10 mtu 2180
user@switch# set xe-0/0/11 mtu 2180
```

6. Assign the Ethernet interfaces to the FCoE VLAN:

```
[edit vlans blue interface]
user@switch# set xe-0/0/6.0
user@switch# set xe-0/0/7.0
user@switch# set xe-0/0/8.0
user@switch# set xe-0/0/9.0
user@switch# set xe-0/0/10.0
user@switch# set xe-0/0/11.0
```

7. Disable IGMP snooping on the FCoE VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan blue disable
```

8. Configure the FCoE VLAN interface and port mode for the FCoE traffic:

```
[edit interfaces]
user@switch# set vlan unit 100 family fibre-channel port-mode f-port
```

9. Define the FCoE VLAN interface as the interface for the FCoE VLAN:

```
[edit vlans]
user@switch# set blue l3-interface vlan.100
```

10. Configure the physical FC interfaces the fabric uses to connect to the FC switch:

```
[edit chassis fpc 0 pic 0]
user@switch# set fibre-channel port-range 0 5
```



**NOTE:** When you configure ports as FC ports, the port designation changes from xe-n/n/n.n format to fc-n/n/n.n format to indicate that the interface is an FC interface. FC interfaces do not support 10-Gbps interface speed but instead conform to FC interface speeds of 2 Gbps, 4 Gbps, or 8 Gbps.

11. Configure the native FC interfaces and port mode:

```
[edit interfaces]
user@switch# set fc-0/0/0 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/1 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/2 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/3 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/4 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/5 unit 0 family fibre-channel port-mode np-port
```

12. Configure the native FC interface port speed:

```
[edit interfaces]
user@switch# set fc-0/0/0 fibrechannel-options speed 4g
user@switch# set fc-0/0/1 fibrechannel-options speed 4g
user@switch# set fc-0/0/2 fibrechannel-options speed 4g
user@switch# set fc-0/0/3 fibrechannel-options speed 4g
user@switch# set fc-0/0/4 fibrechannel-options speed 4g
user@switch# set fc-0/0/5 fibrechannel-options speed 4g
```

13. Configure the FC fabric name and unique ID:

```
[edit fc-fabrics]
user@switch# set fcproxy1 fabric-id 1
```

14. Define the FC fabric as an FCoE-FC gateway:

```
[edit fc-fabrics]
user@switch# set fcproxy1 fabric-type proxy
```

15. Assign the FCoE VLAN interface to the fabric:

```
[edit fc-fabrics]
user@switch# set fcproxy1 interface vlan.100
```

16. Assign the native FC interfaces to the fabric:

```
[edit fc-fabrics]
user@switch# set fcproxy1 interface fc-0/0/0.0
user@switch# set fcproxy1 interface fc-0/0/1.0
user@switch# set fcproxy1 interface fc-0/0/2.0
user@switch# set fcproxy1 interface fc-0/0/3.0
user@switch# set fcproxy1 interface fc-0/0/4.0
user@switch# set fcproxy1 interface fc-0/0/5.0
```

**Results** Display the results of the configuration:

```
user@switch> show configuration
fc-0/0/0 {
 fibrechannel-options {
 speed 4g;
 }
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
}
fc-0/0/1 {
 fibrechannel-options {
 speed 4g;
 }
}
```

```
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
 }
 fc-0/0/2 {
 fibrechannel-options {
 speed 4g;
 }
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
 }
 fc-0/0/3 {
 fibrechannel-options {
 speed 4g;
 }
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
 }
 fc-0/0/4 {
 fibrechannel-options {
 speed 4g;
 }
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
 }
 fc-0/0/5 {
 fibrechannel-options {
 speed 4g;
 }
 unit 0 {
 family fibre-channel {
 port-mode np-port;
 }
 }
 }
 xe-0/0/6 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
 members blue;
 }
 native-vlan-id 1;
 }
 }
 }
}
```

```
 }
 }
 xe-0/0/7 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
 members blue;
 }
 native-vlan-id 1;
 }
 }
 }
 xe-0/0/8 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
 members blue;
 }
 native-vlan-id 1;
 }
 }
 }
 xe-0/0/9 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
 members blue;
 }
 native-vlan-id 1;
 }
 }
 }
 xe-0/0/10 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
 members blue;
 }
 native-vlan-id 1;
 }
 }
 }
 xe-0/0/11 {
 mtu 2180;
 unit 0 {
 family ethernet-switching {
 port-mode tagged-access;
 vlan {
```



```

 members blue;
 }
 native-vlan-id 1;
}
}
vlan {
 unit 100 {
 family fibre-channel {
 port-mode f-port;
 }
 }
}
fc-fabrics {
 fcproxy1 {
 fabric-id 1
 fabric-type proxy
 interface {
 vlan.100
 fc-0/0/0.0;
 fc-0/0/1.0;
 fc-0/0/2.0;
 fc-0/0/3.0;
 fc-0/0/4.0;
 fc-0/0/5.0;
 }
 }
}
protocols {
 igmp-snooping {
 vlan blue {
 disable;
 }
 }
}
vlans {
 blue {
 vlan-id 100
 interface {
 xe-0/0/6.0;
 xe-0/0/7.0;
 xe-0/0/8.0;
 xe-0/0/9.0;
 xe-0/0/10.0;
 xe-0/0/11.0;
 }
 l3-interface vlan.100
 }
 native {
 vlan-id 1;
 }
}

```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the native FC interfaces and FCoE VLAN interface have been created, added to the FC fabric, and are operating properly, perform these tasks:

- [Verifying That the Native FC Interfaces and the FCoE VLAN Interface Have Been Created on page 5110](#)
- [Verifying That the FCoE VLAN Includes the Correct Ethernet Interfaces on page 5111](#)
- [Verifying That the FC Fabric Includes the Correct Interfaces on page 5111](#)
- [Verifying Native FC Interface Operation on page 5112](#)
- [Verifying That IGMP Snooping Has Been Disabled on the FCoE VLAN on page 5112](#)

### *Verifying That the Native FC Interfaces and the FCoE VLAN Interface Have Been Created*

**Purpose** Verify that the six native FC interfaces and the FCoE VLAN interface have been created on the switch and are configured in the correct mode.

**Action** List all of the FC interfaces configured on the switch using the `show fibre-channel interfaces` command:

```
user@switch> show fibre-channel interfaces
```

| Interface  | Idx | Type | Native<br>Fabric-id | NPIV | Config<br>Mode | Oper<br>Mode | State |
|------------|-----|------|---------------------|------|----------------|--------------|-------|
| fc-0/0/0.0 | 70  | FC   | 1                   | YES  | NP             | NP           | up    |
| fc-0/0/1.0 | 71  | FC   | 1                   | YES  | NP             | NP           | up    |
| fc-0/0/2.0 | 72  | FC   | 1                   | YES  | NP             | NP           | up    |
| fc-0/0/3.0 | 73  | FC   | 1                   | YES  | NP             | NP           | up    |
| fc-0/0/4.0 | 74  | FC   | 1                   | YES  | NP             | NP           | up    |
| fc-0/0/5.0 | 75  | FC   | 1                   | YES  | NP             | NP           | up    |
| vlan.100   | 67  | FCoE | 1                   | YES  | F              | F            | up    |

**Meaning** The `show fibre-channel interfaces` command lists all native FC interfaces and FCoE VLAN interfaces configured on the switch. The command output shows that the FC interfaces `fc-0/0/0.0`, `fc-0/0/1.0`, `fc-0/0/2.0`, `fc-0/0/3.0`, `fc-0/0/4.0`, and `fc-0/0/5.0` have been created and that those six interfaces:

- Are native Fibre Channel interfaces (type **FC**).
- Belong to the FC fabric with a configured fabric ID of 1.
- Are capable of N\_Port ID virtualization (NPIV).
- Have a configured mode and an operational mode of proxy N\_Port (**NP**), which means that they should be connected to an FCF or an FC switch, not to an FCoE device, and that they carry native FC traffic.
- Show an operational state of **up**.

The command output also shows that the FCoE VLAN interface **vlan.100** has been created and that interface:

- Is an FCoE VLAN interface (type **FCOE**).
- Belongs to the FC fabric with a configured fabric ID of 1.
- Is capable of N\_Port ID virtualization (NPIV).
- Has a configured mode and an operational mode of F\_Port (F), which means that its interfaces connect to FCoE devices and carry FCoE traffic.
- Shows an operational state of **up**.

#### *Verifying That the FCoE VLAN Includes the Correct Ethernet Interfaces*

**Purpose** Verify that the FCoE VLAN **blue** has been created with the correct VLAN tag (**100**) and with the correct Ethernet interfaces.

**Action** List all of the interfaces configured on the switch in VLAN **blue** using the **show vlans** command:

```
user@switch> show vlans blue
Name Tag Interfaces
blue 100
 xe-0/0/6.0, xe-0/0/7.0, xe-0/0/8.0, xe-0/0/9.0, xe-0/0/10.0
 xe-0/0/11.0
```

**Meaning** The **show vlans blue** command lists the interfaces that are members of the FCoE VLAN **blue**. The command output shows that the **blue** VLAN has a tag ID of 100 and includes the interfaces **xe-0/0/6.0**, **xe-0/0/7.0**, **xe-0/0/8.0**, **xe-0/0/9.0**, **xe-0/0/10.0**, and **xe-0/0/11.0**.

#### *Verifying That the FC Fabric Includes the Correct Interfaces*

**Purpose** Verify that the FC fabric configuration is configured on the switch with the correct native FC and FCoE VLAN interfaces.

**Action** List all of the interfaces configured on FC fabrics on the switch using the **show fibre-channel fabric** command:

```
user@switch> show fibre-channel fabric
Name Fabric-id Type Interfaces
fcproxy1 1 PROXY
 fc-0/0/0.0
 fc-0/0/1.0
 fc-0/0/2.0
 fc-0/0/3.0
 fc-0/0/4.0
 fc-0/0/5.0
 vlan.100
```

**Meaning** The **show fibre-channel fabric** command lists the interfaces that are members of each FC fabric. The command output shows that the only fabric configured on the switch is

named **fcproxy1**, has a fabric-id of 1, and is a **proxy** fabric in an FCoE-FC gateway. The command output also shows that the native FC interfaces **fc-0/0/0.0**, **fc-0/0/1.0**, **fc-0/0/2.0**, **fc-0/0/3.0**, **fc-0/0/4.0**, and **fc-0/0/5.0**, and the FCoE VLAN interface **vlan.100** belong to **fcproxy1**.

#### *Verifying Native FC Interface Operation*

**Purpose** Verify that the native FC interfaces are online and display the number of FC sessions on each interface.

**Action** List all of the native FC NP\_Port interface states and sessions by FC fabric using the **show fibre-channel proxy np-port** command:

```
user@switch> show fibre-channel proxy np-port
Fabric: fcproxy1, Fabric-id: 1
NP-Port State Sessions LB state LB weight
fc-0/0/0.0 online 3 ON 4
fc-0/0/1.0 online 3 ON 4
fc-0/0/2.0 online 2 ON 4
fc-0/0/3.0 online 2 ON 4
fc-0/0/4.0 online 2 ON 4
fc-0/0/5.0 online 2 ON 4
```

**Meaning** The **show fibre-channel proxy np-port** command lists the interfaces that are configured as native FC proxy N\_Port interfaces. The command output shows:

- The fabric name is **fcproxy1** and its fabric ID is 1.
- The interfaces are **online**.
- The number of FC sessions (virtual links) running on each interface.
- The load-balancing (LB) state is **ON** for all of the interfaces.
- The LB weight reflects the port speed of each interface, which is 4 Gbps.

#### *Verifying That IGMP Snooping Has Been Disabled on the FCoE VLAN*

**Purpose** Verify that IGMP snooping is disabled on the FCoE VLAN.

**Action** List the IGMP snooping protocol information for the FCoE VLAN using the **show configuration protocols igmp-snooping vlan blue** command:

```
user@switch> show configuration protocols igmp-snooping vlan blue
disable;
```

**Meaning** The **show configuration protocols igmp-snooping vlan blue** command lists the IGMP snooping configuration for the FCoE VLAN. The command output shows that IGMP snooping is disabled on the FCoE VLAN.

**Related Documentation**

- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway on page 5189](#)

- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Disabling VN2VF\\_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

## Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS flag in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

To configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic.
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.
- Configure the bandwidth scheduling for the FCoE forwarding class output queue.
- Create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- Configure the bandwidth scheduling for the FCoE priority group.
- Apply the scheduling to the egress interfaces.



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding classes for the lossless fcoe forwarding class. If you explicitly configure default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 5114](#)
- [Overview on page 5114](#)
- [Configuration on page 5115](#)
- [Verification on page 5119](#)

## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to:

- Assign FCoE traffic to the FCoE priority at the ingress.
- Create and apply CoS for the FCoE traffic using ETS (hierarchical port scheduling).
- Apply PFC to the FCoE traffic.
- Apply the configuration to ingress and egress interfaces.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Each interface in this example is configured as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and port scheduling are applied to all of the interfaces.

## Topology

[Table 405 on page 5114](#) shows the configuration components for this example.

**Table 405: Components of the PFC for FCoE Traffic Configuration Topology**


| Component                                                                                                   | Settings                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                                                    | QFX3500 switch                                                                                                                                              |
| Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point) | Code point <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |

Table 405: Components of the PFC for FCoE Traffic Configuration Topology (*continued*)

| Component                                  | Settings                                                                                                                |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| PFC congestion notification profile        | <b>fcoe-cnp:</b><br>Code point <b>011</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>      |
| FCoE queue scheduler                       | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>               |
| Forwarding class-to-scheduler mapping      | Scheduler map <b>fcoe-map:</b><br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>                           |
| Forwarding class set (FCoE priority group) | <b>fcoe-pg:</b><br>Forwarding class <b>fcoe</b><br>Egress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |
| Traffic control profile                    | <b>fcoe-tcp:</b><br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b>       |

Figure 196 on page 5115 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 196: PFC for FCoE Traffic Configuration Components Block Diagram



### Configuration

#### CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```

[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier

```

```

set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp

```

### Step-by-Step Procedure

To configure the FCoE forwarding class (priority), ingress classifier, output queue scheduling, forwarding class set (priority group) and its output port scheduling, PFC application, and interfaces to set up PFC for FCoE traffic:

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority
low code-points 011

```

2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc

```

3. Apply the PFC configuration to the ingress interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp

```

4. Assign the classifier to the ingress interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier

```

5. Configure output scheduling for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```

6. Map the FCoE forwarding class to the FCoE scheduler:

```

[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched

```

7. Configure the forwarding class set for the FCoE traffic:

```

[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe

```



8. Define the traffic control profile for the FCoE forwarding class set:

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

9. Apply the FCoE forwarding class set and traffic control profile to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

### Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class):

```
user@switch> show configuration class-of-service
classifiers {
 ieee-802.1 fcoe-classifier {
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 }
}
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/31 {
 congestion-notification-profile fcoe-cnp;
```

```
forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
}
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
}
xe-0/0/32 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
xe-0/0/33 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
xe-0/0/34 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
```

```

schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}

```



**TIP:** To quickly configure the interfaces, issue the load merge terminal command and then copy the hierarchy and paste it into the switch terminal window.

### Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5119](#)
- [Verifying the Ingress Interface PFC Configuration on page 5120](#)

#### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |
| 100      |                     |
|          | 4                   |
| 101      |                     |
|          | 5                   |
| 110      |                     |

```
111 6
 7
```

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point 011 for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

#### *Verifying the Ingress Interface PFC Configuration*

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
```

**Meaning** The **show configuration class-of-service interfaces** commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

### Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX3500 switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree protocol (STP).

You can use an MC-LAG to provide a redundant aggregation layer for Fiber Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX3500 switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.



**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX3500 switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX3500 switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see [“Example: Configuring Multichassis Link Aggregation” on page 2462](#). However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.

QFX3500 Node devices support MC-LAGs. QFabric system Node devices do not support MC-LAGs.

This topic describes:

- [Requirements on page 5122](#)
- [Overview on page 5122](#)
- [Configuration on page 5127](#)
- [Verification on page 5135](#)

## Requirements

---

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 Switches that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX3500 Switches that provide FCoE server access in transit switch mode, and connect to the MC-LAG switches. These switches can be standalone QFX3500 switches or they can be Node devices in a QFabric system.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 12.2 or later for the QFX Series.

## Overview

---

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX3500 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Disable IGMP snooping on the FCoE VLAN.
- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

This example does not show how to configure the MC-LAG itself. This example does include a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

## Topology

QFX3500 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 197 on page 5123](#):

Figure 197: Supported Topology for an MC-LAG on an FCoE Transit Switch

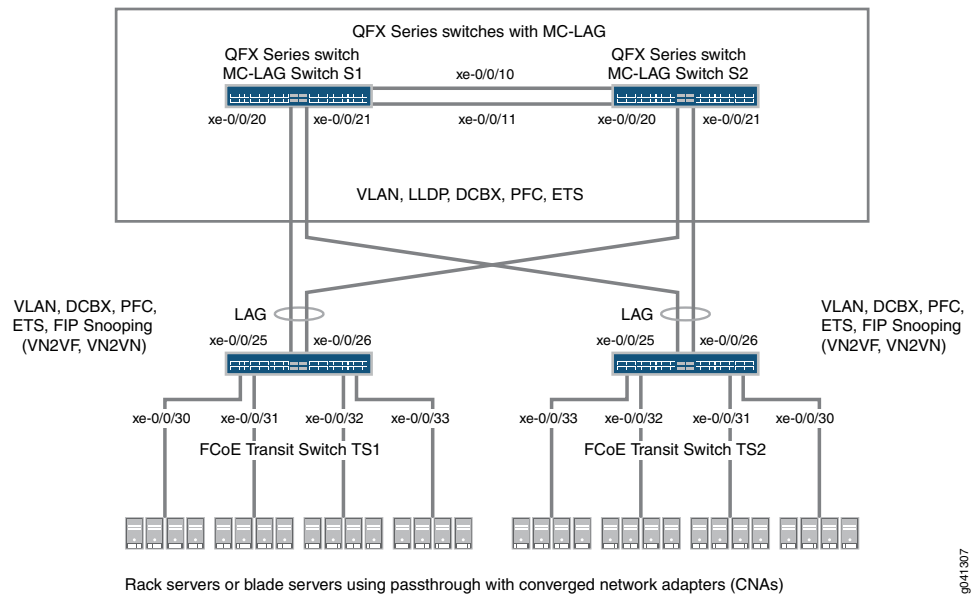


Table 406 on page 5123 shows the configuration components for this example.

Table 406: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

| Component                                                                  | Settings                                                                                                                |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                   | Four QFX3500 switches (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access) |
| Forwarding class (all switches)                                            | Default <b>fc0e</b> forwarding class.                                                                                   |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.                                                          |

Table 406: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (*continued*)

| Component                                                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAGs and MC-LAG                                          | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                      | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Forwarding class-to-scheduler mapping (all switches)     | Scheduler map <b>fcoe-map</b> :<br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Forwarding class set (FCoE priority group, all switches) | <b>fcoe-pg:</b><br>Forwarding class <b>fcoe</b><br><br>Egress interfaces: <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Traffic control profile (all switches)                   | <b>fcoe-tcp:</b><br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



**Table 406: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)**

| Component                                          | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC congestion notification profile (all switches) | <p><b>fcoe-cnp:</b><br/>Code point 011</p> <p>Ingress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG ae0 and MC-LAG ae1</li> <li>• S2—LAG ae0 and MC-LAG ae1</li> <li>• TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> <li>• TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> </ul>                                                                                                                                                                   |
| FCoE VLAN name and tag ID                          | <p>Name—<b>fcoe_vlan</b><br/>ID—100</p> <p>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.</p> <p>Disable IGMP snooping on the FCoE VLAN on all four switches.</p>                                                                                                                                                                                                                                                                                                                       |
| FIP snooping                                       | <p>Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.</p> <p>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.</p> |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode and tagged access mode ports if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is **queue 3**.



**NOTE:** In Junos OS Release 12.2, traffic mapped to explicitly configured forwarding classes, even lossless forwarding classes such as **fcoe**, is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment. To receive lossless treatment in release 12.2, traffic must use one of the default lossless forwarding classes (**fcoe** or **no-loss**).

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.

- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk and tagged-access port modes. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (011) to the FCoE forwarding class. You can configure the BA classifier if you wish, but you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point 011 in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers and how to disable IGMP snooping on the FCoE VLAN. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required

to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. ([“Example: Configuring Multichassis Link Aggregation” on page 2462](#) describes how to configure MC-LAGs.)
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. ([“Configuring Link Aggregation” on page 2537](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

### Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [Configuring MC-LAG Switches S1 and S2 on page 5129](#)
- [Configuring FCoE Transit Switches TS1 and TS2 on page 5130](#)
- [Results on page 5132](#)

#### CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the [edit] hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the [edit] hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
set ethernet-switching-options secure-access-port vlan fcoe_vlan examine-fip examine-vn2v2
beacon-period 90000
```

**Configuring MC-LAG Switches S1 and S2**

**Step-by-Step Procedure** To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:
 

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):
 

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```
6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:
 

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
7. Apply the PFC configuration to the LAG and MC-LAG interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):
 

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
9. Disable IGMP snooping on the FCoE VLAN:
 

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG between the two MC-LAG switches:

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```

11. Add the member interfaces to the MC-LAG:

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```

12. Configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and MC-LAG interfaces (2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes):

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```

14. Set the LAG and MC-LAG interfaces as FCoE trusted ports (ports that connect to other switches should be trusted and should not perform FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

### *Configuring FCoE Transit Switches TS1 and TS2*

#### **Step-by-Step Procedure**

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:

- ```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:


```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
 5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```
 6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:


```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
 7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile
fcoe-cnp
```
 8. Configure the VLAN for FCoE traffic (**fcoe_vlan**):


```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
 9. Disable IGMP snooping on the FCoE VLAN:


```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```
 10. Add the member interfaces to the LAG:


```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```
 11. On the LAG (**ae1**), configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe_vlan**):

```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

12. On the FCoE access interfaces (xe-0/0/30, xe-0/0/31, xe-0/0/32, xe-0/0/33), configure the port mode as **tagged-access** and membership in the FCoE VLAN (**fcoe_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and FCoE access interfaces (2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes):

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

14. Set the LAG interface as an FCoE trusted port (ports that connect to other switches should be trusted and should not perform FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

15. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN_Port FIP snooping; the example is equally valid if you use VN2VF_Port FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port vlan fcoe_vlan
examine-fip examine-vn2v2 beacon-period 90000
```

Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are similar):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
```



```

}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  ae0 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  ae1 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}

```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see [“Example: Configuring Multichassis Link Aggregation” on page 2462](#).

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are similar):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  xe-0/0/30 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/31 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/32 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/33 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
```

```

}
ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
}

```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 5136](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created on page 5136](#)
- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created on page 5137](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5137](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 5138](#)
- [Verifying That the Interfaces Are Correctly Configured on page 5140](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 5142](#)
- [Verifying That the FIP Snooping Mode is Correct on FCoE Transit Switches TS1 and TS2 on page 5143](#)
- [Verifying That IGMP Snooping Is Disabled on the FCoE VLAN on page 5143](#)

Verifying That the Output Queue Schedulers Have Been Created

Purpose Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

Action List the scheduler map using the operational mode command **show class-of-service scheduler-map fcoe-map**:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      <default-drop-profile>
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created

Purpose Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

Action List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
Shaping rate: 100 percent
```

```
Scheduler map: fcoe-map
Guaranteed rate: 3000000000
```

Meaning The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

Verifying That the Forwarding Class Set (Priority Group) Has Been Created

Purpose Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

Action List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class      Index
  fcoe                  1
```

Meaning The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

Action List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled   2500
  100      Disabled
  101      Disabled
```

110	Disabled
111	Disabled
Type: Output	
Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	
	7

Meaning The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying That the Interface Class of Service Configuration Has Been Created

Purpose Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

Action List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}

ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
```

Meaning The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



NOTE: Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

Verifying That the Interfaces Are Correctly Configured

Purpose Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

Action List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/20 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
```



```

        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```

user@switch> show configuration interfaces
xe-0/0/25 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/26 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/30 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/31 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/32 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {

```

```
        port-mode tagged-access;
        vlan {
            members fcoe_vlan;
        }
    }
}

ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
```

Meaning The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The port mode (**trunk** mode for interfaces that connect two switches, **tagged-access** mode for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe_vlan**)

Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

Purpose Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Switch TS1 and TS2 FCoE access ports.

Action List the port security configuration on FCoE transit switches TS1 and TS2 using the operational mode command **show configuration ethernet-switching-options secure-access-port**:

```
user@switch> show configuration ethernet-switching-options secure-access-port
interface ae1.0 {
    fcoe-trusted;
}
vlan fcoe_vlan {
```

```
    examine-fip;
}
```

Meaning The `show configuration ethernet-switching-options secure-access-port` command lists port security information, including whether a port is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- FIP snooping is enabled (**examine-fip**) on the FCoE VLAN (**fcoe_vlan**). On Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

Verifying That the FIP Snooping Mode is Correct on FCoE Transit Switches TS1 and TS2

Purpose Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Switch TS1 and TS2 FCoE access ports.

Action List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,    Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
...
```



NOTE: The output has been truncated to show only the relevant information.

Meaning The `show fip snooping brief` command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe_vlan**
- The FIP snooping mode is VN2VN_Port FIP snooping (**VN2VN Snooping**)

Verifying That IGMP Snooping Is Disabled on the FCoE VLAN

Purpose Verify that IGMP snooping is disabled on the FCoE VLAN on all four switches.

Action List the IGMP snooping protocol information on each of the four switches using the **show configuration protocols igmp-snooping** command:

```
user@switch> show configuration protocols igmp-snooping
vlan fcoe_vlan {
    disable;
}
```

Meaning The **show configuration protocols igmp-snooping** command lists the IGMP snooping configuration for the VLANs configured on the switch. The command output shows that IGMP snooping is disabled on the FCoE VLAN (**fcoe_vlan**).

- Related Documentation**
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
 - [Configuring Link Aggregation on page 2537](#)
 - [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
 - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Understanding Multichassis Link Aggregation on page 2435](#)
 - [Understanding MC-LAGs on an FCoE Transit Switch on page 5047](#)

Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



NOTE: LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The QFX Series handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the QFX Series exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).

- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 5145](#)
- [Overview on page 5145](#)
- [Configuration on page 5149](#)
- [Verification on page 5150](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device
- Junos OS Release 12.1 or later for the QFX Series

Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



NOTE: DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 5204](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.
- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 407 on page 5146](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 408 on page 5146](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

Table 407: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

Table 408: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low

Table 408: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier) (*continued*)

Code Point	Forwarding Class	Loss Priority
110	best-effort	low
111	best-effort	low

Topology

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



NOTE: You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 409 on page 5147 shows the configuration components for this example.

Table 409: Components of DCBX Application Protocol Exchange Configuration Topology

Component	Settings
Hardware	QFX Series device
LLDP	Enabled by default on Ethernet interfaces
DCBX	Enabled by default on Ethernet interfaces
iSCSI application (Layer 4)	Application name— iscsi protocol— TCP destination-port— 3260 code-points— 111
PTP application (Layer 2)	Application name— ptp ether-type— 0x88F7 code-points— 001, 101

Table 409: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)

Component	Settings
FCoE application (Layer 2)	<p>Application name—fcoe</p> <p>ether-type—0x8906</p> <p>code-points—011</p> <p>NOTE: You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map.</p>
Application maps	<p>dcbx-iscsi-fcoe-app-map—Maps the iSCSI and FCoE applications to IEEE 802.1p code points</p> <p>dcbx-iscsi-ptp-app-map—Maps iSCSI and PTP applications to IEEE 802.1p code points</p>
Interfaces	<p>xe-0/0/10—Configured to exchange FCoE and iSCSI application TLVs (uses application map dcbx-iscsi-fcoe-app-map, carries FCoE traffic, and has PFC enabled on the FCoE priority)</p> <p>xe-0/0/11—Configured to exchange iSCSI and PTP application TLVs (uses application map dcbx-iscsi-ptp-app-map)</p>
PFC congestion notification profile for FCoE application exchange	<p>fcoe-cnp:</p> <ul style="list-style-type: none"> Code point—011 Interface—xe-0/0/10
Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point)	<p>fcoe-iscsi-cl1:</p> <ul style="list-style-type: none"> Maps the fcoe forwarding class to the IEEE 802.1p code point used for the FCoE application (011) and a loss priority of high Maps the network-control forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of high Applied to interface xe-0/0/10 <p>iscsi-ptp-cl2:</p> <ul style="list-style-type: none"> Maps the network-control forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of low Maps the best-effort forwarding class to the IEEE 802.1p code points used for the PTP application (001 and 101) and a loss priority of low Applied to interface xe-0/0/11



NOTE: This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

Configuration

CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

Configuring DCBX Application Protocol TLV Exchange

Step-by-Step Procedure

To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.


```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.


```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.


```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```

4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
```

5. Create the congestion notification profile to enable PFC on the FCoE code point (**011**), and apply the congestion notification profile to interface **xe-0/0/10**.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```

6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority
high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control
loss-priority high code-points 111
```

7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
```

8. Apply the classifiers to the appropriate interfaces.

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ntp-cl2
```

Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 5150](#)
- [Verifying the Application Map Configuration on page 5151](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 5151](#)
- [Verifying the PFC Configuration on page 5152](#)
- [Verifying the Classifier Configuration on page 5153](#)

Verifying the Application Configuration

Purpose Verify that DCBX applications have been configured.

Action List the applications by using the configuration mode command **show applications**:

```
user@switch# show applications
```

```

application iSCSI {
    protocol tcp;
    destination-port 3260;
}

application fcoe {
    ether-type 0x8906;
}

application ptp {
    ether-type 0x88F7;
}

```

Meaning The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

Verifying the Application Map Configuration

Purpose Verify that the application maps have been configured.

Action List the application maps by using the configuration mode command **show policy-options application-maps**:

```

user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
    application iSCSI code-points 111;
    application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
    application iSCSI code-points 111;
    application PTP code-points [001 101];
}

```

Meaning The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the FCoE application, which is mapped to IEEE 802.1p code point **011**.

The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the PTP application, which is mapped to IEEE 802.1p code points **001** and **101**.

Verifying DCBX Application Protocol Exchange Interface Configuration

Purpose Verify that the application maps have been applied to the correct interfaces.

Action List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
    application-map dcbx-iscsi-ptp-app-map;
}
```

Meaning The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.

Verifying the PFC Configuration

Purpose Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

Action Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
    input {
        ieee-802.1 {
            code-point 011 {
                pfc;
            }
        }
    }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
    congestion-notification-profile fcoe-cnp;
}
```



NOTE: The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

Meaning The `show class-of-service congestion-notification-profile` configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point **011** (the default FCoE code point).

The `show class-of-service interfaces` configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

Verifying the Classifier Configuration

Purpose Verify that the classifiers have been configured and applied to the correct interfaces.

Action Display the classifier configuration by using the configuration mode command `show class-of-service`:

```
user@switch# show class-of-service
classifiers {
  ieee-802.1 fcoe-iscsi-cl1 {
    import default;
    forwarding-class network-control {
      loss-priority high code-points 111;
    }
    forwarding-class fcoe {
      loss-priority high code-points 011;
    }
  }
  ieee-802.1 iscsi-ptp-cl2 {
    import default;
    forwarding-class network-control {
      loss-priority low code-points 111;
    }
    forwarding-class best-effort {
      loss-priority low code-points [ 001 101 ];
    }
  }
}
interfaces {
  xe-0/0/10 {
    congestion-notification-profile fcoe-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-iscsi-cl1;
      }
    }
  }
  xe-0/0/11 {
    unit 0 {
      classifiers {
        ieee-802.1 iscsi-ptp-cl2;
      }
    }
  }
}
```



NOTE: The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

Meaning The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ptp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ptp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ptp-cl2** is mapped to interface **xe-0/0/11.0**.

Related Documentation

- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [show dcbx on page 5299](#)
- [show dcbx neighbors on page 5300](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Using DCBX Protocol to Lower Costs](#)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

This example shows how to configure VN_Port to VN_Port (VN2VN_Port) FIP snooping when the hosts are directly connected to the same FCoE transit switch.

VN2VN_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN_Port FIP snooping is conceptually similar to VN2VN_Port FIP snooping between VN_Ports and VF_Ports, but VN2VN_Port FIP snooping does not require traffic between VN_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN_Port communicates transparently through the transit switch on a virtual link that emulates a direct connection to the VN_Port at the other end of the virtual link.

To configure VN2VN_Port FIP snooping when the hosts are directly connected to the same FCoE transit switch, you must follow these configuration rules:

- VN2VN_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN_Port FIP snooping must use that FCoE VLAN. You cannot mix VN2VN_Port FIP snooping traffic with VN2VF_Port FIP snooping traffic in the same FCoE VLAN.



NOTE: An FCoE VLAN can support either VN2VF_Port FIP snooping or VN2VN_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF_Port FIP snooping traffic and for VN2VN_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN_Port FIP snooping VLANs, VN_Port to VF_Port (FIP snooping) traffic is dropped.

- ENode-facing ports must be set in **tagged-access** port mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** port mode.
- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable VN2VF_Port FIP snooping, the system snoops VN_Port to VF_Port packets and enforces security only on VN_Port to VF_Port virtual links. When you enable VN2VN_Port FIP snooping, the system snoops VN_Port to VN_Port packets and enforces security only on VN_Port to VN_Port virtual links.

The transit switch applies VN2VN_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN_Port FIP snooping when the FCoE hosts are directly connected to the same transit switch:

- [Requirements on page 5156](#)
- [Overview on page 5156](#)
- [Configuration on page 5157](#)
- [Verification on page 5158](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch used as a transit switch
- Junos OS Release 12.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

Overview

This example shows you how to:

- Set the correct interface port modes on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN_Port FIP snooping.
- Configure the dedicated FCoE VLAN for VN2VN_Port FIP snooping traffic.
- Enable VN2VN_Port FIP snooping on the FCoE VLAN and configure the beacon period.

Topology


[Table 410 on page 5156](#) shows the configuration components for this example.

Table 410: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

Component	Settings
Hardware	QFX3500 switch (FCoE transit switch TS1) Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)
Interfaces and port modes	<ul style="list-style-type: none"> • Interface xe-0/0/20, port mode tagged-access, connects directly to the FCoE host with ENode1. • Interface xe-0/0/21, port mode tagged-access, connects directly to the FCoE host with ENode2.
Interface VLAN membership	Both interfaces use VLAN vlan200 .
VN2VN_Port FIP snooping VLAN	VLAN name— vlan200 VLAN ID—200
FIP snooping mode and beacon period	Set examine-vn2vn (VN2VN_Port FIP snooping) Beacon period—90000 ms

Figure 198 on page 5157 shows the network topology for this example.

Figure 198: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Same Transit Switch) Topology



Configuration

CLI Quick Configuration

To quickly configure VN2VN_Port FIP snooping for FCoE hosts connected directly to the same transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000
```

Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

Step-by-Step Procedure

To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:

1. Configure the port modes of the interfaces that connect directly to the FCoE host ENodes:

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode tagged-access
```
2. Configure the interface VLAN membership so that the interfaces connected to the ENodes are members of the dedicated VN2VN_Port VLAN (vlan200):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```
3. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```
4. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2 beacon-period 90000
```

Verification

To verify that the VN2VN_Port FIP snooping configuration has been created and is operating properly, perform these tasks:

- [Verifying That VN2VN_Port FIP Snooping is Enabled on the FCoE VLAN on page 5158](#)
- [Verifying the Interface Port Mode on page 5159](#)

Verifying That VN2VN_Port FIP Snooping is Enabled on the FCoE VLAN

Purpose Verify that VN2VN_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and the correct interfaces (**xe-0/0/20** and **xe-0/0/21**) are members of the VLAN.

Action List the FIP snooping information using the operational mode command **show fip snooping detail**.

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
```

Meaning The **show fip snooping detail** command lists all of the transit switch information about VN2VN_Port FIP snooping and VN2VF_Port FIP snooping. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF_Port FIP snooping.)
- The beacon period is **90000**.
- The interfaces for the ENodes are **xe-0/0/20** and **xe-0/0/21**.

In addition, this useful command shows information about the ENodes and the VN2VN_Port sessions.

Verifying the Interface Port Mode

Purpose Verify that the interface port modes are **tagged-access**.

Action List the Ethernet switching interfaces to confirm the port mode using the **show ethernet-switching interfaces detail** operational command.

Use the operational mode commands **show ethernet-switching interfaces xe-0/0/20.0 detail** and **show ethernet-switching interfaces xe-0/0/21.0 detail** to list the Ethernet switching interface information. The output is truncated to show only the relevant portions:

```
user@switch> show ethernet-switching interfaces xe-0/0/20.0 detail
Interface: xe-0/0/20.0, Index: 75, State: up, Port mode: Tagged-Access
.
.
.

user@switch> show ethernet-switching interfaces xe-0/0/21.0 detail
Interface: xe-0/0/21.0, Index: 83, State: up, Port mode: Tagged-Access
.
.
.
```

Meaning The **show ethernet-switching interfaces detail** command lists the port mode as **tagged-access** for both interfaces.

- Related Documentation**
- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
 - [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5167](#)
 - [Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5202](#)
 - [Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)

This example shows how to configure VN_Port to VN_Port (VN2VN_Port) FIP snooping when the hosts are directly connected to different FCoE transit switches, and the transit switches are directly connected to each other.

VN2VN_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN_Port FIP snooping is conceptually similar to VN2VF_Port FIP snooping between VN_Ports and VF_Ports, but VN2VN_Port FIP snooping does not require traffic between VN_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a

VN_Port communicates transparently through one or more transit switches on a virtual link that emulates a direct connection to the VN_Port at the other end of the virtual link.

To configure VN2VN_Port FIP snooping when the hosts are directly connected to different FCoE transit switches, and the transit switches are directly connected to each other, you must follow these configuration rules:

- VN2VN_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN_Port FIP snooping must use that FCoE VLAN. The FCoE VLAN must be configured on each transit switch. You cannot mix VN2VN_Port FIP snooping traffic with VN2VF_Port FIP snooping traffic in the same FCoE VLAN.



NOTE: An FCoE VLAN can support either VN2VF_Port FIP snooping or VN2VN_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF_Port FIP snooping traffic and for VN2VN_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN_Port FIP snooping VLANs, VN2VF_Port traffic is dropped.

- ENode-facing ports must be set in **tagged-access** port mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** port mode.
- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable VN2VF_Port FIP snooping, the system snoops VN_Port to VF_Port packets and enforces security only on VN_Port to VF_Port virtual links. When you enable VN2VN_Port FIP snooping, the system snoops VN_Port to VN_Port packets and enforces security only on VN_Port to VN_Port virtual links.

The transit switch applies VN2VN_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN_Port FIP snooping when the FCoE hosts are directly connected to different transit switches, and the transit switches are directly connected to each other:

- [Requirements on page 5160](#)
- [Overview on page 5161](#)
- [Configuration on page 5162](#)
- [Verification on page 5164](#)

Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 Switches used as transit switches

- Junos OS Release 12.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

Overview

This example shows you how to:

- Set the correct interface port modes on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN_Port FIP snooping.
- Configure the network-facing interfaces as FCoE trusted interfaces.
- Configure the dedicated FCoE VLAN for VN2VN_Port FIP snooping traffic.
- Enable VN2VN_Port FIP snooping on the FCoE VLAN and configure the beacon period.


Topology

Table 411 on page 5161 shows the configuration components for this example.

Table 411: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to Different FCoE Transit Switches)

Component	Settings
Hardware	Two QFX3500 switches (FCoE transit switch TS1 and FCoE transit switch TS2) Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)
Interfaces and port modes	<ul style="list-style-type: none">• Interface xe-0/0/20, port mode tagged-access, connects directly from transit switch TS1 to the FCoE host with ENode1.• Interface xe-0/0/21, port mode trunk, connects directly from transit switch TS1 to transit switch TS2.• Interface xe-0/0/31, port mode trunk, connects directly from transit switch TS2 to transit switch TS1.• Interface xe-0/0/30, port mode tagged-access, connects directly from transit switch TS2 to the FCoE host with ENode2.
Interface VLAN membership	The interfaces on both transit switches use VLAN vlan200 .
VN2VN_Port FIP snooping VLAN	VLAN name (both transit switches)— vlan200 VLAN ID—200
FIP snooping mode and beacon period	Set examine-vn2vn (VN2VN_Port FIP snooping) Beacon period—90000 ms

Figure 199 on page 5162 shows the network topology for this example.

Figure 199: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Different Transit Switches) Topology

Configuration

To configure VN2VN_Port FIP snooping for VN_Ports that are directly connected to different transit switches (and the transit switches are directly connected to each other), perform these tasks:

- [Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS1 on page 5163](#)
- [Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS2 on page 5163](#)

CLI Quick Configuration

The configuration for each FCoE transit switch is shown separately.

To quickly configure VN2VN_Port FIP snooping for FCoE hosts connected directly to different transit switches, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS1:

```

set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set ethernet-switching-options secure-access-port interface xe-0/0/21 fcoe-trusted
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000

```

To quickly configure VN2VN_Port FIP snooping for FCoE hosts connected directly to different transit switches, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS2:

```

set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port interface xe-0/0/31 fcoe-trusted
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000

```

Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS1

Step-by-Step Procedure To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:

1. Configure the port modes of the interfaces that connect directly to the FCoE host with ENode1 (xe-0/0/20) and to FCoE transit switch TS2 (xe-0/0/21):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN_Port VLAN (vlan200):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the network-facing port (xe-0/0/21) as an FCoE trusted port:

```
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/21 fcoe-trusted
```

4. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

5. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2 beacon-period 90000
```

Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS2

Step-by-Step Procedure To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:

1. Configure the port modes of the interfaces that connect directly to the FCoE host with ENode2 (xe-0/0/30) and to FCoE transit switch TS1 (xe-0/0/31):

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN_Port VLAN (vlan200):

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the network-facing port (xe-0/0/31) as an FCoE trusted port:

```
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/31 fcoe-trusted
```

4. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

5. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip  
examine-vn2v2 beacon-period 90000
```

Verification

To verify that the VN2VN_Port FIP snooping configuration has been created and is operating properly on both switches, perform these tasks:

- [Verifying That VN2VN_Port FIP Snooping is Enabled on the FCoE VLAN \(Transit Switches TS1 and TS2\) on page 5164](#)
- [Verifying the Interface Port Mode on page 5166](#)

Verifying That VN2VN_Port FIP Snooping is Enabled on the FCoE VLAN (Transit Switches TS1 and TS2)

Purpose Verify that VN2VN_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and that the correct interfaces (**xe-0/0/20** and **xe-0/0/21** on TS1, and **xe-0/0/30** and **xe-0/0/31** on TS2) are members of the VLAN.

Action List the FIP snooping information on transit switch TS1 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
```

List the FIP snooping information on transit switch TS2 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/30
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/31
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
```

Meaning The **show fip snooping detail** command lists all of the transit switch information about VN2VN_Port FIP snooping and VN2VF_Port FIP snooping on each transit switch. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF_Port FIP snooping.)

- The beacon period is **90000**.
- The interfaces connected to the ENodes are **xe-0/0/20** and **xe-0/0/21** on transit switch TS1, and **xe-0/0/30** and **xe-0/0/31** on transit switch TS2. Because the transit switches are transparent passthrough switches, the network-facing trunk ports “see” the FCoE host ENodes at the far end of the VN2VN_Port virtual link.

In addition, this useful command shows information about the ENodes and the VN2VN_Port sessions.

Verifying the Interface Port Mode

Purpose Verify that the interface port modes are **tagged-access** for ENode-facing ports and **trunk** for network-facing ports on each transit switch.

Action List the Ethernet switching interfaces to confirm the port mode using the **show ethernet-switching interfaces detail** operational command.

Use the operational mode commands **show ethernet-switching interfaces xe-0/0/20.0 detail** and **show ethernet-switching interfaces xe-0/0/21.0 detail** to list the Ethernet switching interface information on FCoE transit switch TS1. The output is truncated to show only the relevant portions:

```
user@switch> show ethernet-switching interfaces xe-0/0/20.0 detail
Interface: xe-0/0/20.0, Index: 75, State: up, Port mode: Tagged-Access
.
.
.
```

```
user@switch> show ethernet-switching interfaces xe-0/0/21.0 detail
Interface: xe-0/0/21.0, Index: 83, State: up, Port mode: Trunk
.
.
.
```

List the Ethernet switching interface information on FCoE transit switch TS2 using the operational mode commands **show ethernet-switching interfaces xe-0/0/30.0 detail** and **show ethernet-switching interfaces xe-0/0/31.0 detail**:

```
user@switch> show ethernet-switching interfaces xe-0/0/30.0 detail
Interface: xe-0/0/30.0, Index: 56, State: up, Port mode: Tagged-Access
.
.
.
```

```
user@switch> show ethernet-switching interfaces xe-0/0/31.0 detail
Interface: xe-0/0/31.0, Index: 59, State: up, Port mode: Trunk
.
.
.
```

Related Documentation

- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\)](#) on page 5155
- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\)](#) on page 5167

- [Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5202](#)
- [Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)

This example shows how to configure VN_Port to VN_Port (VN2VN_Port) FIP snooping when the hosts are indirectly connected through an aggregation layer FCoE transit switch. Each FCoE host ENode is directly connected to an FCoE transit switch, but the FCoE transit switches are not directly connected to each other. The FCoE transit switches are both connected to a third FCoE transit switch that acts as an aggregation layer switch.

VN2VN_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN_Port FIP snooping is conceptually similar to VN2VN_Port FIP snooping between VN_Ports and VF_Ports, but VN2VN_Port FIP snooping does not require traffic between VN_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN_Port communicates transparently through one or more transit switches on a virtual link that emulates a direct connection to the VN_Port at the other end of the virtual link.

To configure VN2VN_Port FIP snooping when the hosts are indirectly connected, you must follow these configuration rules:

- VN2VN_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN_Port FIP snooping must use that FCoE VLAN. The FCoE VLAN must be configured on each transit switch. You cannot mix VN2VN_Port FIP snooping traffic with VN2VF_Port FIP snooping traffic in the same FCoE VLAN.



NOTE: An FCoE VLAN can support either VN2VF_Port FIP snooping or VN2VN_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF_Port FIP snooping traffic and for VN2VN_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN_Port FIP snooping VLANs, VN_Port to VF_Port traffic is dropped.

- ENode-facing ports must be set in **tagged-access** port mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** port mode.
- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable FIP snooping, the system snoops VN_Port to VF_Port packets and enforces security only on VN_Port to VF_Port virtual links. When you enable VN2VN_Port FIP snooping, the system snoops VN_Port to VN_Port packets and enforces security only on VN_Port to VN_Port virtual links.

The transit switch applies VN2VN_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN_Port FIP snooping when the FCoE hosts are indirectly connected across an aggregation layer FCoE transit switch:

- [Requirements on page 5168](#)
- [Overview on page 5168](#)
- [Configuration on page 5169](#)
- [Verification on page 5172](#)

Requirements

This example uses the following hardware and software components:

- Three Juniper Networks QFX3500 Switches used as transit switches
- Junos OS Release 12.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

Overview

This example shows you how to:

- Set the correct interface port modes on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN_Port FIP snooping.
- Configure the network-facing interfaces as FCoE trusted interfaces.
- Configure the dedicated FCoE VLAN for VN2VN_Port FIP snooping traffic.
- Enable VN2VN_Port FIP snooping on the FCoE VLAN and configure the beacon period.

Topology

[Table 412 on page 5168](#) shows the configuration components for this example.

Table 412: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Indirectly Connected Across an Aggregation Layer FCoE Transit Switch)

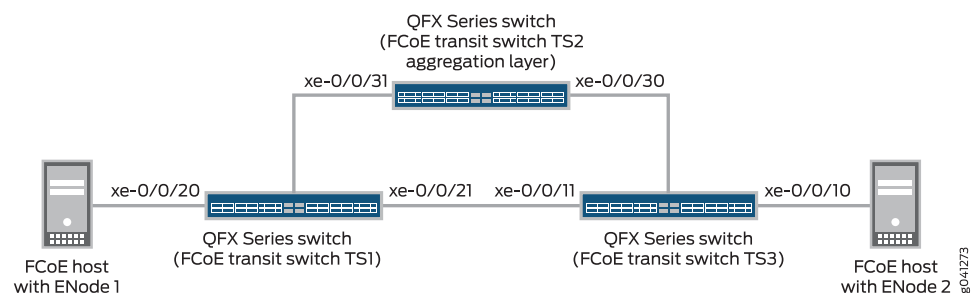
Component	Settings
Hardware	<p>Three QFX3500 switches, two of which are FCoE transit switches that are directly attached to the FCoE hosts (transit switches TS1 and TS2) and one of which is an aggregation layer FCoE transit switch (TS3)</p> <p>Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)</p>

Table 412: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Indirectly Connected Across an Aggregation Layer FCoE Transit Switch) (continued)

Component	Settings
Interfaces and port modes	<ul style="list-style-type: none"> Interface xe-0/0/20, port mode tagged-access, connects directly from transit switch TS1 to the FCoE host with ENode1. Interface xe-0/0/21, port mode trunk, connects directly from transit switch TS1 to aggregation layer transit switch TS2. Interface xe-0/0/31, port mode trunk, connects directly from aggregation layer transit switch TS2 to transit switch TS1. Interface xe-0/0/30, port mode trunk, connects directly from aggregation layer transit switch TS2 to transit switch TS3. Interface xe-0/0/11, port mode trunk, connects directly from transit switch TS3 to aggregation layer transit switch TS2. Interface xe-0/0/10, port mode tagged-access, connects directly from transit switch TS3 to the FCoE host with ENode2.
Interface VLAN membership	The interfaces on all three switches use VLAN vlan200 .
VN2VN_Port FIP snooping VLAN	VLAN name (all three switches)— vlan200 VLAN ID—200
FIP snooping mode and beacon period	Set examine-vn2vn (VN2VN_Port FIP snooping) Beacon period—90000 ms

Figure 200 on page 5169 shows the network topology for this example.

Figure 200: VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected) Topology



Configuration

To configure VN2VN_Port FIP snooping for VN_Ports that are indirectly connected across an aggregation layer FCoE transit switch, perform these tasks:

- [Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS1 on page 5170](#)
- [Configuring VN2VN_Port FIP Snooping on Aggregation Layer FCoE Transit Switch TS2 on page 5171](#)
- [Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS3 on page 5172](#)

CLI Quick Configuration

The configuration for each FCoE transit switch is shown separately.

To quickly configure VN2VN_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS1:

```
set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set ethernet-switching-options secure-access-port interface xe-0/0/21 fcoe-trusted
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000
```

To quickly configure VN2VN_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS2:

```
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port interface xe-0/0/30 fcoe-trusted
set ethernet-switching-options secure-access-port interface xe-0/0/31 fcoe-trusted
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000
```

To quickly configure VN2VN_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS3:

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set ethernet-switching-options secure-access-port interface xe-0/0/11 fcoe-trusted
set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2
beacon-period 90000
```

Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS1

Step-by-Step Procedure

To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:

1. Configure the port modes of the interfaces that connect directly to the FCoE host with ENode1 (xe-0/0/20) and to aggregation layer FCoE transit switch TS2 (xe-0/0/21):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching port-mode
tagged-access
set interfaces xe-0/0/21 unit 0 family ethernet-switching port-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members
vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the network-facing port (**xe-0/0/21**) as an FCoE trusted port:

```
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/21
fcoe-trusted
```

4. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

5. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip
examine-vn2v2 beacon-period 90000
```

Configuring VN2VN_Port FIP Snooping on Aggregation Layer FCoE Transit Switch TS2

Step-by-Step Procedure

To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, set the network-facing ports as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:

1. Configure the port modes of the interfaces that connect directly to FCoE transit switches TS1 (**xe-0/0/31**) and TS3 (**xe-0/0/30**). Both interfaces are network-facing and must be configured as trunk interfaces:

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members
vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the network-facing ports (**xe-0/0/30** and **xe-0/0/31**) as FCoE trusted ports:

```
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/30
fcoe-trusted
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/31
fcoe-trusted
```

4. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

5. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip
examine-vn2v2 beacon-period 90000
```

Configuring VN2VN_Port FIP Snooping on FCoE Transit Switch TS3

- Step-by-Step Procedure** To configure interface port modes, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN_Port FIP snooping:
1. Configure the port modes of the interfaces that connect directly to the FCoE host with ENode2 (**xe-0/0/10**) and to aggregation layer FCoE transit switch TS2 (**xe-0/0/11**):


```
user@switch# set interfaces xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
```
 2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN_Port VLAN (**vlan200**):


```
user@switch# set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members vlan200
```
 3. Configure the network-facing port (**xe-0/0/11**) as an FCoE trusted port:


```
user@switch# set ethernet-switching-options secure-access-port interface xe-0/0/11 fcoe-trusted
```
 4. Configure the FCoE VLAN dedicated to VN2VN_Port FIP snooping:


```
user@switch# set vlans vlan200 vlan-id 200
```
 5. Enable VN2VN_Port FIP snooping on the VLAN and configure the beacon period:


```
user@switch# set ethernet-switching-options secure-access-port vlan vlan200 examine-fip examine-vn2v2 beacon-period 90000
```

Verification

To verify that the VN2VN_Port FIP snooping configuration has been created and is operating properly on all three switches, perform these tasks:

- [Verifying That VN2VN_Port FIP Snooping Is Enabled on the FCoE VLAN \(All Three Transit Switches\)](#) on page 5172
- [Verifying the Interface Port Mode](#) on page 5174

Verifying That VN2VN_Port FIP Snooping Is Enabled on the FCoE VLAN (All Three Transit Switches)

- Purpose** Verify that VN2VN_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and that the correct interfaces (**xe-0/0/20** and **xe-0/0/21** on TS1, **xe-0/0/30** and **xe-0/0/31** aggregation layer TS2, and **xe-0/0/10** and **xe-0/0/11** on TS3) are members of the VLAN.

Action List the FIP snooping information on transit switch TS1 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
```

List the FIP snooping information on aggregation layer transit switch TS2 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/30
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/31
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
```

List the FIP snooping information on transit switch TS3 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
```

```
Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/10
Active VN_Ports : 1
VN_Port Information
VN-Port MAC: 0e:fd:00:00:0b:01
Active Sessions : 1
Session Information
Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/11
Active VN_Ports : 1
VN_Port Information
VN-Port MAC: 0e:fd:00:00:0a:01
Active Sessions : 1
Session Information
Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
```

Meaning The **show fip snooping detail** command lists all of the transit switch information about VN2VN_Port FIP snooping and VN2VF_Port FIP snooping on each transit switch. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF_Port FIP snooping.)
- The beacon period is **90000**.
- The interfaces connected to the ENodes are **xe-0/0/20** and **xe-0/0/21** on transit switch TS1, **xe-0/0/30** and **xe-0/0/31** on aggregation layer transit switch TS2, and **xe-0/0/10** and **xe-0/0/11** on transit switch TS3. Because the transit switches are transparent passthrough switches, the network-facing trunk ports “see” the FCoE host ENodes at the far end of the VN2VN_Port virtual link.

In addition, this useful command shows information about the ENodes and the VN2VN_Port sessions.

Verifying the Interface Port Mode

Purpose Verify that the interface port modes are **tagged-access** for ENode-facing ports and **trunk** for network-facing ports on each transit switch.

Action List the Ethernet switching interfaces to confirm the port mode using the **show ethernet-switching interfaces detail** operational command for each interface. The output is truncated to show only the relevant portions.

List the Ethernet switching interface information on FCoE transit switch TS1 using the operational mode commands **show ethernet-switching interfaces xe-0/0/20.0 detail** and **show ethernet-switching interfaces xe-0/0/21.0 detail**:

```
user@switch> show ethernet-switching interfaces xe-0/0/20.0 detail
Interface: xe-0/0/20.0, Index: 75, State: up, Port mode: Tagged-Access
.
.
.
user@switch> show ethernet-switching interfaces xe-0/0/21.0 detail
```

```
Interface: xe-0/0/21.0, Index: 83, State: up, Port mode: Trunk
.
.
.
```

List the Ethernet switching interface information on aggregation layer FCoE transit switch TS2 using the operational mode commands **show ethernet-switching interfaces xe-0/0/30.0 detail** and **show ethernet-switching interfaces xe-0/0/31.0 detail**:

```
user@switch> show ethernet-switching interfaces xe-0/0/30.0 detail
Interface: xe-0/0/30.0, Index: 71, State: up, Port mode: Trunk
.
.
.
```

```
user@switch> show ethernet-switching interfaces xe-0/0/31.0 detail
Interface: xe-0/0/31.0, Index: 73, State: up, Port mode: Trunk
.
.
.
```

List the Ethernet switching interface information on FCoE transit switch TS3 using the operational mode commands **show ethernet-switching interfaces xe-0/0/10.0 detail** and **show ethernet-switching interfaces xe-0/0/11.0 detail**:

```
user@switch> show ethernet-switching interfaces xe-0/0/10.0 detail
Interface: xe-0/0/10.0, Index: 56, State: up, Port mode: Tagged-Access
.
.
.
```

```
user@switch> show ethernet-switching interfaces xe-0/0/11.0 detail
Interface: xe-0/0/11.0, Index: 59, State: up, Port mode: Trunk
.
.
.
```

Related Documentation

- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5155](#)
- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
- [Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5202](#)
- [Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)

Example: Configuring Automated Fibre Channel Interface Load Rebalancing

Automated Fibre Channel (FC) interface (NP_Port) load rebalancing configures the switch to rebalance the session loads on the native FC interfaces automatically on a load-rebalancing trigger event. (Alternatively, you can rebalance the link load on the FC interfaces on demand so that you control when the link load is rebalanced.) Rebalancing the FC link load is a disruptive action that causes some or all of the current sessions to log out, then log in again to be placed on the active FC links in a balanced manner.

This example shows you how to configure and verify automated FC link load rebalancing on an FCoE-FC gateway local FC fabric.

- [Requirements on page 5176](#)
- [Overview on page 5176](#)
- [Configuration on page 5177](#)
- [Verification on page 5178](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

Overview

When a load rebalancing trigger occurs, the switch automatically rebalances the link loads by redistributing the sessions across the active NP_Port links.

There are three possible load-rebalancing triggers:

1. When you enable automated load rebalancing, the switch checks the load balance on the existing NP_Port links. If the links are already balanced, the switch does not rebalance the link load. If the links are not balanced, the switch rebalances the link loads using the configured load-balancing algorithm.
2. When a new NP_Port link comes up on a local FCoE-FC gateway fabric, the switch rebalances the link load using the configured load-balancing algorithm if automated load balancing is enabled.
3. When the port speed is changed (unless the port speed change does not change the actual port speed, for example, changing the port speed from auto to 8 Gbps).

Automated load rebalancing logs out sessions in accordance with the configured load-balancing algorithm. Disabling automated load rebalancing is not disruptive because the link load is already balanced.

Use automated load rebalancing if you want link loads to be rebalanced automatically instead of at times of your choosing. Keep in mind that load rebalancing is a disruptive event (sessions are logged out).

Topology

This example configures automated load rebalancing on a local FC fabric on an FCoE-FC gateway. This example does not show you how to configure the load-balancing algorithm or any other load-balancing characteristics. The load-balancing configuration for this example is:

- FC fabric name—`fc_fabric_100`
- FC fabric ID—100

- FC fabric type—Proxy
- FC fabric interfaces—fc-0/0/0, fc-0/0/1, fc-0/0/42, fc-0/0/43, vlan.100, vlan.20
- Load-balancing algorithm—Simple
- No fabric WWN verify—Configured
- Traceoptions—Configured to log in file fc_fabric_100_proxy.log

Configuration

To configure automated load balancing on a local FC fabric, perform this task:

- [\[xref target has no title\]](#)
- [Results on page 5177](#)

CLI Quick Configuration

To quickly configure automated load balancing, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit]
set fc-fabrics fc_fabric_100 proxy auto-load-rebalance
```

Step-by-Step Procedure

- Configure automated load balancing on FC fabric fc_fabric_100:
user@switch# set fc-fabrics fc_fabric_100 proxy auto-load-rebalance

Results

Display the results of the configuration:

```
user@switch> show configuration fc-fabrics
fc_fabric_100 {
    fabric-id 100;
    fabric-type proxy;
    interface {
        fc-0/0/0.0;
        fc-0/0/1.0;
        vlan.100;
        vlan.20;
        fc-0/0/42.0;
        fc-0/0/43.0;
    }
    proxy {
        traceoptions {
            file fc_fabric_100_proxy.log size 20m;
            flag all;
        }
        load-balance-algorithm simple;
        auto-load-rebalance;
        no-fabric-wwn-verify;
    }
}
```

Verification

Verifying That Automated Load Rebalancing Is Enabled

Purpose Verify that automated load rebalancing is configured on local FC fabric **fc_fabric_100**.

Action Verify the results of the automated load-rebalancing configuration using the operational mode command **show fibre-channel proxy fabric-state fabric fc_fabric_100**:

```
user@switch> show fibre-channel proxy fabric-state fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
Proxy load balance algorithm: Simple, Fabric WWN verification: No
Auto load rebalance enabled : Yes
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : None
Last rebalance trigger-time : Mon Sep 10 21:42:30 2012 usec: 814602
Last rebalance trigger-result: None
```

Meaning The **show fibre-channel proxy fabric-state fabric fc_fabric_100** operational command displays information about the specified local FC fabric. The output shows that the **Auto load rebalance enabled** field value is **Yes**, which indicates that automated load rebalancing is enabled on fabric **fc_fabric_100**.

Related Documentation

- [Defining the Proxy Load-Balancing Algorithm on page 5190](#)
- [Simulating On-Demand Fibre Channel Link Load Rebalancing \(Dry Run Test\) on page 5191](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

Configuration Tasks

- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
- [Disabling the Fabric WWN Verification Check on page 5180](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Deleting a Fibre Channel Interface on page 5188](#)
- [Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway on page 5189](#)
- [Defining the Proxy Load-Balancing Algorithm on page 5190](#)
- [Simulating On-Demand Fibre Channel Link Load Rebalancing \(Dry Run Test\) on page 5191](#)
- [Enabling and Disabling CoS OxID Hash Control on page 5192](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Setting the Maximum Number of FIP Login Sessions per ENode on page 5195](#)

- [Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197](#)
- [Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198](#)
- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)
- [Disabling VN2VF_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)
- [Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5202](#)
- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)

Configuring an FCoE-FC Gateway Fibre Channel Fabric

Fibre Channel (FC) fabric configuration consists of creating a unique name and identifier for each FC fabric you want to create and configuring it as an FCoE-FC gateway.

You can create a maximum of 12 FC fabrics on a QFX3500 switch. After you create a fabric, you can create and assign interfaces to the fabric, configure FIP parameters for the fabric, and set proxy traceoptions.

To configure an FC fabric using the CLI, specify a unique name and identification number for the fabric:

1. Configure the fabric name and fabric ID:

```
[edit]
user@switch# set fc-fabrics fabric-name fabric-id fabric-id
```



NOTE: Changing the fabric name or the fabric ID causes all logins to drop and forces the ENodes to log in again.

For example, to configure an FC fabric with the name **fab_ulous** and the fabric ID 10 (the range of **fabric-id** values is 1 through 4095):

```
[edit]
user@switch# set fc-fabrics fab_ulous fabric-id 10
```

2. Configure the fabric as a gateway fabric:

```
[edit fc-fabrics fabric-name]
user@switch# set fabric-type proxy
```

For example, to configure the FC fabric with the name **fab_ulous** as a gateway fabric:

```
[edit fc-fabrics fab_ulous]
user@switch# set fabric-type proxy
```

Related Documentation

- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)

Disabling the Fabric WWN Verification Check

When a QFX Series NP_Port sends a fabric login (FLOGI) request to a Fibre Channel (FC) switch, the FLOGI accept (FLOGI-ACC) reply from the FC switch contains the SAN fabric worldwide name (WWN). The QFX Series uses the SAN fabric WWN in the multicast discovery advertisement (MDA) that the QFX Series sends to the ENodes in the FCoE network.

However, some FC switches substitute their own WWN (often the FC switch's virtual WWN) for the SAN fabric WWN in the FLOGI-ACC message. In this case, different NP_Ports that log in to the same FC fabric might receive different fabric WWNs in the FLOGI-ACC messages if the NP_Ports are connected to different FC switches in the SAN fabric.

If the QFX Series receives different fabric WWNs on NP_Ports that are connected to the same SAN fabric, the QFX Series uses the first fabric WWN it receives in the MDA it sends to the ENodes. The QFX Series isolates the NP_Ports that receive a different fabric WWN from other FC switches in that SAN fabric. No ENode sessions are assigned to the isolated NP_Ports. FC traffic is assigned only to NP_Ports that receive a fabric WWN in the FLOGI-ACC message that matches the fabric WWN received by the first NP_Port to log in to the FC fabric. (If an NP_Port receives a fabric WWN that does not match the fabric WWN received by the first NP_Port to log in to the FC fabric, it does not carry traffic to the SAN fabric.)

To prevent ENodes from being isolated due to a mismatched fabric WWN, you can disable the gateway fabric WWN verification check. Disabling the fabric WWN verification check enables all NP_Ports connected to a SAN fabric are used to carry traffic between the gateway and the FC switch, regardless of the fabric WWN the NP_Port receives in the FLOGI-ACC message.



NOTE: Disabling or enabling the fabric WWN verification check logs out all FCoE sessions.

To disable the fabric WWN verification check:

- `[edit fc-fabrics fabric-name proxy]`
`user@switch# set no-fabric-wwn-verify`

- Related Documentation**
- [Understanding FCoE-FC Gateway Functions on page 4979](#)
 - [show fibre-channel proxy fabric-state on page 5364](#)

Configuring a Physical Fibre Channel Interface

When you configure the switch as an FCoE-FC gateway, you must configure either 6 or 12 of the physical interfaces as native FC interfaces. Native FC interfaces connect to the storage area network (SAN) FC switch.

You can configure ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to create blocks of native FC interfaces. You cannot individually configure a single port as a native FC interface. Within these port blocks, you cannot mix FC interfaces with Ethernet interfaces. All of the ports in a block must be either native FC interfaces or Ethernet interfaces.

You can configure:

- Six native FC interfaces by configuring either ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, or ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- Twelve native FC interfaces by configuring ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- No native FC interfaces by leaving ports xe-0/0/0 through xe-0/0/5 and ports xe-0/0/42 through xe-0/0/47 in their default state as Ethernet interfaces.
- To configure physical FC interfaces using the CLI, specify the physical port block you want to configure on the switch as native FC interfaces:

```
[edit chassis]
user@switch# set fpc fpc pic pic fibre-channel port-range port-range-low port-range-high
```

For example, to configure six native FC interfaces, you can configure ports 0 through 5 as physical FC interfaces:

```
[edit chassis]
user@switch# set fpc 0 pic 0 fibre-channel port-range 0 5
```

To configure 12 native FC interfaces requires two separate statements:

```
[edit chassis]
user@switch# set fpc 0 pic 0 fibre-channel port-range 0 5
user@switch# set fpc 0 pic 0 fibre-channel port-range 42 47
```

- Related Documentation**
- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
 - [Configuring an FCoE VLAN Interface on page 5185](#)
 - [Configuring a Fibre Channel Interface on page 5182](#)
 - [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
 - [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)

- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Configuring a Fibre Channel Interface

When a QFX3500 acts as an FCoE-FC gateway, native Fibre Channel (FC) traffic flows between the switch and the storage area network (SAN) FC switch. When you configure a port as an FC interface, it transports only FC traffic. It does not transport Ethernet traffic.

You can configure ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to create blocks of native FC interfaces.

Each of these blocks of ports must be configured either as all Ethernet ports or as all native FC ports. Within each block of ports, you cannot mix FC and Ethernet interfaces. This means that you can configure 0, 6, or 12 ports as native FC ports. “[Configuring a Physical Fibre Channel Interface](#)” on page 5181 describes how to configure the port blocks as physical FC interfaces.



NOTE: Do not configure ports that you want to use for native FC traffic as part of an Ethernet VLAN or as Ethernet ports.

Configure a port as an FC interface when the port connects to the F_Port of an FC switch.

FC interface configuration includes:

- Explicitly specifying one or more ports as an FC family interface in NP_Port mode (mandatory).
- Configuring the FC interface options port speed and buffer-to-buffer credit state change number (BB_SC_N) (optional).
- Configuring the interface as a loopback interface (optional).

The buffer-to-buffer state change number feature prevents the loss of buffer-to-buffer credits between the two interfaces on either end of an FC link. The state change number determines the number of frames and receiver ready (R_RDY) primitives the interfaces exchange between the state change send (BB_SCs) and the state change receive (BB_SCr) primitives used to track these transactions.

Enabling BB_SC_N by configuring BB_SC_N on both of the FC link interfaces:

- Requests that $2^{BB_SC_N}$ number of frames be sent between two consecutive BB_SCs primitives, and
- Requests that $2^{BB_SC_N}$ number of R_RDY primitives be sent between two consecutive BB_SCr primitives.

When the number of R_RDY primitives received equals $2^{BB_SC_N}$, the R_RDY counter resets to zero. When the number of frames received equals $2^{BB_SC_N}$, the frame counter resets

to zero. The interfaces calculate the number of buffer-to-buffer credits lost based on counter discrepancies and take corrective action to recover the lost credits.

If you enable BB_SC_N, the recommended BB_SC_N setting is eight. Setting the BB_SC_N number to zero (0) disables the feature. If either of the two connected FC interfaces is configured with zero as the BB_SC_N value, then both interfaces disable the feature. If the two connected FC interfaces have different nonzero BB_SC_N numbers configured, both interfaces use the higher number.

For the port to transport FC traffic, you must also set the physical port as an FC port using the **port-range** command.

To configure an FC interface using the CLI:

1. Specify the interface as family FC and set the port mode to NP_Port (setting the port mode to NP_Port is a mandatory configuration):

```
[edit]
user@switch# set interfaces interface-name unit unit family fibre-channel port-mode np-port
```

For example, to configure the interface **fc-0/0/3** as an FC interface and set the port mode to **np-port**:

```
[edit]
user@switch# set interfaces fc-0/0/3 unit 0 family fibre-channel port-mode np-port
```

2. Configure the FC interface speed option:

```
[edit]
user@switch: set interfaces interface-name fibrechannel-options speed (auto-negotiation | 2g | 4g | 8g)
```

For example, to set the FC interface speed option to **8g** for the interface **fc-0/0/3**:

```
[edit]
user@switch: set interfaces fc-0/0/3 fibrechannel-options speed 8g
```

The default port mode is **auto-negotiation**, which sets the port speed to match the speed of the attached FC F_Port interface (2 Gbps, 4 Gbps, or 8 Gbps).

3. Configure the optional buffer-to-buffer credit state change number:

```
[edit]
user@switch: set interfaces interface-name fibrechannel-options bb-sc-n 0..15
```

For example, to set the FC interface buffer-to-buffer credit state change number to **8** for the interface **fc-0/0/3**:

```
[edit]
user@switch: set interfaces fc-0/0/3 fibrechannel-options bb-sc-n 8
```

After you configure one or more FC interfaces, assign them and an FCoE VLAN to an FC fabric.

Related Documentation

- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Deleting a Fibre Channel Interface on page 5188](#)

- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)

Configuring an FCoE VLAN Interface

When you configure the switch as an FCoE-FC gateway, a Layer 3 FCoE VLAN interface transmits and receives Fibre Channel over Ethernet (FCoE) traffic between the gateway and FCoE-capable servers on the Ethernet network. Configuring a Layer 3 FCoE VLAN interface on the switch creates virtual fabric port (VF_Port) interfaces facing the FCoE server virtual node ports (VN_Ports).

The FCoE VLAN interface is the interface for the dedicated VLAN the FCoE servers use for FCoE traffic. Each FC fabric requires at least one dedicated FCoE VLAN and at least one Layer 3 FCoE VLAN interface to transport FCoE traffic. On QFabric systems, the FCoE VLAN interface, the FCoE VLAN, and the interfaces that are members of the FCoE VLAN must be on the same Node device.



NOTE: FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

Before you configure an FCoE VLAN interface, create the FCoE VLAN and assign 10-Gigabit Ethernet interfaces configured in tagged-access port mode to the VLAN. These 10-Gigabit Ethernet interfaces are the physical interfaces that transport the FCoE traffic to and from the FCoE devices in the Ethernet network.

Each Ethernet interface that connects to FCoE devices must also include the native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets. The FCoE VLAN must carry only FCoE traffic. A VLAN cannot transport a mix of FCoE and standard Ethernet traffic.

FCoE VLAN interface configuration includes:

- Configuring a VLAN to use as a dedicated FCoE VLAN.
- Configuring a native VLAN for FIP traffic.
- Configuring member interfaces for the FCoE VLAN.
- Configuring the FCoE VLAN as a Fibre Channel (family) VLAN and setting the port mode value to **f-port**. Explicitly configuring the FCoE VLAN interface in F_Port mode is mandatory. The switch interface with which the FCoE server VN_Ports communicate must present a VF_Port to the servers.
- Configuring the FCoE VLAN interface as the Layer 3 interface for FCoE traffic.

To configure an FCoE VLAN interface:

1. Configure a dedicated FCoE VLAN:

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

For example, to configure a VLAN named **fcoe_vlan** with a VLAN ID of **100** as the FCoE VLAN:

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

2. Configure a native VLAN for FIP traffic:

```
[edit vlans]
user@switch# set native vlan-id vlan-id
```

For example, to configure the native VLAN with a VLAN ID of 1:

```
[edit vlans]
user@switch# set native vlan-id 1
```

3. Configure member interfaces for the FCoE VLAN (use **ethernet-switching** as the family and **tagged-access** as the port mode):

```
[edit interfaces]
user@switch# set interface-name unit unit family family port-mode mode vlan members
vlan-name
```

For example, to configure the interface **xe-0/0/10** as a member of the FCoE VLAN **fcoe_vlan**:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access vlan
members fcoe_vlan
```

4. Assign member interfaces to the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit unit family family native-vlan-id vlan-id
```

For example, to configure the interface **xe-0/0/10** as a member of the native VLAN with the native VLAN ID 1:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
```

5. Assign the Ethernet interfaces to the FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name interface interface-name
```

For example, to assign the interface **xe-0/0/10.0** to the FCoE VLAN named **fcoe_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan interface xe-0/0/10.0
```

6. Define an interface as an FCoE VLAN interface in F_Port mode (to present a VF_Port to the FCoE servers):

```
[edit interfaces]
user@switch# set vlan unit unit family fibre-channel port-mode f-port
```

For example, to configure VLAN unit **100** as an FCoE VLAN interface and set the port mode to **f-port**:

```
[edit interfaces]
user@switch# set vlan unit 100 family fibre-channel port-mode f-port
```

7. Define the Layer 3 FCoE VLAN interface:

```
[edit vlans]
user@switch# set vlan-name l3-interface vlan-interface-name
```

For example, to configure VLAN interface unit **100** (the FCoE VLAN interface defined earlier in this example) as the Layer 3 FCoE VLAN interface for FCoE VLAN **fcoe_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan l3-interface vlan.100
```

Related Documentation

- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Configuring a Fibre Channel Interface on page 5182](#)
- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Disabling VN2VF_Port FIP Snooping on an FCoE-FC Gateway Switch Interface on page 5201](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)

Assigning Interfaces to a Fibre Channel Fabric

When you configure the switch as an FCoE-FC gateway, you assign one or more (up to 12) native Fibre Channel (FC) interfaces and at least one FCoE VLAN interface to each FC fabric. FC interfaces transport native FC traffic between the proxy gateway and the storage area network (SAN) FC switch. FCoE VLAN interfaces transport FCoE traffic between FCoE-capable servers and the gateway.

Each FC fabric needs both types of interfaces to transport traffic between FCoE servers on the Ethernet network and FC storage devices in the core FC network behind the FC switch. FCoE traffic between the FCoE servers and the gateway must travel in a dedicated FCoE VLAN. Native FC traffic passes between the gateway and the FC switch on the native FC interfaces.

You must configure the FC interfaces and the FCoE VLAN interfaces that you assign to a particular fabric on the same Juniper Networks QFX3500 Switch. Traffic between an FCoE device and the FC switch must ingress and egress the same gateway.

To assign core-facing native FC interfaces and a server-facing FCoE VLAN interface to an FC fabric, configure a fabric and then specify the interfaces:

1. Assign the native FC interfaces to the FC fabric:

```
[edit fc-fabrics fabric-name]
user@switch: set interface interface-name
user@switch: set interface interface-name
user@switch: set interface interface-name
...
```

2. Assign an FCoE VLAN interface to the FC fabric:

```
[edit fc-fabrics fabric-name]
user@switch: set interface vlan-name
```

For example, to assign the native FC interfaces **fc-0/0/0.0**, **fc-0/0/1.0**, and **fc-0/0/2.0** and the FCoE VLAN interface **vlan.100** to an FC fabric named **san_tana**:

```
user@switch: set fc-fabrics san_tana interface fc-0/0/0.0
user@switch: set fc-fabrics san_tana interface fc-0/0/1.0
user@switch: set fc-fabrics san_tana interface fc-0/0/2.0
user@switch: set fc-fabrics san_tana interface vlan.100
```

Related Documentation

- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)

Deleting a Fibre Channel Interface

Before you delete a Fibre Channel (FC) interface, you must first delete the interface from the FC fabric configuration. This prevents configuration errors that would result if you deleted an FC interface from the **[edit interfaces]** hierarchy level but did not delete the interface from the FC fabric.

When you configure the switch as an FCoE-FC gateway, FC interfaces transmit and receive native FC traffic between the gateway and the FC switch. You can configure ports **xe-0/0/0** through **xe-0/0/5** as **fc-0/0/0** through **fc-0/0/5** and ports **xe-0/0/42** through **xe-0/0/47** as **fc-0/0/42** through **fc-0/0/47** to create one or two blocks of six native FC interfaces.

To delete an FC interface using the CLI:

1. Delete the FC interface from the FC fabric to which it belongs:

```
[edit]
user@switch# delete fc-fabrics fabric-name interface interface-name
```

For example, to delete the FC interface **fc-0/0/3.0** from an FC fabric named **sanfab1**:

```
[edit]
user@switch# delete fc-fabrics sanfab1 interface fc-0/0/3.0
```

2. Delete the FC interface from the switch **[edit interfaces]** hierarchy:

```
[edit]
user@switch: delete interfaces interface-name
```

For example, to delete the interface **fc-0/0/3.0** from the switch:

```
[edit]
user@switch: delete interfaces fc-0/0/3.0
```

The FC interface has been deleted from the FC fabric and from the switch.

Related Documentation

- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)

- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway

Storm control is not supported on the FCoE interfaces of an FCoE-FC gateway VLAN. Enabling storm control on an FCoE-FC gateway VLAN interface may cause FCoE packet loss. Storm control is disabled by default on all interfaces. However, if you enabled storm control globally on all switch interfaces or on any interfaces that are part of the FCoE VLAN interface, you must disable storm control on the Ethernet interfaces of the FCoE VLAN.

If storm control is enabled on only a few interfaces of the FCoE VLAN, you can disable storm control on individual interfaces by including the **set ethernet-switching-options storm-control interface *interface-name*** statement in the configuration, where *interface-name* is the name of the interface on which you want to disable storm control.

If storm control is enabled globally on the switch when the switch is acting as an FCoE-FC gateway, it is often easiest to disable storm control on all interfaces, then enable storm control only on Ethernet interfaces that are not part of the FCoE VLAN interface.

If storm control is enabled globally, you can disable storm control in either of two ways:

- Disable storm control on all interfaces, then enable storm control on the interfaces you want to use storm control. (From the default configuration, you cannot disable storm control on individual interfaces because the default configuration enables storm control on **all** interfaces, not on individual interfaces.)

For example, if you want interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22 to use storm control, disable storm control on all interfaces, then enable storm control on those three interfaces:

1. Disable storm control on all interfaces:

```
user@switch# delete ethernet-switching-options storm-control interface all
```

2. Enable storm control on interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22:

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/20
```

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/21
```

```
user@switch# set ethernet-switching-options storm-control interface xe-0/0/22
```

- Disable storm control for all unknown unicast traffic on all interfaces by including the following statement in your configuration:

```
user@switch# set ethernet-switching-options storm-control interface all no-unknown-unicast
```

Related Documentation

- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)

Defining the Proxy Load-Balancing Algorithm

When the QFX Series is configured as an FCoE-FC gateway, it balances the FCoE session load assigned to each NP_Port link between the gateway and the FC switch in the FC SAN to avoid overloading or underutilizing each link. The QFX Series supports three types of load-balancing mechanisms:

- **Simple load balancing**—Load balancing is based on the weighted utilization (session load) of the NP_Ports connected to an FC fabric. The session load is the sum of the FLOGI and FDISC sessions on each link. Each new ENode fabric login (FLOGI) or VN_Port fabric discovery (FDISC) session is assigned to the least-loaded link, so an FDISC session initiated by the VN_Port on an ENode might not be assigned to the same link as the parent ENode's FLOGI session. Simple load balancing is the default algorithm. Simple load balancing is the default load-balancing algorithm. Rebalancing the link load disrupts only selected sessions to minimize the impact (the switch uses an algorithm to log out only the sessions that need to be moved to other links to balance the load when those sessions log in again).
- **ENode-based load balancing**—Load balancing is based on the weighted utilization (session load) of the NP_Ports connected to an FC fabric. The session load is the sum of the FLOGI and FDISC sessions on each link. However, when an ENode logs in to the fabric, the switch places all subsequent VN_Port FDISC sessions associated with that ENode on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. The switch calculates the link load based on the combined total of FLOGIs and FDISCs on each NP_Port link. Rebalancing the link load disrupts all sessions (all sessions log out and then log in again).
- **FLOGI-based load balancing**—Load balancing is based on the weighted utilization (session load) of the NP_Ports connected to an FC fabric. The session load is the sum of the FLOGI sessions on each link. FDISC sessions are not counted. When an ENode logs in to the fabric, the switch places all subsequent VN_Port FDISC sessions associated with that ENode on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. Rebalancing the link load disrupts only selected sessions to minimize the impact (the switch uses an algorithm to log out only the sessions that need to be moved to other links to balance the load when those sessions log in again).

To define the proxy load-balancing algorithm for a proxy fabric on the FCoE-FC gateway, set the algorithm as **enode-based**, **simple**, or **flogi-based**:

- `[edit fc-fabrics fabric-name proxy]`
`user@switch# set load-balance-algorithm (enode-based | simple | flogi-based)`

For example, to configure a gateway fabric named **san_fab1** to use **enode-based** load balancing:

```
user@switch# set fc-fabrics san_fab1 proxy load-balance-algorithm enode-based
```

Related Documentation

- [Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175](#)
- [Simulating On-Demand Fibre Channel Link Load Rebalancing \(Dry Run Test\) on page 5191](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)

- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test)

On-demand Fibre Channel (FC) link load rebalancing is a disruptive action that causes sessions to log out of the network, then log back in to be placed on FC links (NP_Ports) in a balanced manner. The number of sessions logged out to rebalance the links depends on the load-balancing algorithm used (simple, ENode-based, or FLOGI-based) and whether or not the load is already balanced. (If the link load is already balanced, the switch does not rebalance the loads when you request on-demand load rebalancing.)

You can use the **dry-run** option to list the sessions that might be affected (logged out to be redistributed among the active FC interface links) by on-demand load rebalancing *before* you actually rebalance the link load. (Because new sessions might log in between the time you perform a dry run and the time you request on-demand load rebalancing, the affected sessions may change. Therefore, the sooner that you perform an on-demand load rebalance after you perform a dry run, the more accurate the dry run results are likely to be.)

To request a link load rebalancing dry run:

```
user@switch> request fibre-channel proxy load-rebalance dry-run fabric fabric-name
```

For example, to request a dry run on an FC fabric named *fc_fabric_100* to display a list of sessions that might be disrupted if you request an actual link load rebalance:

```
user@switch> request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
F-Port      FCID      Port-WWN      NP-Port
vlan.100    0x8a013a  02:01:00:64:00:00:2a  fc-0/0/1.0
vlan.100    0x8a013c  02:01:00:64:00:00:2b  fc-0/0/1.0
vlan.100    0x8a0146  02:01:00:64:00:00:2e  fc-0/0/1.0
vlan.100    0x8a014c  02:01:00:64:00:00:2f  fc-0/0/1.0
```

Related Documentation

- [request fibre-channel proxy load-rebalance on page 5286](#)
- [Defining the Proxy Load-Balancing Algorithm on page 5190](#)
- [Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

Enabling and Disabling CoS OxID Hash Control

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links. You can configure whether or not the switch uses the OxID in the hash computation.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, on the QFX3500, you can assign different IEEE priorities to different lossless FCoE flows as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) to further separate the traffic flows.

OxID hash control is enabled by default.

- To enable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid enable
```

- To disable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid disable
```

Related Documentation

- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on page 5024](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

Configuring FIP on an FCoE-FC Gateway

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices. A virtual link emulates the physical point-to-point link that FC requires between two FC devices.

FIP is enabled by default and uses the default FIP settings on all FCoE interfaces that are part of the gateway FC fabric. You can use the default FIP parameter values, or you can configure FIP parameters globally or on a per-interface basis. Configuring FIP on an individual interface overrides the global FIP configuration.

You can configure the following parameters globally for the fabric and per interface:

- FIP keepalive message transmission interval—This interval is the time period between sending FIP keepalive messages.

- **Priority**—If an FCoE node (ENode) connects to more than one switch, the priority value determines the switch to which the ENode connects. The switch with the lowest priority number has the highest priority.

You can only configure the following parameters globally on an FC fabric:

- **FC-MAP**—The 24-bit FCoE mapped address prefix that identifies the attached FC switch in the SAN fabric. The FC-MAP value is used in the fabric provided MAC address (FPMA) created for each ENode that logs in. This value must be the same for the FC switch and the QFX Series.



NOTE: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

- **FCoE trusted**—You can globally configure all of the Ethernet ports in a specified FC fabric to be FCoE trusted. You might want to configure interfaces as FCoE trusted if the interfaces are connected to a transit switch that is performing FIP snooping. For interfaces that are directly connected to FCoE hosts, FIP snooping should be enabled, and you should not configure the fabric as FCoE trusted.



NOTE: Do not configure interfaces with FIP snooping enabled as FCoE trusted.

Configuring interfaces as FCoE trusted reduces system overhead by eliminating the need for filters. The total number of sessions the system can support is 2500 sessions. Sessions are defined as the combined number of VN_Port to VF_Port sessions and VN_Port to VN_Port sessions. (Although VN2VF and VN2VN sessions run in different FCoE VLANs, the session limit is a system limit, not a per-VLAN limit.)



NOTE: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions. There is no limit to the number of end-to-end storage sessions.



NOTE: Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.

- **Maximum number of FCoE sessions per ENode**—You can globally configure the maximum number of FCoE sessions (FLOGI plus FDISC) permitted from an ENode. The maximum number of sessions per ENode is 2000 sessions. The total number of sessions (VN2VF_Port sessions and VN2VN_Port sessions combined) cannot exceed the gateway fabric's maximum limit of 2500 sessions.

To configure FIP options globally using the CLI:

1. Specify the fabric on which you want to configure FIP:

```
[edit]
user@switch# set fc-fabrics fabric-name protocols fip
```

2. Configure the FIP keepalive message transmission interval in milliseconds to specify the amount of time between periodic FIP discovery advertisements for the fabric interfaces (the default is 8000 ms; the range is 250 through 90000 ms):

```
[edit fc-fabrics fabric-name protocols fip]
user@switch# set fka-adv-period milliseconds
```

3. Configure the priority value the switch advertises to ENodes in the range from 0 through 255; the default value is 128:

```
[edit fc-fabrics fabric-name protocols fip]
user@switch# set priority priority
```

4. Configure the FC-MAP value to match the FC-MAP value of the attached FC switch in the FC SAN fabric; the range of possible values is 0EFC00 through 0EFCFF, and the default value is 0EFC00:

```
[edit fc-fabrics fabric-name protocols fip]
user@switch# set fc-map fc-map
```

5. Configure the interfaces in the FC fabric as FCoE trusted (in this example, we assume that the interfaces have not been enabled for FIP snooping):

```
[edit fc-fabrics fabric-name protocols fip]
user@switch# set fcoe-trusted
```

6. Configure the maximum number of FCoE sessions for each ENode in the fabric:

```
[edit fc-fabrics fabric-name protocols fip]
user@switch# set max-sessions-per-enode
```

For example, to configure all FCoE interfaces associated with an FC fabric called **movieco_san** with a FIP keepalive interval value of **25000** milliseconds, a priority of **70**, an FC-MAP value of **0EFC01**, as FCoE trusted, and with a maximum number of FCoE sessions per ENode of 200 sessions:

```
[edit fc-fabrics movieco_san protocols fip]
user@switch# set fka-adv-period 25000
user@switch# set priority 70
user@switch# set fc-map 0EFC01
user@switch# set fcoe-trusted
user@switch# set max-sessions-per-enode 200
```

To override the global FC fabric FIP configuration for a specific FCoE interface using the CLI:

1. Specify the fabric and interface on which you want to configure FIP:

```
[edit fc-fabrics fabric-name protocols fip interface interface-name]
```

2. Configure the FIP keepalive message transmission interval and priority:

```
[edit fc-fabrics fabric-name protocols fip interface interface-name]
user@switch# set fka-adv-period milliseconds
user@switch# set priority priority
```

- Related Documentation**
- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
 - [Configuring an FCoE VLAN Interface on page 5185](#)
 - [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
 - [Understanding FIP Parameters on an FCoE-FC Gateway on page 4992](#)

Setting the Maximum Number of FIP Login Sessions per ENode

When the switch acts as an FCoE-FC gateway, FCoE node (ENode) devices in the Ethernet network use the gateway to connect to the Fibre Channel (FC) storage area network (SAN). You can limit the maximum number of FIP login sessions permitted on each ENode. Limiting the number of login sessions can prevent login session rejections caused when the connected FC switch port configuration limits the number of FIP login sessions.

The maximum number of FIP sessions per ENode is 2000 sessions (FLOGI plus FDISC sessions). The limit you set applies to every ENode in the specified gateway fabric. Each ENode in the fabric can have up to the maximum number of sessions, but the total number of active sessions cannot exceed the session limits you apply to the fabric or the Node device.

There are also configurable FIP login session limits that you can apply to the gateway FC fabric, to the QFX3500 switch or QFabric system Node device, and to the interfaces in each FC fabric.

- To set a maximum number of FIP login sessions per ENode using the CLI:

```
[edit fc-fabrics fc-fabric-name protocols fip]
user@switch# set max-sessions-per-enode max-login-sessions
```

For example, to configure the ENodes on an FC fabric named **sanfab1** with a maximum FIP login session limit of **250** sessions:

```
[edit fc-fabrics sanfab1]
user@switch# set protocols fip max-sessions-per-enode 250
```

- Related Documentation**
- [Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196](#)
 - [Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197](#)
 - [Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198](#)
 - [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Setting the Maximum Number of FIP Login Sessions per FC Interface

When the switch acts as an FCoE-FC gateway, NP_Ports are the native FC interfaces the gateway uses to connect to the FC switch. You can limit the maximum number of FIP login sessions permitted on an NP_Port interface. Limiting the number of login sessions on an interface can prevent login session rejections caused when the connected FC switch port configuration limits the number of FIP login sessions.



TIP: A good practice is to configure a maximum number of login sessions on each NP_Port that is less than or equal to the maximum number of login sessions permitted on the connected FC switch port.

The maximum number of FIP sessions is 2500 sessions. (This is the combined total of all VN2VF_Port and VN2VN_Port sessions on the system.)

There are also configurable FIP login session limits that you can apply to the gateway FC fabric, to the QFX3500 switch or QFabric system Node device, and to the ENodes in each FC fabric. To prevent unexpected FIP login rejections, the sum of the maximum FIP login sessions on all of the NP_Port interfaces that belong to an FC fabric should not exceed the maximum number of sessions the FC fabric supports or the device supports.

- To set a maximum number of FIP login sessions on an NP_Port using the CLI:

```
[edit fc-fabrics fc-fabric-name interface interface-name]  
user@switch# set max-login-sessions max-login-sessions
```

For example, to configure NP_Port interface **fc-0/0/5** with a maximum FIP login session limit of **500** sessions on an FC fabric named **sanfab1**:

```
[edit fc-fabrics sanfab1]  
user@switch# set interface fc-0/0/5 max-login-sessions 500
```

Related Documentation

- [Setting the Maximum Number of FIP Login Sessions per ENode on page 5195](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197](#)
- [Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Setting the Maximum Number of FIP Login Sessions per FC Fabric

When the QFX Series acts as an FCoE-FC gateway, you configure at least one local FC fabric on the gateway. A gateway FC fabric creates associations that connect FCoE devices on an Ethernet network to an FC switch on a Fibre Channel network. Each FC fabric on a gateway includes native FC interfaces (NP_Ports) that connect the gateway to the FC switch. When FCoE devices want to log in to the FC switch, the gateway sends the FIP login requests to the FC switch on the NP_Port links.

You can limit the maximum number of FIP login sessions permitted on a gateway FC fabric. If a QFX3500 switch or QFabric system Node device has more than one FC fabric, limiting the number of login sessions on an FC fabric can prevent one FC fabric from using all of the login sessions available on the device.

The maximum number of FIP sessions is 2500 sessions. (This is the combined total of all VN2VF_Port and VN2VN_Port sessions on the system.)

There are also configurable FIP login session limits that you can apply to the FC fabric NP_Port interfaces, to the QFX3500 switch or QFabric system Node device, and to the ENodes in each FC fabric. To prevent unexpected FIP login rejections:

- The sum of the maximum FIP login sessions on all of the NP_Port interfaces that belong to an FC fabric should not exceed the maximum number of sessions the FC fabric supports or the device supports.
- The sum of the maximum FIP login sessions on all of the FC fabrics on a device should not exceed the maximum number of sessions per device.
- To set a maximum number of FIP login sessions on an FC fabric using the CLI:

```
[edit fc-fabrics fc-fabric-name]  
user@switch# set max-login-sessions max-login-sessions
```

For example, to configure an FC fabric named **sanfab1** with a maximum FIP login session limit of **2000** sessions:

```
[edit fc-fabrics sanfab1]  
user@switch# set fc-fabrics sanfab1 max-login-sessions 2000
```

Related Documentation

- [Setting the Maximum Number of FIP Login Sessions per ENode on page 5195](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196](#)
- [Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Setting the Maximum Number of FIP Login Sessions per Node Device

When a QFX3500 switch or QFabric system Node device acts as an FCoE-FC gateway, it connects FCoE devices on an Ethernet network to an FC switch in a Fibre Channel network. You can limit the maximum number of FIP login sessions for the FCoE devices on each Node device.

For QFX3500 switches, the maximum limit means that the sum of the FIP login sessions on all of the local FC fabrics on that QFX3500 switch cannot exceed the device maximum.

For the QFabric system, the limit applies to each Node device in the QFabric system. For example, if you configure a maximum FIP login session value of 2000 sessions, each Node device in the QFabric system can have a total of up to 2000 FIP login sessions running on its FC fabrics.

The maximum number of FIP sessions a device can support is 2500 sessions. (This is the combined total of all VN2VF_Port and VN2VN_Port sessions on the system.)

There are also configurable FIP login session limits that you can apply to the FC fabrics on the devices, to the NP_Port interfaces in each FC fabric, and to the ENodes in each FC fabric. To prevent unexpected FIP login rejections:

- The sum of the maximum FIP login sessions for all of the FC fabrics on a device should not exceed the maximum number of sessions per device.
- The sum of the maximum FIP login sessions on all of the NP_Port interfaces that belong to an FC fabric should not exceed the maximum number of sessions the FC fabric supports or the device supports.
- To set a maximum number of FIP login sessions for Node devices using the CLI:

[edit **fc-options**]

```
user@switch# set max-login-sessions-per-node max-login-sessions-per-node
```

For example, to configure a maximum FIP login limit of **2000** sessions on a QFX3500 switch or on all Node devices in a QFabric system:

[edit **fc-options**]

```
user@switch# set max-login-sessions-per-node 2000
```

Related Documentation

- [Setting the Maximum Number of FIP Login Sessions per ENode on page 5195](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196](#)
- [Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch

VN_Port to VF_Port (VN2VF_Port) Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping uses information gathered during FIP discovery and login to create firewall filters that provide security against unauthorized access to the FC switch or FCoE forwarder (FCF) through the EX4500 or QFX Series when the switch is acting as an FCoE transit switch. The firewall filters allow only FCoE devices that succeed at logging in to the FC fabric to access the FCF through the transit switch. VN2VF_Port FIP snooping provides security for the point-to-point virtual links that connect host FCoE Nodes (ENodes) and FCFs in the FCoE VLAN by denying access to any device that does not successfully log in to the FCF.

VN2VF_Port FIP snooping is disabled by default. You enable VN2VF_Port FIP snooping on a per-VLAN basis for VLANs that carry FCoE traffic. Ensure that a VLAN that carries FCoE traffic carries only FCoE traffic, because enabling VN2VF_Port FIP snooping denies access for all other Ethernet traffic.



NOTE: All of the transit switch ports are untrusted by default. If an ENode on an FCoE device logs in to an FCF before you enable VN2VF_Port FIP snooping on the VLAN and you then enable VN2VF_Port FIP snooping, the transit switch denies traffic from the ENode because the transit switch has not snooped (learned) the ENode state. The following process automatically logs the ENode back in to the FCF to reestablish the connection:

1. VN2VF_Port FIP snooping is enabled on an FCoE VLAN on the switch.
2. The switch denies existing connections between servers and the FCF on the FCoE VLAN by filtering the FCoE traffic and FIP traffic, so no keepalive messages from the ENodes reach the FCF.
3. The FCF port timer for each ENode and for each VN_Port on each ENode expires.
4. The FCF sends each ENode whose port timer has expired a Clear Virtual Links (CVL) message.
5. The CVL message causes the ENode to log in again.

Because the FCF is a trusted source, you configure interfaces that connect to the FCF as trusted interfaces. VN2VF_Port FIP snooping continues to run on trusted interfaces so that the switch learns the FCF state.



NOTE: Do not configure ENode-facing interfaces both with FIP snooping enabled and as trusted interfaces. FCoE VLANs with interfaces that are directly connected to FCoE hosts should be configured with FIP snooping enabled and the interfaces should not be trusted interfaces. Ethernet interfaces that are connected to an FCF should be configured as trusted interfaces and should not have FIP snooping enabled. Interfaces that are connected to a transit switch that is performing FIP snooping can be configured as trusted interfaces if the FCoE VLAN is not enabled for FIP snooping.

Optionally, you can specify an FC-MAP value for each FCoE VLAN. On a given FCoE VLAN, the switch learns only FCFs that have a matching FC-MAP value. The default FC-MAP value is 0EFC00h for all FC devices. (Enter hexadecimal values for FC-MAP preceded by the hexadecimal indicator “0x”—for example, 0x0EFC00.) If you change the FC-MAP value of an FCF, change the FC-MAP value for the FCoE VLAN it belongs to on the switch and on the servers you want to communicate with the FCF. An FCoE VLAN can have one and only one FC-MAP value.

To enable VN2VF_Port FIP snooping:

- To enable VN2VF_Port FIP snooping on a single VLAN and specify the optional FC-MAP value:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan vlan-name examine-fip fc-map fc-map-value
```

For example, to enable VN2VF_Port FIP snooping on a VLAN named **san1_vlan** and change the FC-MAP value to **0x0EFC03**:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan san1_vlan examine-fip fc-map 0x0EFC03
```



NOTE: Changing the FC-MAP value causes all logins to drop and forces ENodes to log in again.

- To enable VN2VF_Port FIP snooping on all VLANs and use the default FC-MAP value:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all examine-fip
```

- To configure an interface as a trusted interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface interface-name fcoe-trusted
```

For example, to configure interface **xe-0/0/30** as a trusted interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface xe-0/0/30 fcoe-trusted
```

Related Documentation

- Example: Configuring an FCoE Transit Switch*
- Understanding FIP Snooping*

- [Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)

Disabling VN2VF_Port FIP Snooping on an FCoE-FC Gateway Switch Interface

When the switch acts as an FCoE-FC gateway, the FCoE-network-facing Ethernet interfaces in the FCoE VLAN are automatically enabled for VN_Port to VF_Port (VN2VF_Port) FIP snooping. You can disable VN2VF_Port FIP snooping on an individual Ethernet interface or you can disable VN2VF_Port FIP snooping globally for all Ethernet interfaces in a gateway Fibre Channel (FC) fabric.

Disable VN2VF_Port FIP snooping on an Ethernet interface by configuring it as an FCoE trusted interface. Disable VN2VF_Port FIP snooping on all Ethernet interfaces in an FC fabric by configuring the FC fabric as FCoE trusted.

Do not disable VN2VF_Port FIP snooping on an interface unless you are certain that the interface is connected to a trusted device. Do not disable VN2VF_Port FIP snooping on an FC fabric unless all of the FCoE-network-facing interfaces in the fabric are either connected to a transit switch that is performing VN2VF_Port FIP snooping on the FCoE devices as they log in to the FC network or all of the interfaces are connected to trusted devices.

VN2VF_Port FIP snooping installs firewall filters that block FIP and FCoE frames from sources that have not logged in to the switch and prevents unauthorized access to the network. Disabling VN2VF_Port FIP snooping disables these firewall filters and permits access to all FIP and FCoE frames transported on that interface.

- To disable VN2VF_Port FIP snooping on an FCoE-device-facing Ethernet interface in an FCoE VLAN, configure that interface as a trusted interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface interface-name fcoe-trusted
```

For example, to configure interface *xe-0/0/7* as a trusted FC interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface xe-0/0/7 fcoe-trusted
```

- To disable VN2VF_Port FIP snooping on all FCoE-device-facing interfaces in a gateway FC fabric, configure that fabric as a trusted fabric:

```
[edit]
user@switch# set fc-fabrics fabric-name protocols fip fcoe-trusted
```

For example, to configure an FC fabric named *santastic* as an FCoE trusted fabric:

```
[edit]
user@switch# set fc-fabrics santastic protocols fip fcoe-trusted
```

Related Documentation

- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)
- [Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch on page 5025](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)
- [Understanding an FCoE-FC Gateway on page 4975](#)

Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch

VN_Port to VN_Port (VN2VN_Port) FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN_Port FIP snooping is conceptually similar to VN2VF_Port FIP snooping between VN_Ports and VF_Ports, but VN2VN_Port FIP snooping does not require traffic between VN_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN_Port communicates transparently through the transit switch on a virtual link that emulates a direct connection to the VN_Port at the other end of the virtual link.

VN2VN_Port FIP snooping is disabled by default. You enable VN2VN_Port FIP snooping on a per-VLAN basis on VLANs that carry VN2VN_Port FCoE traffic. Ensure that the VLAN carries only FCoE traffic between VN_Ports, because enabling VN2VN_Port FIP snooping denies access for all other traffic, including VN2VF_Port FIP snooping traffic.

All ENodes that you want to communicate using VN2VN_Port FIP snooping must use an FCoE VLAN dedicated to VN2VN_Port traffic. You cannot mix VN2VN_Port FIP snooping traffic with VN2VF_Port FIP snooping traffic in the same FCoE VLAN.



NOTE: An FCoE VLAN can support either VN2VF_Port FIP snooping or VN2VN_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF_Port FIP snooping traffic and for VN2VN_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN_Port FIP snooping VLANs, VN2VF_Port traffic is dropped.

The *beacon period* is conceptually similar to the FIP keepalive period (timer) for VN2VF_Port FIP snooping virtual link maintenance. The beacon period performs virtual link maintenance for VN2VN_Port FIP snooping. It is the time interval between messages that verify the connection is still valid and the device at the other end of the virtual link is still reachable. You set the beacon period value for each FCoE VLAN that you configure to do VN2VN_Port FIP snooping.



NOTE: In addition to enabling VN2VN_Port FIP snooping and configuring the beacon period, you must also configure a dedicated FCoE VLAN for the VN2VN_Port traffic, and set the FCoE transit switch ports in the proper port mode and trusted or untrusted state (interfaces are untrusted by default). See the VN2VN_Port FIP snooping configuration example topics for complete configurations of several common network topologies.

To enable VN2VN_Port FIP snooping and set the beacon period on an FCoE VLAN that is dedicated to VN2VN_Port traffic:

- [edit ethernet-switching-options secure-access-port]
user@switch# **set vlan *vlan-name* examine-fip examine-vn2vn beacon-period *milliseconds***

For example, to enable VN2VN_Port FIP snooping on a VLAN named **vlan200** and set the beacon period to **90000** milliseconds:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan vlan200 examine-fip examine-vn2vn beacon-period 90000
```

Related Documentation

- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5155](#)
- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5159](#)
- [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5167](#)
- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)
- [Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch on page 5031](#)

Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. QFX Systems support three DCBX modes:

- Autonegotiation—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- IEEE DCBX—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- DCBX Version 1.01—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric Node devices come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



NOTE: QFX Systems do not support pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00. If a QFX Series interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

Related Documentation

- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Understanding DCBX on page 5051](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [show dcbx neighbors on page 5300](#)

Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



NOTE: LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default

forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



NOTE: If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE-FC gateway, it does not send or receive FCoE Initialization Protocol (FIP) packets.
- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]
user@switch# set enhanced-transmission-selection no-recommendation-tlv

Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)

- *Using DCBX Protocol to Lower Costs*

Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



NOTE: Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- `[edit protocols dcbx interface interface-name]`
`user@switch# set enhanced-transmission-selection no-recommendation-tlv`

Related Documentation

- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Understanding DCBX on page 5051](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



.....

NOTE: In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

.....

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [show dcbx neighbors on page 5300](#)

Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



NOTE: In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [ aliases ] [ bit-patterns ]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[ 001 101 ]
```

Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5300](#)

Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



NOTE: In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

```
[edit protocols]
```

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5300](#)

Configuration Statements

- [application \(Application Maps\) on page 5212](#)
- [application \(Applications\) on page 5213](#)
- [application-map on page 5214](#)
- [application-maps on page 5215](#)
- [applications \(Applications\) on page 5216](#)
- [applications \(DCBX\) on page 5217](#)

- [auto-load-rebalance](#) on page 5217
- [bb-sc-n](#) on page 5218
- [beacon-period](#) on page 5219
- [code-points \(Application Maps\)](#) on page 5220
- [dcbx](#) on page 5221
- [dcbx-version](#) on page 5222
- [description \(Fibre Channel Fabrics\)](#) on page 5223
- [destination-port \(Applications\)](#) on page 5224
- [disable \(DCBX\)](#) on page 5225
- [enhanced-transmission-selection](#) on page 5226
- [ether-type](#) on page 5227
- [examine-fip](#) on page 5228
- [examine-vn2vn](#) on page 5229
- [fabric-id](#) on page 5230
- [fabric-type](#) on page 5230
- [family fcoe](#) on page 5231
- [fc2](#) on page 5231
- [fc-fabrics](#) on page 5232
- [fc-map](#) on page 5234
- [fc-options](#) on page 5235
- [fcoe-trusted](#) on page 5236
- [fibre-channel \(Family Interfaces\)](#) on page 5237
- [fibre-channel \(Port\)](#) on page 5238
- [fibrechannel-options](#) on page 5238
- [fip](#) on page 5239
- [fka-adv-period](#) on page 5240
- [interface \(DCBX\)](#) on page 5241
- [interface \(Fibre Channel Fabric\)](#) on page 5242
- [interface \(FIP\)](#) on page 5243
- [load-balance-algorithm](#) on page 5244
- [loopback \(Fibre Channel Interface\)](#) on page 5245
- [max-login-sessions](#) on page 5246
- [max-login-sessions-per-node](#) on page 5247
- [max-sessions-per-enode](#) on page 5248
- [no-fabric-wwn-verify](#) on page 5249
- [oxid](#) on page 5250
- [policy-options](#) on page 5251

- [port-mode \(Fibre Channel Interfaces\) on page 5252](#)
- [port-range on page 5253](#)
- [priority \(FIP\) on page 5254](#)
- [priority-flow-control on page 5255](#)
- [protocol \(Applications\) on page 5256](#)
- [protocols \(FIP\) on page 5257](#)
- [proxy \(Fibre Channel\) on page 5258](#)
- [recommendation-tlv on page 5258](#)
- [speed \(Fibre Channel Interfaces\) on page 5259](#)
- [traceoptions \(FC-2 Fibre Channel\) on page 5260](#)
- [traceoptions \(Fibre Channel\) on page 5262](#)
- [traceoptions \(FIP Protocol Fibre Channel\) on page 5265](#)
- [traceoptions \(Proxy Fibre Channel\) on page 5267](#)

application (Application Maps)

Syntax	<pre>application <i>application-name</i> { <i>code-points</i> [<i>aliases</i>] [<i>bit-patterns</i>]; }</pre>
Hierarchy Level	[edit policy-options application-maps <i>application-map-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Add an application to an application map and define the application's code points.
Options	<i>application-name</i> —Name of the application. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application (Applications)

Syntax	<pre> application <i>application-name</i> { <i>destination-port</i> <i>port-value</i>; <i>protocol</i> (tcp udp); <i>ether-type</i> <i>type</i>; } </pre>
Hierarchy Level	[edit applications]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure properties to define an application.
Options	<p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Defining an Application for DCBX Application Protocol TLV Exchange on page 5207 • Example: Configuring DCBX Application Protocol TLV Exchange on page 5144 • Example: Configuring DCBX to Support an iSCSI Application • Understanding DCBX Application Protocol TLV Exchange on page 5060 • Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application-map

Syntax	<code>application-map <i>application-map-name</i>;</code>
Hierarchy Level	[edit protocols dcbx interface interface-name]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify an application map to apply to an interface.
Options	<i>application-map-name</i> —Name of the application map.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show dcbx neighbors on page 5300• Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application-maps

Syntax	<pre> application-maps <i>application-map-name</i> { application <i>application-name</i> { code-points [<i>aliases</i>] [<i>bit-patterns</i>]; } } </pre>
Hierarchy Level	[edit policy-options]
Release Information	<p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Define an application map by specifying the applications that belong to the application map.
Options	<p><i>application-map-name</i>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209 • Example: Configuring DCBX Application Protocol TLV Exchange on page 5144 • Example: Configuring DCBX to Support an iSCSI Application • Understanding DCBX Application Protocol TLV Exchange on page 5060 • Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

applications (Applications)

Syntax	<pre>applications { application application-name { destination-port port-value; protocol (tcp udp); ether-type type; } }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Define applications that DCBX advertises.
Options	The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Defining an Application for DCBX Application Protocol TLV Exchange on page 5207• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

applications (DCBX)

Syntax	<code>applications { no-auto-negotiation; }</code>
Hierarchy Level	[edit <code>protocols dcbx interface interface-name</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.1 for the EX Series
Description	Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 5300 • Understanding DCB Features and Requirements on page 4965

auto-load-rebalance

Syntax	<code>auto-load-rebalance;</code>
Hierarchy Level	[edit <code>fc-fabrics fabric-name proxy</code>]
Release Information	Command introduced in Junos OS Release 12.3 for the QFX Series.
Description	Configure the system to rebalance NP_Port link loads automatically on an FCoE-FC gateway proxy fabric if the link loads become unbalanced. Load rebalancing is a disruptive action that forces some or all sessions (depending on the configured load-balancing algorithm) to log out and then log in again. When sessions log in again, they are placed on NP_Port interfaces so that the link loads are balanced.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Defining the Proxy Load-Balancing Algorithm on page 5190 • Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175 • Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test) on page 5191 • Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008 • Monitoring Fibre Channel Interface Load Balancing on page 5269

bb-sc-n

Syntax	<code>bb-sc-n <i>bb-sc-n</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> fibrechannel-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the buffer-to-buffer credit state change number to prevent the permanent loss of Fibre Channel credits over time (buffer-to-buffer credit recovery).
Options	<p><i>bb-sc-n</i>—Number of buffer-to-buffer state change credits.</p> <p>Range: 0 through 15</p> <p>Default: 0 (disabled)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show fibre-channel interfaces on page 5357• Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099• Configuring a Fibre Channel Interface on page 5182• Configuring a Physical Fibre Channel Interface on page 5181• Understanding Interfaces on an FCoE-FC Gateway on page 4996

beacon-period

Syntax	<code>beacon-period <i>milliseconds</i></code>
Hierarchy Level	[edit ethernet-switching options secure-access-port vlan (all <i>vlan-name</i>) examine-fip examine-vn2vn]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	<p>Set the interval between periodic beacons. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements.</p> <p>The ENode sends periodic beacons every 90 seconds on behalf of the VN_Port. Each received beacon resets the session timer for the virtual link connection to the other VN_Port. If the FCF does not receive a beacon before the beacon timer expires, the VN_Port is considered as “down” and the virtual link is terminated. The beacon timer expires in 2.5 times the configured beacon timer value.</p>
Options	<p><i>milliseconds</i>—Time in milliseconds between beacons.</p> <p>Range: 250 through 90000 milliseconds</p> <p>Default: 8000 milliseconds</p>
Required Privilege Level	<p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch) on page 5155 • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches) on page 5159 • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch) on page 5167

code-points (Application Maps)

Syntax	<code>code-points [<i>aliases</i>] [<i>bit-patterns</i>];</code>
Hierarchy Level	[edit policy-options application-maps <i>application-map-name</i> application <i>application-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Define one or more code-point aliases or bit sets for an application.
Options	<i>aliases</i> —Name of the alias or aliases. <i>bit-patterns</i> —Value of the code-point bits, in decimal form.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

dcbx

Syntax	<pre> dcbx { disable; interface (interface-name all) { disable; application-map application-map-name; applications { no-auto-negotiation; } enhanced-transmission-selection { no-auto-negotiation; no-recommendation-tlv; recommendation-tlv { no-auto-negotiation; } } dcbx-version (auto-negotiate ieee-dcbx dcbx-version-1.01); priority-flow-control { no-auto-negotiation; } } } </pre>
Hierarchy Level	[edit protocols]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p>mode and recommendation-tlv statements introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Configure DCBX properties.
Options	The statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 5300 • Understanding DCB Features and Requirements on page 4965 • Configuring DCBX Autonegotiation on page 5204 • <i>Understanding DCB Features and Requirements on EX Series Switches</i> • <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i>


dcbx-version

Syntax	<code>dcbx-version (auto-negotiate ieee-dcbx dcbx-version-1.01);</code>
Hierarchy Level	[edit protocols dcbx interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	<p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p>
Default	The default DCBX mode is autonegotiation.
Options	<p>auto-negotiate—Automatically negotiate the DCBX version with the connected peer.</p> <p>ieee-dcbx—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p>dcbx-version-1.01—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show dcbx neighbors on page 5300• Configuring DCBX Autonegotiation on page 5204• Understanding DCBX on page 5051

description (Fibre Channel Fabrics)

Syntax	<code>description <i>description</i></code>
Hierarchy Level	[edit fc-fabrics <i>fabric-name</i>]
Description	Text string that describes the Fibre Channel fabric. The text string has no effect on the operation of the fabric.
Options	<i>description</i> —Text that describes the fabric. Text can include letters, numbers, and hyphens (-) and can be up to 255 characters in length. If the text includes spaces, enclose the entire text string in quotation marks.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fabric on page 5322

destination-port (Applications)

Syntax	<code>destination-port <i>port-value</i>;</code>
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with protocol to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml for a list of assigned port numbers.</p>
<hr/>	
<div> NOTE: To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>.</div> <hr/>	
Options	<i>port-value</i> —Identifier for the port.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Defining an Application for DCBX Application Protocol TLV Exchange on page 5207• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches


disable (DCBX)

Syntax	disable
Hierarchy Level	[edit protocols dcbx] [edit protocols dcbx interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.
Default	DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces. DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring DCBX Autonegotiation on page 5204 • <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i> • Understanding DCB Features and Requirements on page 4965 • <i>Understanding DCB Features and Requirements on EX Series Switches</i>

enhanced-transmission-selection

Syntax	<pre>enhanced-transmission-selection { no-auto-negotiation; no-recommendation-tlv; recommendation-tlv { no-auto-negotiation; } }</pre>
Hierarchy Level	[edit protocols dcbx interface interface-name]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p>
Options	<p>no-auto-negotiation—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p>no-recommendation-tlv—Disable automatic negotiation of the ETS Recommendation TLV</p> <p>recommendation-tlv—Enable automatic negotiation of ETS Recommendation TLV</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show dcbx neighbors on page 5300• Configuring DCBX Autonegotiation on page 5204• Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606• Understanding DCB Features and Requirements on page 4965

ether-type

Syntax	<code>ether-type <i>ether-type</i>;</code>
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See http://standards.ieee.org/develop/regauth/ethertype/eth.txt for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.
<div>  NOTE: To create a FIP application, use the EtherType 0x8914. </div>	
Options	<i>type</i> —Identifier for the EtherType.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Defining an Application for DCBX Application Protocol TLV Exchange on page 5207 • Example: Configuring DCBX Application Protocol TLV Exchange on page 5144 • Understanding DCBX Application Protocol TLV Exchange on page 5060

examine-fip

Syntax	<pre>examine-fip { examine-vn2vn { beacon-period milliseconds; } fc-map fc-map-value; }</pre>
Hierarchy Level	[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)]
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement examine-vn2vn introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	<p>Enable FIP snooping on a specified VLAN. Ensure that the VLAN is a dedicated FCoE VLAN that transports only FCoE traffic.</p> <p>(QFX Series only) Enable VN_Port to VN_Port (VN2VN_Port) FIP snooping on a specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only VN2VN_Port traffic. One FCoE VLAN cannot support both VN_Port to VF_Port (VN2VF_Port) FIP snooping and VN2VN_Port FIP snooping. Configure separate, dedicated FCoE VLANs for VN2VN_Port FIP snooping and VN2VN_Port FIP snooping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>vlan</i>• <i>Example: Configuring an FCoE Transit Switch</i>• Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199

examine-vn2vn

Syntax	<pre>examine-vn2vn { beacon-period milliseconds; }</pre>
Hierarchy Level	[edit ethernet-switching options secure-access-port vlan (all <i>vlan-name</i>) examine-fip]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	<p>Enable VN_Port to VN_Port (VN2VN) FIP snooping on a specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only FCoE traffic. A VLAN cannot support VN2VN FIP snooping and VN_Port to VF_Port FIP snooping (VN2VF) simultaneously. Configure separate VLANs for VN2VN FIP snooping and VN2VF FIP snooping.</p> <p>When you enable VN2VN FIP snooping on a VLAN, the VN2VF session filters are removed and the all existing VN2VF sessions are terminated.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch) on page 5155 • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches) on page 5159 • Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch) on page 5167

fabric-id

Syntax	<code>fabric-id <i>fc-fabric-id</i>;</code>
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a unique identifier for the FC fabric.



NOTE: Changing the ID of an FC fabric causes all logins to drop and forces the ENodes to log in again.

Options	<code><i>fc-fabric-id</i></code> —Unique identifier of the FC fabric.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fabric on page 5322• Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179• Understanding an FCoE-FC Gateway on page 4975

fabric-type

Syntax	<code>fabric-type proxy;</code>
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify that the FC fabric be an FCoE-FC gateway fabric.
Options	<code>proxy</code> —Specify that the switch be an FCoE-FC gateway fabric.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fabric on page 5322• Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179• Understanding an FCoE-FC Gateway on page 4975

family fcoe

Syntax	family fcoe { oxid (enable disable); }
Hierarchy Level	[edit forwarding-options hash-key]
Release Information	Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Configure whether or not to use the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.
Options	The statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling CoS OxID Hash Control on page 5192 • Understanding OxID Hash Control for FCoE Traffic Load Balancing on page 5024

fc2

Syntax	fc2 { traceoptions { file <i>filename</i> <replace> <size <i>size</i> > <files <i>number</i> > <no-stamp>; <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } }
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i>]
Description	Fibre Channel network layer (FC2) configuration.
Options	The statements are explained separately.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.

fc-fabrics

```
Syntax  fc-fabrics {
        fc-fabric-name {
            description
            fabric-id fc-fabric-id;
            fabric-type proxy;
            interface {
                interface-name {
                    max-login-sessions max-login-sessions;
                }
                interface-name {
                    max-login-sessions max-login-sessions;
                }
                <...>;
                max-login-sessions max-login-sessions;
            }
            vlan.interface-name;
        }
        fc2 {
            traceoptions {
                file filename <replace> <size size> <files number> <no-stamp>;
                <world-readable | no-world-readable>;
                flag flag <flag-modifier>;
            }
        }
        max-login-sessions max-login-sessions;
        protocols {
            fip {
                fcoe-trusted;
                fc-map fc-map-value;
                fka-adv-period milliseconds;
                interface {
                    interface-name {
                        fka-adv-period milliseconds;
                        priority priority;
                    }
                }
                max-sessions-per-enode max-sessions-per-enode;
                priority priority;
                traceoptions {
                    file filename <replace> <size size> <files number> <no-stamp>;
                    <world-readable | no-world-readable>;
                    flag flag <flag-modifier> <disable>;
                }
            }
        }
        proxy {
            auto-load-rebalance
            load-balance-algorithm (simple | enode-based | flogi-based);
            no-fabric-wwn-verify;
            traceoptions {
                file filename <replace> <size size> <files number> <no-stamp>;
                <world-readable | no-world-readable>;
            }
        }
    }
```

```

        flag flag <flag-modifier> <disable>;
    }
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure an FC fabric. You can configure a maximum of 12 FC fabrics, one per native FC port.



NOTE: Changing the name of an FC fabric causes all logins to drop and forces the ENodes to log in again.

Options *fc-fabric-name* —Unique name of the FC fabric.


The other statements are explained separately.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [show fibre-channel fabric on page 5322](#)
- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Assigning Interfaces to a Fibre Channel Fabric on page 5187](#)
- [Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)

fc-map

Syntax	<code>fc-map <i>fc-map-value</i>;</code>
Hierarchy Level	<p>[edit ethernet-switching options secure-access-port vlan (all <i>vlan-name</i>) examine-fip]</p> <p>For the QFX Series only:</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).</p> <p>You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN_Port to VF_Port (VN2VF_Port) FIP snooping (0x0EFC00) than for VN_Port to VN_Port (VN2VN_Port) FIP snooping.</p> <p>The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN_Port MAC address.</p> <p>When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.</p>
	<p> NOTE: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.</p>
Options	<p><i>fc-map-value</i>—FC-MAP value, hexadecimal value preceded by “0x”.</p> <p>Range: 0x0EFC00 through 0x0EFCFF</p> <p>Default: 0x0EFC00 for VN2VF_Port FIP snooping 0x0EFD00 for VN2VN_Port FIP snooping</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • examine-fip on page 4813 • show fip snooping on page 5377 • <i>Example: Configuring an FCoE Transit Switch</i>

- [Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199](#)

fc-options

Syntax `fc-options`
 `max-login-sessions-per-node` *max-login-sessions-per-node*;
 `traceoptions` {
 `file` *filename* <replace> <size *size*> <files *number*> <no-stamp>;
 <world-readable | no-world-readable>;
 `flag` *flag* <*flag-modifier*> <disable>;
 }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set Fibre Channel options.

Required Privilege storage—To view this statement in the configuration.

Level storage-control—To add this statement to the configuration.

fcoe-trusted

Syntax	fcoe-trusted;
Hierarchy Level	[edit ethernet-switching-options secure-access-port interface <i>interface-name</i>] For the QFX Series only: [edit fc-fabrics <i>fc-fabric-name</i> protocols fip]
Release Information	Statement introduced in Junos OS Release 10.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.
Description	<p>Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>(QFX Series only) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the fcoe-trusted configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fip snooping on page 5377• <i>Example: Configuring an FCoE Transit Switch</i>• Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199

fibre-channel (Family Interfaces)

Syntax	fibre-channel { <code>port-mode</code> (f-port np-port); }
Hierarchy Level	[edit <code>interfaces</code> vlan <code>unit</code> <i>logical-unit-number</i> <code>family</code>], [edit <code>interfaces</code> <i>interface-name</i> <code>unit</code> <i>logical-unit-number</i> <code>family</code>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the port mode for FCoE VLAN interfaces and native FC interfaces.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099 • Configuring an FCoE VLAN Interface on page 5185 • Configuring a Fibre Channel Interface on page 5182 • show fibre-channel interfaces on page 5357 • show vlans on page 2251 • Understanding Interfaces on an FCoE-FC Gateway on page 4996

fibre-channel (Port)

Syntax	<pre>fibre-channel { port-range { port-range-low port-range-high; } }</pre>
Hierarchy Level	[edit chassis fpc fpc-id pic pic-id]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify a range of ports to carry FC traffic when the switch is configured as an FCoE-FC gateway.
Options	The statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099• Configuring a Physical Fibre Channel Interface on page 5181• show fibre-channel interfaces on page 5357• Understanding Interfaces on an FCoE-FC Gateway on page 4996

fibrechannel-options

Syntax	<pre>fibrechannel-options { bb-sc-n (loopback no-loopback); speed (auto-negotiation 2g 4g 8g); }</pre>
Hierarchy Level	[edit interfaces interface-name]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure FC interface properties such as speed and loopback mode.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel interfaces on page 5357• Configuring a Fibre Channel Interface on page 5182

fip

```
Syntax  fip {
        fcoe-trusted;
        fc-map fc-map-value;
        fka-adv-period milliseconds;
        interface {
            interface-name {
                fka-adv-period milliseconds;
                priority priority;
            }
        }
        max-sessions-per-enode max-sessions-per-enode;
        priority priority;
        traceoptions {
            file filename <replace> <size size> <files number> <no-stamp>;
            <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
```

Hierarchy Level [edit [fc-fabrics](#) *fc-fabric-name* [protocols](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure global or interface-specific FIP options. Individual interface settings override global settings.

Options The statements are explained separately.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [show fibre-channel fip on page 5328](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Overview of FIP on page 4960](#)

fka-adv-period

Syntax	fka-adv-period <i>milliseconds</i> ;
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i> protocols fip], [edit fc-fabrics <i>fc-fabric-name</i> protocols fip interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the global or interface-specific interval between periodic FIP keepalive advertisements. An interval set at the interface level overrides the global setting.
Options	<i>milliseconds</i> —Time in milliseconds between FIP keepalive advertisements. Range: 250 through 90000 milliseconds Default: 8000 milliseconds
Required Privilege Level	storage —To view this statement in the configuration. storage-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fip on page 5328• show fibre-channel fip interface on page 5343• Overview of FIP on page 4960

interface (DCBX)

Syntax	<pre> interface (<i>interface-name</i> all) { disable; application-map <i>application-map-name</i>; applications { no-auto-negotiation; } enhanced-transmission-selection { no-auto-negotiation; no-recommendation-tlv; recommendation-tlv { no-auto-negotiation; } } dcbx-version (auto-negotiate ieee-dcbx dcbx-version-1.01); priority-flow-control { no-auto-negotiation; } } </pre>
Hierarchy Level	[edit protocols dcbx]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p>Mode and recommendation-tlv statements introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Configure DCBX properties on an interface.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 5300 • Configuring DCBX Autonegotiation on page 5204 • Example: Configuring DCBX to Support an iSCSI Application • Understanding DCB Features and Requirements on page 4965 • Understanding DCB Features and Requirements on EX Series Switches • Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

interface (Fibre Channel Fabric)

Syntax

```
interface {  
    interface-name {  
        max-login-sessions max-login-sessions;  
    }  
    interface-name {  
        max-login-sessions max-login-sessions;  
    }  
    <...> {  
        max-login-sessions max-login-sessions;  
    }  
    vlan.interface-name;  
}
```

Hierarchy Level [edit [fc-fabrics](#) *fc-fabric-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Associate one or more native Fibre Channel (FC) interfaces with an FC fabric and one VLAN interface for FCoE traffic. An FC interface can be associated with only one FC fabric.

Options *interface-name*—Name of the native FC interface. You can assign one or more FC interfaces to an FC fabric.

vlan.vlan-interface-name—Name of the VLAN interface for FCoE traffic. You can assign one VLAN interface to an FC fabric.

The remaining statement is explained separately.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099](#)
- [Configuring a Fibre Channel Interface on page 5182](#)
- [Configuring a Physical Fibre Channel Interface on page 5181](#)
- [Configuring an FCoE VLAN Interface on page 5185](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

interface (FIP)

Syntax	<pre> interface { interface-name { fka-adv-period milliseconds; priority priority; } } </pre>
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure FIP options on a per-interface basis. (Override global FIP configuration for a specified interface.)
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel fip on page 5328 • Configuring FIP on an FCoE-FC Gateway on page 5192 • Overview of FIP on page 4960

load-balance-algorithm

Syntax	load-balance-algorithm (simple enode-based flogi-based);
Hierarchy Level	[edit fc-fabrics <i>fabric-name</i> proxy]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set the load-balancing algorithm that the QFX Series uses to distribute FCoE sessions (FLOGI and FDISC sessions from the FCoE devices in the Ethernet network) among the NP_Port links to the FC switch.



NOTE: Changing the load-balancing algorithm when FCoE sessions are running forces the FCoE sessions to log out, then log in again.

Options	<p>simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. To further minimize disruption, the algorithm logs out the sessions with the fewest dependencies (for example, FDISC sessions are logged out before FLOGI sessions). When the sessions log in again, they are placed on NP_Port interfaces in a manner that balances the link loads. This is the default load-balancing algorithm.</p> <p>enode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions (VN_Port sessions) associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system logs off all sessions. The sessions log in again and are placed on NP_Port interfaces in a balanced manner.</p> <p>flogi-based—FLOGI-based load balancing is similar to ENode-based load balancing, but the behavior when the loads are rebalanced is different. Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. When the logged out sessions log back in, they are placed on NP_Port interfaces in a manner that balances the link loads.</p>
Required Privilege Level	<p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Defining the Proxy Load-Balancing Algorithm on page 5190• Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175

- [Simulating On-Demand Fibre Channel Link Load Rebalancing \(Dry Run Test\) on page 5191](#)
- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)


loopback (Fibre Channel Interface)

Syntax	(loopback no-loopback);
Hierarchy Level	[edit interfaces <i>interface-name</i> fibrechannel-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Enable or disable loopback mode for FC interfaces.
Default	By default, loopback mode is disabled on FC interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel interfaces on page 5357• Configuring a Fibre Channel Interface on page 5182


max-login-sessions

Syntax	<code>max-login-sessions <i>max-login-sessions</i>;</code>
Hierarchy Level	<code>[edit fc-fabrics <i>fc-fabric-name</i>];</code> <code>[edit fc-fabrics <i>fc-fabric-name</i> interface <i>interface-name</i>];</code>
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	<p>Set the maximum number of FCoE initialization protocol (FIP) session logins permitted for an individual NP_Port interface in an FCoE-FC gateway fabric (FC fabric) or for the entire FCoE-FC gateway fabric. You can set a maximum FIP session limit for each NP_Port interface connected to an FC switch. You can also set a maximum FIP session limit for the entire FC fabric. The sum of the maximum login sessions permitted on the NP_Port interfaces in an FC fabric should not exceed the maximum login sessions configured for that FC fabric.</p> <p>The maximum number of FIP sessions (the combined total of all VN2VF_Port and VN2VN_Port sessions on the system) is 2500 sessions.</p>
Options	<p><i>max-login-sessions</i>—Maximum number of FIP login sessions.</p> <p>Range: 128 through 2500</p> <p>Default: 2500</p>
Required Privilege Level	<p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• max-login-sessions-per-node on page 5247• Setting the Maximum Number of FIP Login Sessions per FC Interface on page 5196• Setting the Maximum Number of FIP Login Sessions per FC Fabric on page 5197• Understanding Interfaces on an FCoE-FC Gateway on page 4996


max-login-sessions-per-node

Syntax	<code>max-login-sessions-per-node</code> <i>max-login-sessions-per-node</i> ;
Hierarchy Level	[edit fc-options]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	<p>Set the maximum number of FCoE initialization protocol (FIP) session logins permitted on a Node device. (This is the combined total of all VN2VF_Port and VN2VN_Port sessions on the Node device.)</p> <p>On a QFX3500 switch, the max-login-sessions-per-node command sets the maximum FIP session login limit for all of the FC fabrics configured on the device. The combined number of FIP sessions on all FC fabrics on the device should not exceed this limit.</p> <p>On a QFabric system, the max-login-sessions-per-node command globally sets the maximum FIP session login limit for each QFabric system Node device in the QFabric system. For example, if you set the Node limit to 2000 login sessions, then each QFabric Node device supports up to 2000 FIP login sessions. The total configured maximum number of login sessions of all of the FC fabrics on a Node device should not exceed the Node session limit.</p>
	<div>  <p>NOTE: FIP login session limits configured at the FC fabric level or at the FC fabric interface level might limit a Node device to fewer total sessions than the configured Node limit.</p> </div>
Options	<p>max-login-sessions-per-node—Maximum number of FIP login sessions.</p> <p>Range: 128 through 2500</p> <p>Default: 2500</p>
Required Privilege Level	<p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • max-login-sessions on page 5246 • Setting the Maximum Number of FIP Login Sessions per Node Device on page 5198 • Understanding Interfaces on an FCoE-FC Gateway on page 4996

max-sessions-per-enode

Syntax	<code>max-sessions-per-enode max-sessions-per-enode;</code>
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set the maximum number of FCoE login sessions (FLOGI plus FDISC) from a single ENode allowed on the gateway FC fabric (the fabric configured on the QFabric system). The maximum number of logins per ENode is 2000 sessions.
<div> NOTE: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions. There is no limit to the number of end-to-end storage sessions.</div>	
Options	<i>max-sessions-per-enode</i> —Maximum number of FCoE sessions a single ENode can establish on the switch. Range: 32 through 2000 Default: 32
Required Privilege Level	storage —To view this statement in the configuration. storage-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• fcoe-trusted on page 4815• show fibre-channel fip on page 5328• Configuring FIP on an FCoE-FC Gateway on page 5192• Understanding FIP Parameters on an FCoE-FC Gateway on page 4992

no-fabric-wwn-verify

Syntax	no-fabric-wwn-verify;
Hierarchy Level	[edit fc-fabrics <i>fabric-name</i> proxy]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable the fabric worldwide name (WWN) verification check in the fabric login accept message (FLOGI-ACC) for implicit FLOGIs. If you enable this option, when a QFX Series NP_Port performs a FLOGI to the FC fabric, the QFX Series does not verify the fabric WWN in the FLOGI-ACC against the current fabric WWN.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: Disabling or enabling the fabric WWN verification check logs out all FCoE sessions.</p> </div> </div>	
Default	Disabled. By default, all implicit FLOGIs from the QFX Series NP_Ports to the FC fabric are verified against the current fabric WWN.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel proxy fabric-state on page 5364 • Understanding FCoE-FC Gateway Functions on page 4979

oxid

Syntax	oxid (enable disable)
Hierarchy Level	[edit forwarding-options hash-key family fcoe]
Release Information	Statement introduced in Junos OS Release 12.3 for the QFX Series.
Description	Enable or disable whether the switch uses the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.
Default	OxID hash control is enabled by default.
Options	oxid (enable disable) —Enable or disable whether the switch uses the OxID hash control field for FCoE traffic load balancing.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling and Disabling CoS OxID Hash Control on page 5192• Understanding OxID Hash Control for FCoE Traffic Load Balancing on page 5024

policy-options

```
Syntax  policy-options
        application-maps application-map-name {
            application application-name {
                code-points [ aliases ] [ bit-patterns ];
            }
        }
        policy-statement policy-name {
            term term-name {
                from {
                    family family-name;
                    match-conditions;
                    policy subroutine-policy-name;
                    prefix-list prefix-list-name;
                    prefix-list-filter prefix-list-name match-type <actions>;
                    route-filter destination-prefix match-type <actions>;
                    source-address-filter source-prefix match-type <actions>;
                }
                to {
                    match-conditions;
                    policy subroutine-policy-name;
                }
                then actions;
            }
        }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for the EX Series.

Description Configure options such as application maps for DCBX application protocol exchange and policy statements.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.


Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

port-mode (Fibre Channel Interfaces)

Syntax	<code>port-mode (f-port np-port);</code>
Hierarchy Level	[edit interfaces <i>vlan</i> unit <i>unit</i> family <i>fibre-channel</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>fibre-channel</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the FCoE VLAN interface port mode to F_Port to connect the switch to FCoE initiators, or configure the native FC interface port mode to proxy N_Port (NP_Port) to connect the switch to an FC switch fabric port (F_Port).
Options	f-port —Configure an FCoE VLAN interface to connect to FCoE initiator Virtual N_Ports (VN_Ports). np-port —Configure a native FC port to connect to an FC switch F_Port.
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel interfaces on page 5357• Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099• Configuring an FCoE VLAN Interface on page 5185• Configuring a Fibre Channel Interface on page 5182• Understanding Interfaces on an FCoE-FC Gateway on page 4996

port-range

Syntax	<code>port-range <i>port-range-low</i> <i>port-range-high</i>;</code>
Hierarchy Level	[edit chassis fpc <i>fpc-id</i> pic <i>pic-id</i> fibre-channel]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Configure a contiguous block of ports as FC ports. You can configure the FC-capable ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to create blocks of native FC interfaces. You cannot individually configure a single port as a native FC interface. Within these port blocks, you cannot mix FC interfaces with Ethernet interfaces. All of the ports in a block must be either native FC interfaces or Ethernet interfaces.</p> <p>You can configure:</p> <ul style="list-style-type: none"> • Six native FC interfaces by configuring either ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, or ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47. • Twelve native FC interfaces by configuring ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47. • No native FC interfaces by leaving ports xe-0/0/0 through xe-0/0/5 and ports xe-0/0/42 through xe-0/0/47 in their default state as Ethernet interfaces.
Options	<p><i>port-range-low</i>—Lowest-numbered port in the block of native FC interfaces, either 0 or 42.</p> <p><i>port-range-high</i>—Highest-numbered port in the block of native FC interfaces. The value is 5 if the <i>port-range-low</i> value is 0. The value is 47 if the <i>port-range-low</i> value is 42.</p>
<div>  <p>NOTE: Only a complete block of ports, xe-0/0/0 through xe-0/0/5, xe-0/0/42 through xe0/0/47, or both, can be configured as FC ports.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel interfaces on page 5357 • Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099 • Configuring a Physical Fibre Channel Interface on page 5181 • Understanding Interfaces on an FCoE-FC Gateway on page 4996


priority (FIP)

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	[edit fc-fabrics <i>fc-fabric-name</i> protocols fip], [edit fc-fabrics <i>fc-fabric-name</i> protocols fip interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Sets the global or interface-specific priority value associated with the switch FCF-MAC. CNAs use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. The switch advertises this value to the server ENodes on the FCoE network. A priority value set at the interface level overrides the global setting.
Options	<i>priority</i> —Value that determines the FCF an ENode selects to perform FIP FLOGI. The lower the priority number, the higher the priority of the FCF. Range: 0 through 255 Default: 128
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fip on page 5328• show fibre-channel fip interface on page 5343• Overview of FIP on page 4960• Configuring FIP on an FCoE-FC Gateway on page 5192

priority-flow-control

Syntax	<code>priority-flow-control { no-auto-negotiation; }</code>
Hierarchy Level	[edit protocols dcbx interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.
Options	no-auto-negotiation —Disable automatic negotiation of PFC.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 5300 • Configuring CoS PFC (Congestion Notification Profiles) on page 5795 • Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) • Configuring DCBX Autonegotiation on page 5204 • Example: Configuring CoS PFC for FCoE Traffic on page 5113 • Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches • Understanding Priority-Based Flow Control • Understanding DCB Features and Requirements on page 4965

protocol (Applications)

Syntax	<code>protocol (tcp udp);</code>
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Networking protocol type, which combines with destination-port to identify an application type.
<div> NOTE: To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number 3260.</div>	
Options	<code>tcp</code> —Transmission Control Protocol <code>udp</code> —User Datagram Protocol
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Defining an Application for DCBX Application Protocol TLV Exchange on page 5207• Example: Configuring DCBX Application Protocol TLV Exchange on page 5144• Example: Configuring DCBX to Support an iSCSI Application• Understanding DCBX Application Protocol TLV Exchange on page 5060• Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

protocols (FIP)

```
Syntax protocols {
    fip {
        fcoe-trusted;
        fc-map fc-map-value;
        fka-adv-period milliseconds;
        interface {
            interface-name {
                fka-adv-period milliseconds;
                priority priority;
            }
        }
        max-sessions-per-enode max-sessions-per-enode;
        priority priority;
        traceoptions {
            file filename <replace> <size size> <files number> <no-stamp>;
            <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
}
```

Hierarchy Level [edit [fc-fabrics](#) *fc-fabric-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure global or interface-specific FC protocol options. Individual interface settings override global settings.

Options The statements are explained separately.

Required Privilege Level storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- [show fibre-channel fip on page 5328](#)
- [Configuring FIP on an FCoE-FC Gateway on page 5192](#)
- [Overview of FIP on page 4960](#)

proxy (Fibre Channel)

Syntax	<pre>proxy { auto-load-rebalance load-balance-algorithm (simple enode-based flog-based); no-fabric-wwn-verify; traceoptions { file <i>filename</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } }</pre>
Hierarchy Level	[edit fc-fabrics <i>fabric-name</i>]
Description	Configure proxy fabric operations.
Options	The statement is explained separately.
Required Privilege Level	storage—To view this statement in the configuration. storage-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding FCoE-FC Gateway Functions on page 4979

recommendation-tlv

Syntax	<pre>recommendation-tlv { no-auto-negotiation; }</pre>
Hierarchy Level	[edit protocols dcbx interface <i>interface-name</i> enhanced-transmission-selection]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	Disable or enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.)
Default	DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.
Options	no-auto-negotiation —Disable sending of the ETS recommendation TLV.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show dcbx neighbors on page 5300• Configuring DCBX Autonegotiation on page 5204

speed (Fibre Channel Interfaces)

Syntax	<code>speed (auto-negotiation 2g 4g 8g);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> fibrechannel-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure FC interface speed.
Options	<p>auto-negotiation—Automatically negotiate interface speed to match the speed of the attached link (2 Gbps, 4 Gbps, 8 Gbps).</p> <p>2g—2 Gbps link speed</p> <p>4g—4 Gbps link speed</p> <p>8g—8 Gbps link speed</p> <p>Default: <code>auto-negotiation</code></p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel interfaces on page 5357 • Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099 • Configuring a Fibre Channel Interface on page 5182 • Configuring a Physical Fibre Channel Interface on page 5181 • Understanding Interfaces on an FCoE-FC Gateway on page 4996

traceoptions (FC-2 Fibre Channel)

Syntax `traceoptions {
 file filename <replace> <size size> <files number> <no-stamp>;
 <world-readable | no-world-readable>;
 flag flag <flag-modifier>;
 }`

Hierarchy Level [edit **fc-fabrics** *fabric-name* **fc2**]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set FC-2 protocol tracing options.



NOTE: The **traceoptions** statement is not supported on the QFabric system.

Default Traceoptions is disabled.

Options **file *name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. The traceoption output continues in a second trace file named **trace-file.1**. When **trace-file.1** reaches its maximum size, output continues in a third file named **trace-file.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all operations.
- **error**—Trace all error events
- **normal**—Trace all normal events.

Default: If you do not specify the **normal** option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **rx-frame**—(Optional) Trace received frames.
- **rx-frame-header**—(Optional) Trace received frame headers.

- **tx-frame**—(Optional) Trace transmitted frames.
- **tx-frame-header**—(Optional) Trace transmitted frame headers.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***. Incoming tracefile data is logged in the now empty ***trace-file***. When ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***xk*** to specify KB, ***xm*** to specify MB, or ***xg*** to specify GB

Range: 10 KB through the maximum file size of 4 GB (the maximum is lower if 4 GB is not supported on your system)

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	storage—To view this statement in the configuration.
Level	storage-control—To add this statement to the configuration.

traceoptions (Fibre Channel)

Syntax	<pre>traceoptions { file <i>filename</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; }</pre>
Hierarchy Level	[edit fc-options]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set FC protocol tracing options.



NOTE: The **traceoptions** statement is not supported on the QFabric system.

Default Traceoptions is disabled.

Options **file *name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. The traceoption output continues in a second trace file named **trace-file.1**. When **trace-file.1** reaches its maximum size, output continues in a third file named **trace-file.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all operations.
- **fabric**—Trace virtual fabric events.
- **fc2**—Trace the FC2 (network layer protocols) events.
- **fip**—Trace the Fibre Channel over Ethernet (FCoE) Initialization Protocol events.
- **flogi**—Trace the fabric login server events.
- **forwarding-database**—Trace the forwarding database and next-hop events.
- **interface**—Trace the interface events.

- **krt**—Trace the communication over the routing socket.
- **lib**—Trace library calls.
- **lif**—Trace Fibre Channel logical interface (fc-lif) events.
- **vswitch**—Trace virtual switch events.

The following are the global tracing options:

- **all**—All trace operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.O**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.O** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.O**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	storage—To view this statement in the configuration.
Level	storage-control—To add this statement to the configuration.

traceoptions (FIP Protocol Fibre Channel)

Syntax `traceoptions {
 file filename <replace> <size size> <files number> <no-stamp>
 <world-readable | no-world-readable>;
 flag flag <flag-modifier>
}`

Hierarchy Level [edit [fc-fabrics](#) *fabric-name* [protocols](#) [fip](#)]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set proxy FC protocol tracing options.



NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default Traceoptions is disabled.

Options `file name`—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

`files number` —(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`. The traceoption output continues in a second trace file named `trace-file.1`. When `trace-file.1` reaches its maximum size, output continues in a third file named `trace-file.2`, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

Range: 2 through 1000 files

Default: 1 trace file

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Trace all operations.
- `error`—Trace all error events
- `normal`—Trace all normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- `packet`—Trace packet decoding operations
- `parse`—Trace configuration parsing.
- `state`—Trace state transitions.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	storage —To view this statement in the configuration.
	storage-control —To add this statement to the configuration.

traceoptions (Proxy Fibre Channel)

Syntax `traceoptions {
 file filename <replace> <size size> <files number> <no-stamp>
 <world-readable | no-world-readable>;
 flag flag <flag-modifier>
 }`

Hierarchy Level [edit `fc-fabrics fabric-name proxy`]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set proxy FC protocol tracing options.



NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default Traceoptions is disabled.

Options `file name`—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

`files number`—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`. The traceoption output continues in a second trace file named `trace-file.1`. When `trace-file.1` reaches its maximum size, output continues in a third file named `trace-file.2`, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

Range: 2 through 1000 files

Default: 1 trace file

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Trace all operations.
- `error`—Trace all error events.
- `interface`—Trace the interface events.
- `normal`—Trace all normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- `packet`—Trace packet decoding operations
- `parse`—Trace configuration parsing.

- **state**—Trace state transitions.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

**Required Privilege
Level**

storage—To view this statement in the configuration.

storage-control—To add this statement to the configuration.

Administration

- [Routine Monitoring on page 5269](#)
- [Operational Commands on page 5274](#)

Routine Monitoring

- [Monitoring Fibre Channel Interface Load Balancing on page 5269](#)

Monitoring Fibre Channel Interface Load Balancing

You can use operational mode commands to monitor load balancing when the switch is in FCoE-FC gateway mode:

1. [Monitoring the Interface Load-Balancing State on page 5269](#)
2. [Monitoring the Fabric Load-Balancing Algorithm on page 5270](#)

Monitoring the Interface Load-Balancing State

Purpose Monitor the number of sessions, whether load balancing is enabled or disabled, and the load-balancing weight for each native Fibre Channel (FC) interface.



NOTE: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Action To monitor the load-balancing state of the native FC interfaces in the CLI, enter the following CLI command:

```
user@switch> show fibre-channel proxy np-port
```

For example:

```
user@switch> show fibre-channel proxy np-port
Fabric: sanfab1, Fabric-id: 10
NP-Port   State      Sessions    LB state    LB weight
fc-0/0/0.0 online      5           ON          4
fc-0/0/1.0 online      5           ON          4
fc-0/0/2.0 online     10          ON          8
```

```
Fabric: fc_fab2, Fabric-id: 200
```

```

NP-Port      State      Sessions      LB state      LB weight
fc-0/0/44.0  isolated    0             OFF           0

Fabric: fc_fabric_100, Fabric-id: 100
NP-Port      State      Sessions      LB state      LB weight
fc-0/0/46.0  online     1             ON            8

```

Meaning [Table 413 on page 5270](#) summarizes key output fields for the FC interface load-balancing state.

Table 413: Summary of Key FC Interface Load-Balancing Output Fields

Field	Values
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.
NP-Port	NP_Port interface connected to the FCoE forwarder (FCF) or the FC switch.
State	FCID state of the NP_Port interface: <ul style="list-style-type: none"> • online—The port is online and connected to the FC switch. FCoE devices can log in to the FC switch using this port. • isolated—The port is isolated and is not part of the load-balancing function. FCoE devices cannot log in to the FC switch using this port. • offline—The port is offline.
Sessions	Number of active sessions on the NP_Port interface.
LB state	Load-balancing state: <ul style="list-style-type: none"> • On—Load balancing is on • Off—Load balancing is off.
LB weight	Load-balancing weight, which reflects the port speed: <ul style="list-style-type: none"> • 2—Port speed is 2 Gbps. • 4—Port speed is 4 Gbps. • 8—Port speed is 8 Gbps.

The gateway determines the least-loaded interface using the following weighted round-robin (WRR) algorithm:

$$(\text{number-of-sessions} * \text{max-weight}) / \text{weight}$$

where *max-weight* is an internal constant. If the load on the FC interfaces is equal, the session is assigned to the interface with the highest link speed (the greatest weight).

Monitoring the Fabric Load-Balancing Algorithm

Purpose Monitor the type of load-balancing algorithm (simple, ENode-based, or FLOGI-based) used on the native FC interfaces, whether or not automated load rebalancing is enabled, and the load rebalancing state of the fabric.

Action To monitor the load-balancing algorithm used on the native FC interfaces and the load rebalancing state in the CLI, enter the following CLI command:

```
user@switch> show fibre-channel proxy fabric-state
```

For example:

```
user@switch> show fibre-channel proxy fabric-state
Fabric: sanfab1, Fabric-id: 10
Proxy load balance algorithm: Simple, Fabric WWN verification: Yes
Auto load rebalance enabled : No
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Link-up
Last rebalance trigger-time  : Mon Sep 10 21:42:30 2012 usec: 814602
Last rebalance trigger-result: Not-configured

Fabric: fc_fab2, Fabric-id: 200
Proxy load balance algorithm: ENode based, Fabric WWN verification: Yes
Auto load rebalance enabled : No
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Link-up
Last rebalance trigger-time  : Mon Sep 17 17:23:35 2012 usec: 619684
Last rebalance trigger-result: Not-configured

Fabric: fc_fabric_100, Fabric-id: 100
Proxy load balance algorithm: FLOGI based, Fabric WWN verification: No
Auto load rebalance enabled : Yes
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Config-CLI
Last rebalance trigger-time  : Fri Nov 2 08:56:16 2012 usec: 004487
Last rebalance trigger-result: Not-required
```

Meaning You can configure each local FC fabric on an FCoE-FC gateway to use one of three types of load-balancing algorithms, *simple*, *ENode-based*, or *FLOGI-based*. All of the native FC interfaces (NP_Ports) in a particular gateway FC fabric use the same load-balancing algorithm (the load-balancing algorithm is applied on a per-fabric basis).

[Table 414 on page 5271](#) summarizes key output fields for the FC interface load-balancing algorithm and state.

Table 414: show fibre-channel proxy fabric-state Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.

Table 414: show fibre-channel proxy fabric-state Output Fields (*continued*)

Field Name	Field Description
Proxy load balance algorithm	<p>Load-balancing algorithm used on the FCoE-FC gateway FC fabric:</p> <ul style="list-style-type: none"> Simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. <p>On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner.</p> <ul style="list-style-type: none"> ENode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions (VN_Port sessions) associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. <p>On a link load rebalance, all sessions are logged out. When the sessions log in again, they are placed on active NP_Port interfaces in a balanced manner.</p> <ul style="list-style-type: none"> FLOGI-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions (VN_Port sessions) associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. <p>On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner.</p>
Fabric WWN verification	<p>Fabric worldwide name (WWN) verification check state on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> Yes—Fabric WWN verification check is enabled. No—Fabric WWN verification check is disabled.
Auto load rebalance enabled	<p>Automated link load rebalancing configuration for the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> No—Automated load balancing is disabled (default state). Yes—Automated load balancing is enabled.
Last rebalance start-time	<p>Time that the last link load rebalance began on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> Never—The link load has never been rebalanced. Timestamp value—Time the last link load rebalancing started.
Last rebalance end-time	<p>Time that the last link load rebalance ended on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> Never—The link load has never been rebalanced. Timestamp value—Time the last link load rebalancing ended.

Table 414: show fibre-channel proxy fabric-state Output Fields (*continued*)

Field Name	Field Description
Last rebalance trigger	<p>Event that triggered the last link load rebalance on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • None—The link load has never been rebalanced. • Config-CLI—Configure (enable) automated load balancing. • Request-CLI—Rebalance requested from the CLI using the request fibre-channel proxy load-rebalance fabric <i>fabric-name</i> operational command. • Preview-CLI—Rebalancing <i>dry run</i> requested from the CLI using the request fibre-channel proxy load-rebalance dry-run fabric <i>fabric-name</i> operational command. Indicates that the switch completed the dry run. A dry run simulates a link load rebalance and displays a list of sessions that might be affected if you request an actual rebalance. • Link-up—New FC link (NP_Port) up on the FCoE-FC gateway fabric, which causes a rebalance to distribute sessions to the new link. • Restore-complete—If the FC process on the switch restarts, the switch attempts to restore the session state that existed before the restart. When automated rebalance is enabled, restore-complete indicates that the sessions have been restored and rebalanced.
Last rebalance trigger-time	<p>Time that the last link load rebalance was triggered on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—Link load rebalancing has never been triggered. • Timestamp value—Time the last link load rebalancing was triggered.
Last rebalance trigger-result	<p>Result of the last trigger event on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—Link load rebalancing has never been triggered. • Not-configured—Automated rebalancing is not configured on the FCoE-FC gateway fabric. • Not-required—Last rebalance trigger did not require rebalancing the link load (the link load was already balanced across the active NP_Port links). • In-progress—Link load rebalancing is in progress and has not finished yet. • Restore-in-progress—The switch is recovering from an FC process restart and is in the process of restoring the sessions to the active NP_Port links. • Success—Link load rebalancing was successful. • Logged-out-all—All sessions have been logged out. • Preview-complete—The switch has finished simulating a dry run rebalancing request from the CLI (request fibre-channel proxy load-rebalance dry-run fabric <i>fabric-name</i> operational command) and reported the sessions that might be affected if you request an actual link load rebalance. • Fabric-deletion-in-progress—FCoE-FC gateway fabric is in the process of being deleted. <p>NOTE: A trigger event does not necessarily result in a rebalance action. Link load rebalancing only occurs if the NP_Port interface session load is not balanced at the time of the trigger event.</p>

Related Documentation

- [show fibre-channel proxy fabric-state on page 5364](#)
- [show fibre-channel proxy np-port on page 5371](#)

- [Configuring a Fibre Channel Interface on page 5182](#)
- [Defining the Proxy Load-Balancing Algorithm on page 5190](#)
- [Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175](#)
- [Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008](#)
- [Understanding FCoE-FC Gateway Functions on page 4979](#)

Operational Commands

- `clear fibre-channel fc2 statistics`
- `clear fibre-channel fip enode`
- `clear fibre-channel fip statistics`
- `clear fibre-channel fip vn-port`
- `clear fibre-channel flogi statistics`
- `clear fibre-channel proxy statistics`
- `clear fip snooping enode`
- `clear fip snooping statistics`
- `clear fip snooping vlan`
- `clear fip vlan-discovery statistics`
- `request fibre-channel proxy load-rebalance`
- `restart`
- `show dcbx`
- `show dcbx neighbors`
- `show fibre-channel fabric`
- `show fibre-channel fc2 sessions`
- `show fibre-channel fc2 statistics`
- `show fibre-channel fip`
- `show fibre-channel fip enode`
- `show fibre-channel fip fabric`
- `show fibre-channel fip fcf`
- `show fibre-channel fip interface`
- `show fibre-channel fip statistics`
- `show fibre-channel flogi fport`
- `show fibre-channel flogi nport`
- `show fibre-channel flogi statistics`
- `show fibre-channel interfaces`
- `show fibre-channel next-hops`
- `show fibre-channel routes`

- `show fibre-channel proxy fabric-state`
- `show fibre-channel proxy login-table`
- `show fibre-channel proxy np-port`
- `show fibre-channel proxy statistics`
- `show fip snooping`
- `show fip snooping enode`
- `show fip snooping fcf`
- `show fip snooping interface`
- `show fip snooping statistics`
- `show fip snooping vlan`
- `show fip vlan-discovery`
- `show route forwarding-table family fibre-channel`

clear fibre-channel fc2 statistics

Syntax	<code>clear fibre-channel fc2 statistics</code> <code><fabric <i>fabric-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear FC-2 (network layer) Fibre Channel statistics globally or on a specified Fibre Channel fabric.
Options	<code>fabric <i>fabric-name</i></code> —(Optional) Clear FC-2 statistics only on the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fibre-channel fc2 statistics on page 5326• show fibre-channel fc2 sessions on page 5324
List of Sample Output	clear fibre-channel fc2 statistics on page 5276

Sample Output

clear fibre-channel fc2 statistics

```
user@switch> clear fibre-channel fc2 statistics
```


clear fibre-channel fip enode

Syntax	<code>clear fibre-channel fip enode <i>enode-mac</i></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear Fibre Channel over Ethernet (FCoE) node (ENode) information for a specified ENode. This operation deletes the ENode state from the switch database and from the FIP snooping firewall filters, which causes the ENode to lose the connection to the Fibre Channel (FC) fabric and to log in to the fabric again. If you clear an ENode, all VN_Ports associated with that ENode are also cleared and lose their connection to the FC fabric and must log in to the fabric again.
Options	<i>enode-mac</i> —MAC address of the ENode.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel fip enode on page 5333 • clear fibre-channel fip statistics on page 5278 • clear fibre-channel fip vn-port on page 5279
List of Sample Output	clear fibre-channel fip enode on page 5277

Sample Output

clear fibre-channel fip enode

```
user@switch> clear fibre-channel fip enode 00:10:94:00:00:02
```

clear fibre-channel fip statistics

Syntax `clear fibre-channel fip statistics`
 `<fabric fabric-name>`

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Clear Fibre Channel over Ethernet (FCoE) initialization protocol (FIP) statistics.

Options `fabric fabric-name`—(Optional) Clear FIP statistics only on the specified fabric.

Required Privilege Level view

Related Documentation

- [show fibre-channel fip statistics on page 5346](#)
- [show fibre-channel fip on page 5328](#)

List of Sample Output [clear fibre-channel fip statistics on page 5278](#)

Sample Output

clear fibre-channel fip statistics

```
user@switch> clear fibre-channel fip statistics
```

clear fibre-channel fip vn-port

Syntax	<code>clear fibre-channel fip vn-port <i>vn-port</i>--<i>mac</i></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear virtual N_Port (VN_Port) information for a specified VN_Port. This operation deletes the VN_Port state from the switch database and from the FIP snooping firewall filters, which causes the VN_Port to lose its connection to the Fibre Channel fabric and to log in to the fabric again. When you clear a VN_Port, other VN_Ports associated with the same Fibre Channel over Ethernet (FCoE) Node (ENode) are not affected and are not cleared.
Options	<i>vn-port-mac</i> —MAC address of the VN_Port.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel fip enode on page 5333 • clear fibre-channel fip enode on page 5277 • clear fibre-channel fip statistics on page 5278
List of Sample Output	clear fibre-channel fip vn-port on page 5279

Sample Output

clear fibre-channel fip vn-port

```
user@switch> clear fibre-channel fip vn-port 00:10:94:00:00:08
```

clear fibre-channel flogi statistics

Syntax	<code>clear fibre-channel flogi statistics</code> <code><fabric <i>fabric-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear fabric login (FLOGI) statistics globally or on a specified Fibre Channel fabric.
Options	<code>fabric <i>fabric-name</i></code> —(Optional) Clear FLOGI statistics only on the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fibre-channel flogi statistics on page 5354• show fibre-channel flogi fport on page 5350• show fibre-channel flogi nport on page 5352
List of Sample Output	clear fibre-channel flogi statistics on page 5280

Sample Output

clear fibre-channel flogi statistics

```
user@switch> clear fibre-channel flogi statistics
```

clear fibre-channel proxy statistics

Syntax	<code>clear fibre-channel proxy statistics</code> <code><fabric <i>fabric-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear Fibre Channel gateway statistics globally or on a specified Fibre Channel fabric.
Options	<code>fabric <i>fabric-name</i></code> —(Optional) Clear proxy statistics only on the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fibre-channel proxy statistics on page 5374• show fibre-channel proxy login-table on page 5368• show fibre-channel proxy np-port on page 5371
List of Sample Output	clear fibre-channel proxy statistics on page 5281

Sample Output

clear fibre-channel proxy statistics

```
user@switch> clear fibre-channel proxy statistics
```

clear fip snooping enode

Syntax	<code>clear fip snooping enode <i>enode-mac</i></code> <code><vlan <i>vlan-name</i>></code>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear FIP snooping information for the specified FCoE Node (ENode) or (optionally) only on a specified VLAN. This operation deletes the ENode state from the switch database and from the FIP snooping firewall filters, which causes the ENode to lose its connection to the FCoE forwarder (FCF) and to log in to the FCF again.
Options	<i>enode-mac</i> —MAC address of the ENode. <i>vlan vlan-name</i> —(Optional) Name of the VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fip snooping enode on page 5381
List of Sample Output	clear fip snooping enode enode-mac on page 5282

Sample Output

clear fip snooping enode enode-mac

```
user@switch> clear fip snooping enode 00:10:94:00:00:02
```

clear fip snooping statistics

Syntax	<code>clear fip snooping statistics</code> <code><vlan <i>vlan-name</i>></code>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear FIP snooping statistics globally or on a specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fip snooping statistics on page 5391
List of Sample Output	clear fip snooping statistics on page 5283

Sample Output

clear fip snooping statistics

```
user@switch> clear fip snooping statistics
```

clear fip snooping vlan

Syntax	<code>clear fip snooping vlan <i>vlan-name</i></code>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear FIP snooping information for the specified VLAN. This operation deletes all ENode and FCF information for the VLAN from the switch database and causes the ENodes to lose their connections to the FCFs. After clearing a VLAN, the switch relearns all of the FCFs and ENodes on the VLAN, and the ENodes must log in to the FCF again.
Options	<i>vlan-name</i> —Name of the VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fip snooping vlan on page 5394
List of Sample Output	clear fip snooping vlan vlan-name on page 5284

Sample Output

clear fip snooping vlan *vlan-name*

```
user@switch> clear fip snooping vlan fcoevlan1
```


clear fip vlan-discovery statistics

Syntax	clear fip vlan-discovery statistics
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear FIP VLAN discovery statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fip vlan-discovery on page 5398
List of Sample Output	clear fip vlan-discovery statistics on page 5285

Sample Output

clear fip vlan-discovery statistics

```
user@switch> clear fip vlan-discovery statistics
```

request fibre-channel proxy load-rebalance

Syntax	<pre>request fibre-channel proxy load-rebalance <dry-run> fabric <fabric-name> <brief detail></pre>
Release Information	Command introduced in Junos OS Release 12.3 for the QFX Series.
Description	<p>Rebalance the link load on one or more FCoE-FC gateway proxy fabrics (local Fibre Channel fabrics on the gateway) on demand. Load rebalancing is a disruptive action that forces some or all sessions (depending on the configured load-balancing algorithm) to log out and then log in again. When sessions log in again, they are placed on NP_Port interfaces so that the link loads are balanced.</p> <p>Link load rebalancing occurs 10 seconds after you run the rebalancing command, unless another rebalancing trigger occurs before the 10 seconds elapse. If another rebalancing event occurs before the 10-second timer elapses, the timer is extended. Rebalancing occurs a maximum of 30 seconds after you run the rebalancing command, regardless of whether more rebalancing events occur.</p> <p>You can also perform a <i>dry run</i> to see a list of sessions that might be affected (logged out) if you request a load rebalance. A dry run does not rebalance the link loads; it only lists the sessions that might be affected if you rebalance.</p>
Options	<p>dry-run—(Optional) Simulates performing link load rebalancing and displays a list of sessions that might be affected if you rebalance the link loads.</p> <p>fabric <i>fabric-name</i>—Name of the fabric on which you want to rebalance the link loads. If you do not specify a fabric name with the fabric keyword, all fabrics on the FCoE-FC gateway rebalance their link loads.</p> <p>brief detail—(Optional) Display the specified level of output.</p>
Additional Information	Requesting link load rebalancing is a one-time, on-demand operation. You must explicitly request load rebalancing every time you want to rebalance the link loads. Alternatively, you can configure automated load rebalancing if you want the NP_Port links to be rebalanced automatically whenever a load-rebalancing trigger occurs.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Monitoring Fibre Channel Interface Load Balancing on page 5269• Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test) on page 5191• Defining the Proxy Load-Balancing Algorithm on page 5190• Example: Configuring Automated Fibre Channel Interface Load Rebalancing on page 5175• Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric on page 5008

List of Sample Output [request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100 on page 5287](#)

Output Fields [Table 415 on page 5287](#) lists the output fields for the **request fibre-channel proxy load-rebalance dry-run** command. Output fields are listed in the approximate order in which they appear.

Table 415: request fibre-channel proxy load-rebalance dry-run Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.
F-Port	FCoE VLAN interface (VF_Port interface to the FCoE network).
FCID	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet Forwarder (FCoE forwarder) or the Fibre Channel switch.
Port-WWN	Unique worldwide name (WWN) of the VN_Port.
NP-Port	Name of the native Fibre Channel interface.

Sample Output

request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100

```

user@host> request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
F-Port          FCID      Port-WWN          NP-Port
vlan.100        0x8a013a  02:01:00:64:00:00:2a fc-0/0/1.0
vlan.100        0x8a013c  02:01:00:64:00:00:2b fc-0/0/1.0
vlan.100        0x8a0146  02:01:00:64:00:00:2e fc-0/0/1.0
vlan.100        0x8a014c  02:01:00:64:00:00:2f fc-0/0/1.0

```

restart

List of Syntax [Syntax on page 5288](#)

[Syntax \(ACX Series Routers\) on page 5288](#)
[Syntax \(EX Series Switches\) on page 5288](#)
[Syntax \(Routing Matrix\) on page 5289](#)
[Syntax \(J Series Routing Platform\) on page 5289](#)
[Syntax \(TX Matrix Routers\) on page 5289](#)
[Syntax \(TX Matrix Plus Routers\) on page 5289](#)
[Syntax \(MX Series Routers\) on page 5289](#)
[Syntax \(J Series Routers\) on page 5290](#)
[Syntax \(QFX Series\) on page 5290](#)

Syntax restart

```

<adaptive-services | ancpd-service | application-identification | audit-process |
  auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
  class-of-service | clksyncd-service | database-replication | datapath-trace-service
  | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
  | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
  | local-policy-decision-function | mac-validation | mib-process | mobile-ip | mountd-service
  | mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
  packet-triggered-subscribers | peer-selection-service | pgcp-service | pgm |
  pic-services-logging | pki-service | ppp | ppp-service | pppoe |
  protected-system-domain-service | redundancy-interface-process | remote-operations |
  root-system-domain-service | routing <logical-system logical-system-name> | sampling
  | sbc-configuration-process | sdk-service | service-deployment | services | services pgcp
  gateway gateway-name | snmp | soft | static-subscribers | statistics-service |
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
  vrrp | web-management>
<gracefully | immediately | soft>

```

Syntax (ACX Series Routers)

```

restart
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
  class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | immediately | interface-control |
  ipsec-key-management | l2-learning | lacp | link-management | mib-process | mobile-ip |
  mountd-service | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service
  | ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
  sdk-service | secure-neighbor-discovery | service-deployment | services | snmp | soft
  | statistics-service | subscriber-management | subscriber-management-helper | tunnel-oamd
  | vrrp>

```

Syntax (EX Series Switches)

```

restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management

```

	lldpd-service mib-process mountd-service multicast-snooping pgm redundancy-interface-process remote-operations routing secure-neighbor-discovery service-deployment sflow-service snmp vrrp web-management>
Syntax (Routing Matrix)	restart <adaptive-services audit-process chassis-control class-of-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp> <all all-lcc lcc <i>number</i> > <gracefully immediately soft>
Syntax (J Series Routing Platform)	restart <adaptive-services audit-process chassis-control class-of-service dhcp dialer-services dls event-processing firewall interface-control ipsec-key-management isdn-signaling l2-learning l2tp-service mib-process network-access-service pgm ppp pppoe remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp usb-control web-management> <gracefully immediately soft>
Syntax (TX Matrix Routers)	restart <adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp statistics-service> <all-chassis all-lcc lcc <i>number</i> scc> <gracefully immediately soft>
Syntax (TX Matrix Plus Routers)	restart <adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i> > sampling service-deployment snmp statistics-service> <all-chassis all-lcc all-sfc lcc <i>number</i> sfc <i>number</i> > <gracefully immediately soft>
Syntax (MX Series Routers)	restart <adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mobile-ip mountd-service mpls-traceroute msp multicast-snooping named-service nfsd-service

```
packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system logical-system-name> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway gateway-name |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

**Syntax (J Series
Routers)**

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsr | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsr | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |logical-system-name> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsr** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description Restart a Junos OS process.



CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

d datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dlsw—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lACP—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG,

and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management—(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function—(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation—(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation [• Overview of Junos OS CLI Operational Mode Commands on page 75](#)

List of Sample Output [restart interfaces on page 5297](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```


show dcbx

Syntax	show dcbx
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 5300 • Configuring DCBX Autonegotiation on page 5204
Output Fields	Table 416 on page 5299 lists the output fields for the show dcbx command. Output fields are listed in the approximate order in which they appear.

Table 416: show dcbx output fields

Field Name	Field Description
DCBX	Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> • Enabled—DCBX is enabled on the switch or on the specified interface • Disabled—DCBX is disabled on the switch or on the specified interface
Interface	Name of the interface

Sample Output

show dcbx

```

user@switch> show dcbx
DCBX                : Enabled
Interface           DCBX
xe-0/0/9.0          enabled
xe-0/0/32.0         enabled
xe-0/0/36.0         enabled

```

show dcbx neighbors

Syntax	show dcbx neighbors <interface interface-name> <terse>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.3 for EX Series switches.
Description	Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.
Options	none —Display information about all DCBX neighbor interfaces. interface-name —(Optional) Display information for the specified interface. terse —Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring DCBX Autonegotiation on page 5204 • Example: Configuring DCBX Application Protocol TLV Exchange on page 5144 • Example: Configuring an FCoE Transit Switch • Example: Configuring DCBX to Support an iSCSI Application • Understanding DCB Features and Requirements on page 4965 • Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches • dcbx on page 5221
List of Sample Output	show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 5313 show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 5315 show dcbx neighbors terse (QFX Series) on page 5317 show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 5317 show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 5318 show dcbx neighbors (EX4500 Switch: Includes ETS) on page 5319
Output Fields	Table 417 on page 5300 lists the output fields for the show dcbx neighbors command. Output fields are listed in the approximate order in which they appear.

Table 417: show dcbx neighbors Output Fields

Field Name	Field Description
Interface	Name of the interface.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Parent Interface	Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.
Active-application-map	Name of the application map applied to the interface.
Protocol-Mode	<p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> IEEE DCBX Version—The interface uses IEEE DCBX mode. DCBX Version 1.01—The interface uses DCBX version 1.01. <p>NOTE: On interfaces that use the IEEE DCBX mode, the show dcbx neighbors interface <i>interface-name</i> operational command does not include application, PFC, or ETS operational state in the output.</p>
Protocol-State	<p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface. ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.
Local-Advertisement	<p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>
Operational version	Version of the DCBX standard used.
sequence-number	<p>Number of state change messages sent to the peer.</p> <p>If the interface Protocol-State value is in-sync, this number should match the acknowledge-id number in the Peer-Advertisement section.</p> <p>If the interface Protocol-State value is ack-pending, this number does not match the acknowledge-id number in the Peer-Advertisement section.</p>
acknowledge-id	<p>Number of acknowledge messages received from the peer.</p> <p>If the Protocol-State value is in-sync, this number should match the sequence-number value in the Peer-Advertisement section.</p> <p>If the Protocol-State value is ack-pending, this number does not match the sequence-number value in the Peer-Advertisement section.</p>

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Peer-Advertisement	(DCBX Version 1.01 only) Status of advertisements that the peer sends to the local interface.
Operational version	Version of the DCBX standard used.
sequence-number	Number of state change messages the peer sent to the local interface. If this number matches the acknowledge-id number in the Local-Advertisement field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized. If this number does not match the acknowledge-id number in the Local-Advertisement field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.
acknowledge-id	Number of acknowledge messages the peer has received from the local interface. If this number matches the sequence-number value in the Local-Advertisement field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization. If this number does not match the sequence-number value in the Local-Advertisement field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: PFC	Priority-based flow control (PFC) feature DCBX state information.
Protocol-State	(DCBX Version 1.01 only) DCBX protocol state synchronization status: <ul style="list-style-type: none"> • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • not-applicable—PFC autonegotiation is disabled.
Operational State	(DCBX Version 1.01 only) Operational state of the feature: enabled or disabled .
Local-Advertisement	Status of advertisements that the local interface sends to the peer.
Enable	(DCBX Version 1.01 only) State that the local interface advertises to the peer: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> • Yes—The local interface is willing to learn the PFC configuration from the peer. • No—The local interface is not willing to learn the PFC configuration from the peer.
Mac auth Bypass Capability	(IEEE DCBX only) (QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is no .
Error	(DCBX Version 1.01 only) Configuration compatibility error status: <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Operational State	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
Maximum Traffic Classes capable to support PFC	<p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> • 6 (EX Series switches) • 8 (QFX Series)
Code Point	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
Admin Mode	<p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled for the code point. • Disabled—PFC is disabled for the code point.
Operational Mode	<p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> • Enable—PFC is enabled on the code point. • Disable—PFC is disabled on the code point.
Peer-Advertisement	<p>Status of advertisements that the peer sends to the local interface.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> • Yes—The peer is willing to learn the PFC configuration from the local interface. • No—The peer is not willing to learn the PFC configuration from the local interface.
Error	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Operational State	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
Mac auth Bypass Capability	<p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> • Yes—The connected peer supports MAC authentication bypass. • No—The connected peer does not support MAC authentication bypass.
Maximum Traffic Classes capable to support PFC	<p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> • 6 (EX Series switches) • 8 (QFX Series)
Code Point	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
Admin Mode	<p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled for the code point. • Disabled—PFC is disabled for the code point.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: Application	State information for the DCBX application.
Protocol-State	<p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • not-applicable—The local interface is set to no-auto-negotiation (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.
Local-Advertisement	<p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to no-auto-negotiation (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> • Yes—The local interface is willing to learn the FCoE interface state from the peer. • No—The local interface is not willing to learn the FCoE interface state from the peer.
Error	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> • No—No error detected. The local and peer configuration are compatible. • Yes—Error detected. The local and peer configuration are not compatible.
Appl-Name	Name of the application:

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Ethernet-Type	<p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, 0x8906 indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
Socket-Number	<p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
Priority-Field or Priority-Map	<p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>
Status	<p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> • Enabled—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • Disabled—The local configuration and the peer configuration do not match. <p>NOTE: If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p>
Peer-Advertisement	<p>Status of advertisements that the peer sends to the local interface.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> • Yes—The peer is willing to learn the FCoE interface state from the local interface. • No—The peer is not willing to learn the FCoE interface state from the local interface.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Error	(DCBX Version 1.01 only) Configuration compatibility error status: <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.
Appl-Name	Name of the application: <ul style="list-style-type: none"> • FCoE—Fibre Channel over Ethernet
Ethernet-Type	Ethernet type (EtherType) of the application. For example, 0x8906 indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.
Socket-Number	Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.
Priority-Field or Priority-Map	Priority assigned to the application.
Status	(DCBX Version 1.01 only) Peer interface status: <ul style="list-style-type: none"> • Enabled—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • Disabled—The local configuration and the peer configuration do not match.

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: ETS	Enhanced Transmission Selection (ETS) DCBX state information.
Protocol-State	(DCBX Version 1.01 only) ETS protocol state synchronization status: <ul style="list-style-type: none"> • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface. • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.
Operational State	(DCBX Version 1.01 only) Operational state of the feature, enabled or disabled .
Local-Advertisement	Status of advertisements that the local interface sends to the peer.
Enable	(DCBX Version 1.01 only) State that the local interface advertises to the peer: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
TLV Type	(IEEE DCBX only) Type of ETS TLV: <ul style="list-style-type: none"> • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration. • Recommendation-or-Configuration—Advertises both TLVs.
Willing	Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise No for this field): <ul style="list-style-type: none"> • Yes—Local interface is willing to learn the ETS state from the peer. • No—Local interface is not willing to learn the ETS state from the peer.
Credit Based Shaper	

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always No .
Error	(DCBX Version 1.01 only) Configuration error status: <ul style="list-style-type: none">• No—No error. This should always be the switch ETS error state.• Yes—Error detected.
Maximum Traffic Classes capable to support PFC	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
Maximum Traffic Classes supported	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
Code Point	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
Priority-Group	Class-of-service (CoS) priority group (forwarding class set) identification number.
Percentage B/W	Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
Transmission Selection Algorithm	(IEEE DCBX only) The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is ETS .
Peer-Advertisement	Status of advertisements that the peer sends to the local interface.
Enable	

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(DCBX Version 1.01 only) State that the peer advertises to the local interface: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
TLV Type	(IEEE DCBX only) Type of ETS TLV: <ul style="list-style-type: none"> • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration. • Configuration/Recommendation—Advertises both TLVs.
Willing	Willingness of the peer to learn the ETS state from the local interface using DCBX: <ul style="list-style-type: none"> • Yes—Peer is willing to learn the ETS state from the local interface. • No—Peer is not willing to learn the ETS state from the local interface.
Credit Based Shaper	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always No .
Error	(DCBX Version 1.01 only) Configuration error status of the peer: <ul style="list-style-type: none"> • No—No error in peer ETS TLV. • Yes—Error in peer ETS TLV.
Maximum Traffic Classes capable to support PFC	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
Maximum Traffic Classes supported	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
Code Point	

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
Priority-Group	CoS priority group (forwarding class set) identification number.
Percentage B/W	Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
Transmission Selection Algorithm	(IEEE DCBX only) Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is ETS .
PFC	(QFX Series, terse option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> • Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled • Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled • Not Advt—Interface does not advertise PFC to the connected peer
ETS	(terse option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> • Advt—Interface advertises ETS TLVs • Disabled—ETS is disabled on the interface (interface does not advertise ETS)
ETS Rec	(terse option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> • Advt—Peer interface advertises ETS TLVs • Not Advt—Peer interface does not advertise ETS <p>NOTE: When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>

Table 417: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Version	<p>(terse option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> • IEEE—The interface uses IEEE DCBX. • 1.01—The interface uses DCBX version 1.01. <p>When the DCBX version used is the result of autonegotiation, the term (Auto) appears next to the version. For example, IEEE (Auto) indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p>

Sample Output

show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
  Operational version: 1
  sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
  Operational version: 1
  sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode      Operational Mode
000              Disabled        Disable
001              Disabled        Disable
010              Disabled        Disable
011              Enabled         Enable
100              Enabled         Enable
101              Disabled        Disable
110              Disabled        Disable
111              Disabled        Disable

Peer-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode
000              Disabled

```

001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001110	Enabled
iSCSI		3260	10000000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906	N/A	00001110	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1

111	7
Priority-Group	Percentage B/W
0	40%
1	5%

show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

user@switch> **show dcbx neighbors interface xe-0/0/0**

Interface : xe-0/0/0.0 - Parent Interface: ae0.0

Active-application-map: app-map-1

Protocol-Mode: IEEE-DCBX Version

Feature: PFC

Local-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application

Local-Advertisement:

Appl-Name	Ethernet-Type	Socket-Number	Priority-field
FCoE	0x8906		00001110
iSCSI		3260	10000000

Peer-Advertisement:

Appl-Name	Ethernet-Type	Socket-Number	Priority-field
-----------	---------------	---------------	----------------

FCoE	0x8906	N/A	00001110
------	--------	-----	----------

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Recommendation

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0

101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

show dcbx neighbors terse (QFX Series)

```
user@switch> show dcbx neighbors terse
```

Interface	Parent Interface	PFC	ETS	ETS	Version Rec
xe-0/0/8.0	-	Enabled	Advt	Advt	IEEE (Auto)
xe-0/0/9.0	-	Disabled	Disabled		1.01
xe-0/0/11.0	ae0.0	Enabled	Advt	Advt	IEEE (Auto)
xe-0/0/12.0	ae0.0	Enabled	Advt	Advt	IEEE (Auto)
xe-0/0/32.0	-	Enabled	Advt	Not Advt	IEEE
xe-0/0/36.0	-	Not Advt	Advt	Advt	IEEE

show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)

```
user@switch> show dcbx neighbors interface xe-0/0/14
```

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
Operational version: 0
sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
Operational version: 0
sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
Enable: Yes, Willing: No, Error: No
Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Status	Appl-Name	Ethernet-Type	Socket-Number	Priority-Map
Enabled	FCoE	0x8906		00001000

show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

show dcbx neighbors (EX4500 Switch: Includes ETS)

user@switch> show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0
 Protocol-State: in-sync
 Active-application-map: map_iscsi

Local-Advertisement:

Operational version: 0

sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Enabled
001	Enabled
010	Disabled
011	Disabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Enabled
001	Disabled
010	Disabled
011	Disabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00000001	Enabled
iscsi		3260	00000010	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00010000	Enabled
iscsi		3260	00010000	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No
Maximum Traffic Classes supported : 3

Code Point	Priority-Group
000	7
001	7
010	7
011	7
100	7
101	7
110	7
111	7
Priority-Group	Percentage B/W
7	100%

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No
Maximum Traffic Classes supported : 8

Code Point	Priority-Group
000	0
001	1
010	0
011	0
100	2
101	0
110	0
111	0
Priority-Group	Percentage B/W
0	30%
1	40%
2	30%

show fibre-channel fabric

Syntax	<code>show fibre-channel fabric</code> <code><extensive summary></code> <code><fabric-name></code> <code><sort-by (name fabric-id)></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel fabric information.
Options	<p>fabric-name—(Optional) Display output only for the specified fabric.</p> <p>extensive summary—(Optional) Display the specified level of output.</p> <p>sort-by (name fabric-id)—(Optional) Sort output by fabric name or fabric ID.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • fc-fabrics on page 5232 • Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179
List of Sample Output	show fibre-channel fabric on page 5323 show fibre-channel fabric extensive on page 5323
Output Fields	<p>Table 418 on page 5322 lists the output fields for the show fibre-channel fabric command. Output fields are listed in the approximate order in which they appear.</p>

Table 418: show fibre-channel fabric Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Fabric-ID	Identification number of the fabric.	All
Type	Type of fabric. All fabrics are PROXY fabrics.	All
Interfaces	Native Fibre Channel interfaces and FCoE interfaces assigned to the fabric.	All
Created at	Date and time the fabric was created.	extensive
Internal Index	Fabric index internal to Junos OS.	extensive
Origin	Origin information internal to Junos OS.	extensive
Description	Text description of the fabric.	extensive

Table 418: show fibre-channel fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Fabric WWN	Unique WWN of the fabric generated by the FCF.	extensive
Login sessions	Number of FIP login sessions currently running on the fabric.	extensive
Configured max login sessions	Configured maximum number of FIP login sessions permitted on the fabric.	extensive

Sample Output

show fibre-channel fabric

```

user@switch> show fibre-channel fabric
Fabric          Fabric-ID      Type      Interfaces
-----
proxy2          200            PROXY     fc-0/0/0.0
                                     fc-0/0/1.0

```

show fibre-channel fabric extensive

```

user@switch> show fibre-channel fabric extensive
Fabric: proxy2, Created at: Mon Apr 19 14:02:58 2010
Fabric-ID: 200, Internal index: 2, Origin: Static
Description: srv-fabric, Type: PROXY, Fabric WWN: 10:00:00:05:33:51:d7:cd
Login sessions: 200, Configured max login sessions: 500
      fc-0/0/0.0, (untagged)
      fc-0/0/1.0, (untagged)

```

show fibre-channel fc2 sessions

Syntax `show fibre-channel fc2 sessions`
`<fabric fabric-name>`
`<brief | detail>`

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display Fibre Channel FC-2 information.



NOTE: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Options `fabric fabric-name`—(Optional) Display output only for the specified fabric.
`brief | detail`—(Optional) Display the specified level of output.

Required Privilege Level view

Related Documentation

- [show fibre-channel fc2 statistics on page 5326](#)
- [clear fibre-channel fc2 statistics on page 5276](#)

List of Sample Output [show fibre-channel fc2 sessions on page 5325](#)
[show fibre-channel fc2 sessions detail on page 5325](#)

Output Fields [Table 419 on page 5324](#) lists the output fields for the `show fibre-channel fc2 sessions` command. Output fields are listed in the approximate order in which they appear.

Table 419: show fibre-channel fc2 sessions Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Fabric-id	Identification number of the fabric.	All
Interface Name	Name of the interface.	All
Local FCID	Address of the local end of the connection.	All
Far FCID	Address of the far (remote) end of the connection.	All
# Pending Exchanges	Number of pending exchanges for the session.	All
Flags	Flags internal to Junos OS.	detail

Table 419: show fibre-channel fc2 sessions Output Fields (*continued*)

Field Name	Field Description	Level of Output
Refcount	Reference count internal to Junos OS.	detail
Users	Information internal to Junos OS.	detail

Sample Output

show fibre-channel fc2 sessions

```

user@switch> show fibre-channel fc2 sessions
Fabric: fip-proxy, Fabric-id: 1
Interface      Local    Far      # Pending
Name          FCID     FCID     Exchanges
fc-0/0/0.0    *        0xfffffe 0
fc-0/0/1.0    *        0xfffffe 0
fc-0/0/2.0    *        0xfffffe 0

```

show fibre-channel fc2 sessions detail

```

user@switch> show fibre-channel fc2 sessions detail
Fabric: fip-proxy, Fabric-id: 1
Interface Name  fc-0/0/0.0
Local FCID:    *
Far FCID:      0xfffffe
Exchanges:     0
Flags:         SELF_LOCK USER_SYNCED
Refcount:      2
Users:         1

Interface Name  fc-0/0/1.0
Local FCID:    *
Far FCID:      0xfffffe
Exchanges:     0
Flags:         SELF_LOCK USER_SYNCED
Refcount:      2

Interface Name  fc-0/0/2.0
Local FCID:    *
Far FCID:      0xfffffe
Exchanges:     0
Flags:         SELF_LOCK USER_SYNCED
Refcount:      2
Users:         1

```

show fibre-channel fc2 statistics

Syntax	show fibre-channel fc2 statistics <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel FC-2 statistics.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel fc2 sessions on page 5324 • clear fibre-channel fc2 statistics on page 5276
List of Sample Output	show fibre-channel fc2 statistics on page 5327
Output Fields	Table 420 on page 5326 lists the output fields for the show fibre-channel fc2 statistics command. Output fields are listed in the approximate order in which they appear.

Table 420: show fibre-channel fc2 statistics Output Fields

Field Name	Field Description
Global statistics	Statistics for all fabrics.
Frame buffers allocated	Number of frame buffers currently allocated to all fabrics.
Frame buffers freed	Number of frame buffers freed.
Frames dropped	Number of dropped frames.
Fabric statistics	Fabric-specific statistics.
Fabric	Name of the fabric.
Fabric-id	Identification number of the fabric.
Tx-FRJT	Number of fabric frame rejects (F_RJT).
Tx-PRJT	Number of port frame rejects (P_RJT).
Tx-LSRJT	Number of link service rejections.
Tx-ABTS	Number of abort sequence frames sent.
Rx-Drops	Number of received frames dropped.

Table 420: show fibre-channel fc2 statistics Output Fields (*continued*)

Field Name	Field Description
Rx-ABTS	Number of abort sequence frames received.

Sample Output

show fibre-channel fc2 statistics

```
user@switch> show fibre-channel fc2 statistics
Global statistics:

Frame buffers allocated: 60
Frame buffers freed:    60
Frames dropped:         0

Fabric statistics:

Fabric : fip-proxy, Fabric-id: 1
Tx-FRJT: 0
Tx-PRJT: 0
Tx-LSRJT: 0
Tx-ABTS: 0
Rx-Drops: 0
Rx-ABTS: 0
```

show fibre-channel fip

Syntax	show fibre-channel fip <brief detail>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet Initialization Protocol (FIP) information.
Options	brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring FIP on an FCoE-FC Gateway on page 5192 • show fibre-channel fip enode on page 5333 • show fibre-channel fip fabric on page 5337 • show fibre-channel fip fcf on page 5340 • show fibre-channel fip interface on page 5343 • show fibre-channel fip statistics on page 5346 • clear fibre-channel fip statistics on page 5278
List of Sample Output	show fibre-channel fip on page 5330 show fibre-channel fip detail on page 5331
Output Fields	Table 421 on page 5328 lists the output fields for the show fibre-channel fip command. Output fields are listed in the approximate order in which they appear. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Table 421: show fibre-channel fip Output Fields

Field Name	Field Description	Level of Output
Configured max FIP sessions per Node Device	Configured maximum number of FIP sessions permitted on the Node device. For QFabric systems, this is the maximum number of FIP sessions permitted on each Node device in the fabric. For QFX3500 devices, this is the maximum number of FIP sessions permitted on the device.	detail
Node Device	Node device identifier.	detail
Total FIP sessions	Total number of FIP sessions on the FCoE-FC gateway switch.	detail

Table 421: show fibre-channel fip Output Fields (*continued*)

Field Name	Field Description	Level of Output
Total FCoE filters	Total number of FIP filters on the FCoE-FC gateway switch.	detail
Fabric Name	Name of the fabric and in parentheses the fabric ID.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the fabric.	detail
FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail
MAX-SESSIONS-PER-ENODE	Maximum number of concurrent sessions (FLOGI and FDISC combined) that each ENode can instantiate.	detail
FCoE trusted	Whether ports on the FC fabric are trusted or untrusted: <ul style="list-style-type: none"> Yes—Ports on the FC fabric are trusted; FIP snooping is turned off. No—Ports on the FC fabric are not trusted; FIP snooping is turned on. 	detail
Member	Information about an FCF that is a member of the fabric.	All
	<ul style="list-style-type: none"> FCF-MAC 	All
	<ul style="list-style-type: none"> FKA-ADV-PERIOD 	detail
	<ul style="list-style-type: none"> FKA-ADV-D-BIT 	detail
	<ul style="list-style-type: none"> Type 	detail
	<ul style="list-style-type: none"> Priority 	detail
	<ul style="list-style-type: none"> State 	detail

Table 421: show fibre-channel fip Output Fields (*continued*)

Field Name	Field Description	Level of Output
ENode	Information about a connected FCoE node (ENode).	All
• ENode-MAC	MAC address of the connected ENode.	All
• Enode State	Login state internal to Junos OS.	All
• Configured ENode timer	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running ENode timer	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Active FIP Sessions	Number of active FIP sessions on the ENode.	detail
• VN-Port-MAC	MAC address of a VN_Port on the ENode.	All
• Session State	Session state internal to Junos OS.	detail
• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running FKA-ADV	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in milliseconds. This value is always 90 and is not user-configurable.	detail
• Running VN-Port Timer	Running state of the VN_Port keepalive timer in milliseconds.	detail
• FCID	Fibre Channel ID of the VN_Port.	detail
• WWN	Unique worldwide name of the VN_Port.	detail

Sample Output

show fibre-channel fip

```

user@switch> show fibre-channel fip
Fabric Name : proxy2 (200)
Member
FCF-MAC : 00:30:48:b0:ee:d2 (Interface v1an.100)
Enode
Enode-MAC : 00:10:94:00:00:02      State : Logged-in
Session
VN-Port-MAC      : 0e:fc:00:03:00:02
VN-Port-MAC      : 0e:fc:00:03:00:01

```

```

Enode-MAC : 00:10:94:00:00:03      State : Logged-in
Session
VN-Port-MAC      : 0e:fc:00:03:00:04
VN-Port-MAC      : 0e:fc:00:03:00:03

```

show fibre-channel fip detail

```

user@switch> show fibre-channel fip detail
Configured max FIP sessions per Node Device: 2500
Node Device: 0 Total FIP sessions: 4 Total FCoE filters: 4

Fabric Name : proxy2 (200)
FC-MAP      : 0e:fc:00
FKA-ADV-PERIOD : 90000      MAX-SESSIONS-PER-ENODE : 32
FCoE trusted : No

Member
FCF-MAC: 00:30:48:b0:ee:d2 (Interface vlan.100)
FKA-ADV-PERIOD : 90000      FKA-ADV-D-BIT-bit : Off
Type : VF_Port Capable
Priority : 86                State : Enable

ENode
Enode-MAC : 00:10:94:00:00:02  ENode State : Logged-in
Configured ENode timer: 8000   Running ENode timer: 12226
Active FIP Sessions : 2

Session details
VN-Port-MAC      : 0e:fc:00:03:00:02
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90000
Running VN-Port Timer : 213193
FCID             : 0x2c1a01
WWN              : 10:00:00:00:c9:a4:a3:cf

VN-Port-MAC      : 0e:fc:00:03:00:01
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90000
Running VN-Port Timer : 213632
FCID             : 0x2c1a02
WWN              : 10:00:00:00:d9:b4:e3:df

ENode
Enode-MAC : 00:10:94:00:00:03  ENode State : Logged-in
Configured ENode timer: 8000   Running ENode timer: 12254
Active FIP Sessions : 2

Session details
VN-Port-MAC      : 0e:fc:00:03:00:04
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90000
Running VN-Port Timer : 213480
FCID             : 0x2c1a03
WWN              : 21:00:00:c0:dd:11:09:13

```

VN-Port-MAC	: 0e:fc:00:03:00:03
Session state	: Up
Configured FKA-ADV	: 90000
Running FKA-ADV	: 0
Configured VN-Port Timer	: 90000
Running VN-Port Timer	: 214004
FCID	: 0x2c1a04
WWN	: 21:00:00:c0:df:12:08:14

show fibre-channel fip enode

Syntax	show fibre-channel fip enode <i>enode-mac</i> <brief detail> <vn-port-mac <i>vn-port-mac</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified ENode or a specified VN_Port on an ENode.
Options	brief detail —(Optional) Display the specified level of output. <i>enode-mac</i> —Display information for the ENode specified by the MAC address. vn-port-mac <i>vn-port-mac</i> —(Optional) Display information only for the specified VN_Port.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring FIP on an FCoE-FC Gateway on page 5192 • show fibre-channel fip on page 5328 • show fibre-channel fip fabric on page 5337 • show fibre-channel fip fcf on page 5340 • show fibre-channel fip interface on page 5343 • show fibre-channel fip statistics on page 5346 • clear fibre-channel fip enode on page 5277
List of Sample Output	show fibre-channel fip enode on page 5335 show fibre-channel fip enode detail on page 5336
Output Fields	Table 422 on page 5333 lists the output fields for the show fibre-channel fip enode command. Output fields are listed in the approximate order in which they appear. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Table 422: show fibre-channel fip enode Output Fields

Field Name	Field Description	Level of Output
Fabric Name	Name of the fabric and in parentheses the fabric ID.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the fabric.	detail
FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail

Table 422: show fibre-channel fip enode Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAX-SESSIONS-PER-ENODE	Maximum number of concurrent sessions (FLOGI and FDISC combined) that each ENode can instantiate.	detail
FCoE trusted	Whether ports on the FC fabric are trusted or untrusted: <ul style="list-style-type: none"> Yes—Ports on the FC fabric are trusted; FIP snooping is turned off. No—Ports on the FC fabric are not trusted; FIP snooping is turned on. 	detail
Member	Information about an FCF that is a member of the fabric.	All
• FCF-MAC	MAC address used in discovery advertisements.	All
• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off .	detail
• Type	Type of interface: <ul style="list-style-type: none"> VF_Port Capable—Interface can act as a VF_Port interface. 	detail
• Priority	Priority value associated with the switch FCF-MAC. Converged network adapters (CNAs) use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. Value range: 0 through 255.	detail
• State	FIP state on the fabric: <ul style="list-style-type: none"> Enable—FIP is enabled on the fabric. Disable—FIP is disabled on the fabric. 	detail

Table 422: show fibre-channel fip enode Output Fields (*continued*)

Field Name	Field Description	Level of Output
ENode	Information about a connected FCoE node (ENode).	All
• ENode-MAC	MAC address of the connected ENode.	All
• ENode State	Login state internal to Junos OS.	All
• Configured ENode timer	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running ENode timer	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Active FIP Sessions	Number of active FIP sessions on the ENode.	detail
• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
• Session State	Session state internal to Junos OS.	detail
• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail
• FCID	Fibre Channel ID of the VN_Port.	detail
• WWN	Unique worldwide name of the VN_Port.	detail

Sample Output

show fibre-channel fip enode

```

user@switch> show fibre-channel fip enode 00:10:94:00:00:02
Fabric Name : proxy2 (200)
Member
FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
Enode
Enode-MAC : 00:10:94:00:00:02      State : Logged-in
Session
VN-Port-MAC      : 0e:fc:00:03:00:02
VN-Port-MAC      : 0e:fc:00:03:00:01

```

show fibre-channel fip enode detail

```
user@switch> show fibre-channel fip enode 00:10:94:00:00:02 detail
```

```
Fabric Name : proxy2 (200)
```

```
FC-MAP      : 0e:fc:00
```

```
FKA-ADV-PERIOD : 90000      MAX-SESSIONS-PER-ENODE : 32
```

```
FCoE trusted : No
```

Member

```
FCF-MAC: 00:30:48:b0:ee:d2 (Interface vlan.100)
```

```
FKA-ADV-PERIOD : 90000      FKA-ADV-D-BIT-bit : Off
```

```
Type : VF_Port Capable
```

```
Priority : 86      State : Enable
```

ENode

```
Enode-MAC : 00:10:94:00:00:02      ENode State : Logged-in
```

```
Configured ENode timer: 8000      Running ENode timer: 12226
```

```
Active FIP Sessions : 2
```

Session details

```
VN-Port-MAC      : 0e:fc:00:03:00:02
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV   : 0
```

```
Configured VN-Port Timer : 90000
```

```
Running VN-Port Timer : 213193
```

```
FCID              : 0x2c1a01
```

```
WWN               : 10:00:00:00:c9:a4:a3:cf
```

```
VN-Port-MAC      : 0e:fc:00:03:00:01
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV   : 0
```

```
Configured VN-Port Timer : 90000
```

```
Running VN-Port Timer : 213632
```

```
FCID              : 0x2c1a02
```

```
WWN               : 10:00:00:00:d9:b4:e3:df
```

show fibre-channel fip fabric

Syntax	show fibre-channel fip fabric <i>fabric-name</i> <brief detail>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified Fibre Channel fabric.
Options	brief detail —(Optional) Display the specified level of output. <i>fabric-name</i> —Display information for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring FIP on an FCoE-FC Gateway on page 5192 • show fibre-channel fip on page 5328 • show fibre-channel fip enode on page 5333 • show fibre-channel fip fcf on page 5340 • show fibre-channel fip interface on page 5343 • show fibre-channel fip statistics on page 5346
List of Sample Output	show fibre-channel fip fabric proxy2 on page 5338 show fibre-channel fip fabric detail on page 5339
Output Fields	Table 423 on page 5337 lists the output fields for the show fibre-channel fip fabric command. Output fields are listed in the approximate order in which they appear.

Table 423: show fibre-channel fip fabric Output Fields

Field Name	Field Description	Level of Output
Fabric Name	Name of the fabric and in parentheses the fabric ID.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the fabric.	detail
FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail

Table 423: show fibre-channel fip fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Member	Information about an FCF that is a member of the fabric.	All
• FCF-MAC	MAC address used in discovery advertisements.	All
• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off .	detail
• Type	Type of interface: • VF_Port Capable —Interface can act as a VF_Port interface.	detail
ENode	Information about a connected FCoE node (ENode).	All
• ENode-MAC	MAC address of the connected ENode.	All
• State	Login state internal to Junos OS.	All
• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
• Session State	Session state internal to Junos OS.	detail
• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail

Sample Output

show fibre-channel fip fabric proxy2

```

user@switch> show fibre-channel fip fabric proxy2
Fabric Name : proxy2 (200)
  Member
    FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
  ENode

```

```

Enode-MAC : 00:10:94:00:00:02      State : Logged-in
Enode-MAC : 00:10:94:00:00:03      State : Logged-in

```

show fibre-channel fip fabric detail

```
user@switch> show fibre-channel fip fabric proxy2 detail
```

```
Fabric Name : proxy2 (200)
```

```
FC-MAP      : 0e:fc:00
```

```
FKA-ADV-PERIOD : 90000
```

Member

```
FCF-MAC: 00:30:48:b0:ee:d2 (Interface vlan.100)
```

```
FKA-ADV-PERIOD : 90000      FKA-ADV-D-bit : Off
```

```
Type : VF_Port Capable
```

ENode

```
Enode-MAC : 00:10:94:00:00:02      State : Logged-in
```

Session details

```
VN-Port-MAC      : 0e:fc:00:03:00:02
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV    : 0
```

```
Configured VN-Port Timer : 90
```

```
Running VN-Port Timer  : 0
```

```
VN-Port-MAC      : 0e:fc:00:03:00:01
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV    : 0
```

```
Configured VN-Port Timer : 90
```

```
Running VN-Port Timer  : 0
```

ENode

```
Enode-MAC : 00:10:94:00:00:03      State : Logged-in
```

Session details

```
VN-Port-MAC      : 0e:fc:00:03:00:04
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV    : 0
```

```
Configured VN-Port Timer : 90
```

```
Running VN-Port Timer  : 0
```

```
VN-Port-MAC      : 0e:fc:00:03:00:03
```

```
Session state     : Up
```

```
Configured FKA-ADV : 90000
```

```
Running FKA-ADV    : 0
```

```
Configured VN-Port Timer : 90
```

```
Running VN-Port Timer  : 0
```

show fibre-channel fip fcf

Syntax	show fibre-channel fip fcf <i>fcf-mac</i> <brief detail> <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified FCoE forwarder (FCF).
Options	brief detail —(Optional) Display the specified level of output. fabric <i>fabric-name</i> —(Optional) Display FCF information only for the specified fabric. <i>fcf-mac</i> —Display information for the FCF specified by the MAC address.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring FIP on an FCoE-FC Gateway on page 5192 • show fibre-channel fip on page 5328 • show fibre-channel fip enode on page 5333 • show fibre-channel fip fabric on page 5337 • show fibre-channel fip interface on page 5343 • show fibre-channel fip statistics on page 5346
List of Sample Output	show fibre-channel fip fcf on page 5341 show fibre-channel fip fcf detail on page 5342
Output Fields	Table 424 on page 5340 lists the output fields for the show fibre-channel fip fcf command. Output fields are listed in the approximate order in which they appear.

Table 424: show fibre-channel fip fcf Output Fields

Field Name	Field Description	Level of Output
Fabric Name	Name of the fabric and in parentheses the fabric ID.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the fabric.	detail
FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail

Table 424: show fibre-channel fip fcf Output Fields (*continued*)

Field Name	Field Description	Level of Output
Member	Information about an FCF that is a member of the fabric.	All
• FCF-MAC	MAC address used in discovery advertisements.	All
• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off .	detail
• Type	Type of interface: • VF_Port Capable —Interface can act as a VF_Port interface.	detail
ENode	Information about a connected FCoE node (ENode).	All
• ENode-MAC	MAC address of the connected ENode.	All
• State	Login state internal to Junos OS.	All
• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
• Session State	Session state internal to Junos OS.	detail
• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail

Sample Output

show fibre-channel fip fcf

```

user@switch> show fibre-channel fip fcf 00:30:48:b0:ee:d2
Fabric Name : proxy2 (200)
  Member
    FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
  ENode

```

```
Enode-MAC : 00:10:94:00:00:02      State : Logged-in
Enode-MAC : 00:10:94:00:00:03      State : Logged-in
```

show fibre-channel fip fcf detail

```
user@switch> show fibre-channel fip fcf 00:30:48:b0:ee:d2 detail
Fabric Name : proxy2 (200)
FC-MAP      : 0e:fc:00
FKA-ADV-PERIOD : 90000
```

Member

```
FCF-MAC: 00:30:48:b0:ee:d2 (Interface vlan.100)
FKA-ADV-PERIOD : 90000      FKA-ADV-D-bit : Off
Type : VF_Port Capable
```

ENode

```
Enode-MAC : 00:10:94:00:00:02      State : Logged-in
```

Session details

```
VN-Port-MAC      : 0e:fc:00:03:00:02
Session state     : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0
```

```
VN-Port-MAC      : 0e:fc:00:03:00:01
Session state     : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0
```

ENode

```
Enode-MAC : 00:10:94:00:00:03      State : Logged-in
```

Session details

```
VN-Port-MAC      : 0e:fc:00:03:00:04
Session state     : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0
```

```
VN-Port-MAC      : 0e:fc:00:03:00:03
Session state     : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0
```

show fibre-channel fip interface

Syntax	<pre>show fibre-channel fip interface <i>interface-name</i> <brief detail> <enode <i>enode-mac</i>> <fabric <i>fabric-name</i>> <vn-port <i>vn-port-mac</i>></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified interface.
Options	<p>brief detail—(Optional) Display the specified level of output.</p> <p>enode <i>enode-mac</i>—MAC address of the ENode.</p> <p>fabric <i>fabric-name</i>—(Optional) Display interface information only for the specified fabric.</p> <p>interface-name—Display information for the specified interface.</p> <p>vn-port-mac—MAC address of the VN_Port.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring FIP on an FCoE-FC Gateway on page 5192 • show fibre-channel fip on page 5328 • show fibre-channel fip enode on page 5333 • show fibre-channel fip fabric on page 5337 • show fibre-channel fip fcf on page 5340 • show fibre-channel fip statistics on page 5346 • clear fibre-channel fip vn-port on page 5279
List of Sample Output	<p>show fibre-channel fip interface on page 5344</p> <p>show fibre-channel fip interface detail on page 5345</p>
Output Fields	Table 425 on page 5343 lists the output fields for the show fibre-channel fip interface command. Output fields are listed in the approximate order in which they appear.

Table 425: show fibre-channel fip interface Output Fields

Field Name	Field Description	Level of Output
Fabric Name	Name of the fabric and in parentheses the fabric ID.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the fabric.	detail

Table 425: show fibre-channel fip interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail
Member	Information about an FCF that is a member of the fabric.	All
• FCF-MAC	MAC address used in discovery advertisements.	All
• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off .	detail
• Type	Type of interface: • VF_Port Capable —Interface can act as a VF_Port interface.	detail
ENode	Information about a connected FCoE node (ENode).	All
• ENode-MAC	MAC address of the connected ENode.	All
• State	Login state internal to Junos OS.	All
• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
• Session State	Session state internal to Junos OS.	detail
• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail

Sample Output

show fibre-channel fip interface

```

user@switch> show fibre-channel fip interface vlan.100
Fabric Name : proxy2 (200)
Member

```

```

FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
ENode
ENode-MAC : 00:10:94:00:00:02      State : Logged-in
ENode-MAC : 00:10:94:00:00:03      State : Logged-in

```

show fibre-channel fip interface detail

```

user@switch> show fibre-channel fip interface vlan.100 detail
Fabric Name : proxy2 (200)
FC-MAP      : 0e:fc:00
FKA-ADV-PERIOD : 90000

Member
FCF-MAC: 00:30:48:b0:ee:d2 (Interface vlan.100)
FKA-ADV-PERIOD : 90000      FKA-ADV-D-bit : Off
Type : VF_Port Capable

ENode
ENode-MAC : 00:10:94:00:00:02      State : Logged-in

Session details
VN-Port-MAC      : 0e:fc:00:03:00:02
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0

VN-Port-MAC      : 0e:fc:00:03:00:01
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0

ENode
ENode-MAC : 00:10:94:00:00:03      State : Logged-in

Session details
VN-Port-MAC      : 0e:fc:00:03:00:04
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0

VN-Port-MAC      : 0e:fc:00:03:00:03
Session state    : Up
Configured FKA-ADV : 90000
Running FKA-ADV   : 0
Configured VN-Port Timer : 90
Running VN-Port Timer : 0

```

show fibre-channel fip statistics

Syntax	show fibre-channel fip statistics <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel over Ethernet Initialization Protocol (FIP) statistics.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel fip on page 5328 • show fibre-channel fip enode on page 5333 • show fibre-channel fip fabric on page 5337 • show fibre-channel fip fcf on page 5340 • show fibre-channel fip interface on page 5343 • clear fibre-channel fip statistics on page 5278
List of Sample Output	show fibre-channel fip statistics on page 5348
Output Fields	Table 426 on page 5346 lists the output fields for the show fibre-channel fip statistics command. Output fields are listed in the approximate order in which they appear.

Table 426: show fibre-channel fip statistics Output Fields

Field Name	Field Description
Fabric name	Name of the fabric.
Interface name	Name of the FCoE VLAN interface.

Table 426: show fibre-channel fip statistics Output Fields (*continued*)

Field Name	Field Description
FIP Message Type	Type of FIP message for the displayed row of statistics..
• MDS	Number of multicast discovery solicitations.
• UDS	Number of unicast discovery solicitations.
• FLOGI	Number of fabric login (FLOGI) messages.
• FDISC	Number of fabric discovery (FDISC) messages.
• LOGO	Number of fabric logout (LOGO) messages.
• ENODE KA	Number of ENode keepalive messages.
• VN_Port KA	Number of VN_Port keepalive messages.
• MDA	Number of multicast discovery advertisements.
• UDA	Number of unicast discovery advertisements.
• FLOGI ACC	Number of fabric login requests accepted.
• FLOGI RJT	Number of fabric login requests rejected.
• FDISC ACC	Number of fabric discovery requests accepted.
• FDISC RJT	Number of fabric discovery requests rejected.
• LOGO ACC	Number of logout requests accepted.
• LOGO RJT	Number of logout requests rejected.
• CVL	Number of clear virtual links (CVL) messages.
• CVL ALL	Number of CVL all messages.
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of receive errors.

Table 426: show fibre-channel fip statistics Output Fields (*continued*)

Field Name		Field Description
Dropped		Number of dropped messages. NOTE: One cause of dropped messages is that the system limits the number of discovery solicitations (MDS and UDS) it accepts to a maximum of 100 outstanding requests at any given time. If the system has 100 discovery solicitations outstanding, the system does not respond to new discovery solicitations. Instead, the system drops new discovery solicitations and reports the number of dropped discovery solicitations in this field. When there are fewer than 100 outstanding discovery solicitations, the system responds to new requests as usual with a discovery advertisement.
General Statistics	Number of frames recvd with invalid src-mac	Number of frames received that have an invalid source media access control (MAC) address.
	Number of frames recvd with invalid version	Number of FIP frames received with an Invalid FIP version.
	Number of frames recvd with invalid opcode	Number of FIP validation descriptors with an invalid opcode received.
	Number of frames recvd with invalid subcode	Number of FIP validation descriptors with an invalid subcode received.
	Number of frames recvd on inactive FCF	Number of frames received on a logical interface if FIP is not active on that logical interface (for example, if a WWN is not allocated to that logical interface).

Sample Output

show fibre-channel fip statistics

```

user@switch> show fibre-channel fip statistics
Fabric name: proxy2

Interface name: vlan.100
FIP Message type    Received    Sent      Rx errors    Dropped
MDS                  22236      0          0           17089
UDS                   0           0          0            0
FLOGI                1257        0          8            0
FDISC                 0           0          0            0
LOGO                  0           0          0            0
ENODE KA             455         0          6            0
VN_Port KA           22          0          0            0
MDA                   0          243        0            0
UDA                   0          5147       0            0
FLOGI ACC             0          376        0            0
FLOGI RJT             0          881        0            0
FDISC ACC             0           0          0            0
FDISC RJT             0           0          0            0
LOGO ACC              0           0          0            0
LOGO RJT              0           0          0            0

```


CVL	0	374	0	0
CVL ALL	0	380	0	0

General Statistics:

Number of frame recvd with invalid src-mac:	0
Number of frame recvd with invalid version:	0
Number of frame recvd with invalid opcode:	0
Number of frame recvd with invalid subcode:	0
Number of frame recvd on inactive FCF:	0

show fibre-channel flogi fport

Syntax	show fibre-channel flogi fport <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel fabric login (FLOGI) F_Port information.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel flogi nport on page 5352 • show fibre-channel flogi statistics on page 5354
List of Sample Output	show fibre-channel flogi fport on page 5350
Output Fields	<p>Table 427 on page 5350 lists the output fields for the show fibre-channel flogi fport command. Output fields are listed in the approximate order in which they appear.</p>

Table 427: show fibre-channel flogi fport Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Interface	Name of the switch VF_Port interface.
Mac-Address	Media access control (MAC) address of the ENode.
State	Interface physical state: up or down .
Logins	Number of logins to the VF_Port.
NPIV	N_Port ID virtualization (NPIV) state: Yes or No .
FLOGI-Port-WWN	Unique worldwide name (WWN) of the VN_Port performing fabric login (FLOGI) to the switch VF_Port.

Sample Output

show fibre-channel flogi fport

```

user@switch> show fibre-channel flogi fport
Fabric: proxy2
Interface      Mac-Address      State  Logins  NPIV  FLOGI-Port-WWN
vlan.100      00:10:94:00:00:02 Up      2      Yes   20:00:10:94:00:01:00:01
vlan.100      00:10:94:00:00:03 Up      2      Yes   20:00:10:94:00:02:00:01

```


show fibre-channel flogi nport

Syntax	show fibre-channel flogi nport <brief detail> <fabric <i>fabric-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel fabric login (FLOGI) VN_Port information.
Options	brief detail —(Optional) Display the specified level of output. fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel flogi fport on page 5350 • show fibre-channel flogi statistics on page 5354
List of Sample Output	show fibre-channel flogi nport on page 5353 show fibre-channel flogi nport detail on page 5353
Output Fields	Table 428 on page 5352 lists the output fields for the show fibre-channel flogi nport command. Output fields are listed in the approximate order in which they appear.

Table 428: show fibre-channel flogi nport Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Virtual-switch	Name of the fabric.	detail
Interface	Name of the VF_Port interface.	All
FCID	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet Forwarder (FCoE forwarder) or the Fibre Channel switch.	All
Port-WWN	Unique worldwide name (WWN) of the VN_Port.	All
Node-WWN	Unique WWN of the node hosting the VN_Port.	All
State or Flogi-state	Login state internal to Junos OS.	All
FLOGI-Port-WWN	Unique worldwide name (WWN) of the VN_Port performing fabric login (FLOGI) to the switch VF_Port.	detail

Sample Output

show fibre-channel flogi nport

```
user@switch> show fibre-channel flogi nport
Fabric: proxy2
Interface    FCID      Port-WWN      Node-WWN      State
vlan.100     0x030001   20:00:10:94:00:01:00:01  10:00:10:94:00:00:00:01 online
vlan.100     0x030002   20:00:10:94:00:01:00:05  10:00:10:94:00:00:00:01 online
vlan.100     0x030003   20:00:10:94:00:02:00:01  10:00:10:94:00:00:00:02 online
vlan.100     0x030004   20:00:10:94:00:02:00:05  10:00:10:94:00:00:00:02 online
```

show fibre-channel flogi nport detail

```
user@switch> show fibre-channel flogi nport detail
Fabric: proxy2
Virtual-switch: proxy2

Interface: vlan.100
Flogi-state: online
FCID: 0x030001
Port-WWN: 20:00:10:94:00:01:00:01
Node-WWN: 10:00:10:94:00:00:00:01
FLOGI-Port-WWN: 20:00:10:94:00:01:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030002
Port-WWN: 20:00:10:94:00:01:00:05
Node-WWN: 10:00:10:94:00:00:00:01
FLOGI-Port-WWN: 20:00:10:94:00:01:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030003
Port-WWN: 20:00:10:94:00:02:00:01
Node-WWN: 10:00:10:94:00:00:00:02
FLOGI-Port-WWN: 20:00:10:94:00:02:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030004
Port-WWN: 20:00:10:94:00:02:00:05
Node-WWN: 10:00:10:94:00:00:00:02
FLOGI-Port-WWN: 20:00:10:94:00:02:00:01
```

show fibre-channel flogi statistics

Syntax	show fibre-channel flogi statistics <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel fabric login (FLOGI) statistics.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel flogi fport on page 5350 • show fibre-channel flogi nport on page 5352 • clear fibre-channel flogi statistics on page 5280
List of Sample Output	show fibre-channel flogi statistics on page 5355
Output Fields	Table 429 on page 5354 lists the output fields for the show fibre-channel flogi statistics command. Output fields are listed in the approximate order in which they appear.

Table 429: show fibre-channel flogi statistics Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
FLOGI-Server Message type	Type of message: <ul style="list-style-type: none"> • FLOGI—Fabric login (FLOGI) messages. • FDISC—Fabric discovery (FDISC) messages. • FLOGO—Fabric logout messages. • FLOGO-LS-ACC—Fabric logout link service accept messages. • LS-Accept—Link service accept messages. • LS-Reject—Link service reject messages. • invalid—Invalid messages.
Received	Number of messages received for a given message type.
Sent	Number of messages sent for a given message type.
Fabric	Name of the fabric.
Rx errors	Number of receive errors for a given type of message.

Table 429: show fibre-channel flogi statistics Output Fields (*continued*)

Field Name	Field Description
General Statistics	<ul style="list-style-type: none"> • Number of FC2 Header Parse Errors Number of errors parsing the FC-2 header. • Number of FLOGI Parse Errors Number of errors parsing fabric login requests. • Number of FDISC Parse Errors Number of errors parsing fabric discovery requests. • Number of FLOGO Parse Errors Number of errors parsing fabric logout requests. • Number of Logins Discarded as Domain-ID not available Number of discarded logins due to unavailability of a domain ID. • Number of Logins Discarded as FCID not available Number of discarded logins due to the unavailability of a Fibre Channel ID. • Number of FCID requests deferred Number of deferred FCID requests. • Number of deferred FCID requests failed Number of deferred FCID requests that failed.

Sample Output

show fibre-channel flogi statistics

```

user@switch> show fibre-channel flogi statistics
Fabric: proxy2

FLOGI-Server Message type  Received      Sent      Rx errors
FLOGI                      2          0          0
FDISC                      2          0          0
FLOGO                      0          0          0
FLOGO-LS-ACC              0          0          0
LS-Accept                  0          4          0
LS-Reject                  0          0          0
invalid                    0          0          0

General Statistics:

Number of FC2 Header Parse Errors:          0
Number of FLOGI Parse Errors:                0
Number of FDISC Parse Errors:                0
Number of FLOGO Parse Errors:                0
Number of Logins Discarded as Domain-ID not available: 0
Number of Logins Discarded as FCID not available: 0
Number of FCID requests deferred:            0
Number of deferred FCID requests failed:      0

```


show fibre-channel interfaces

Syntax	<brief detail> <fabric <i>fabric-name</i>> show fibre-channel interfaces <i>interface-name</i>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about Fibre Channel (FC) interfaces.
Options	brief detail —(Optional) Display the specified level of output. fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric. <i>interface-name</i> —Display output for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 5099 • Configuring a Physical Fibre Channel Interface on page 5181 • Configuring an FCoE VLAN Interface on page 5185 • Configuring a Fibre Channel Interface on page 5182 • Assigning Interfaces to a Fibre Channel Fabric on page 5187
List of Sample Output	show fibre-channel interfaces on page 5358 show fibre-channel interfaces detail on page 5359
Output Fields	Table 430 on page 5357 lists the output fields for the show fibre-channel interfaces command. Output fields are listed in the approximate order in which they appear.

Table 430: show fibre-channel interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the FC interface.	All
Idx or Index	Interface index internal to Junos OS.	All
Type	Type of interface: <ul style="list-style-type: none"> • FC—Native FC interface • FCOE—Fibre Channel over Ethernet interface 	All
Native Fabric-id	Identification number of the QFX Series fabric.	All
NPIV	N_Port ID virtualization (NPIV) state: Yes or No .	All

Table 430: show fibre-channel interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Config-Mode	User-configured port mode: <ul style="list-style-type: none"> F—The port is configured as a VF_Port, an FCoE port connected to FCoE devices. NP—The port is configured as a proxy N_Port (NP_Port), a native FC port connected to an FC switch. 	All
Oper-Mode	Operational port mode: <ul style="list-style-type: none"> F—The port is operating as a VF_Port, an FCoE port connected to FCoE devices. NP—The port is operating as an NP_Port, a native FC port connected to an FC switch or an FCoE forwarder (FCF). 	All
State	Interface state: up or down .	All
WWN	Unique worldwide name (WWN) of the port.	detail
FSM-State	Finite state machine state, internal to Junos OS.	detail
Class ID	Fibre Channel interface class ID, internal to Junos OS.	detail
BB_SC_N	Buffer-to-buffer state change number.	detail
Tx B2B credits	Number of buffer-to-buffer credits advertised by the neighbor switch that is connected to the FC interface.	detail
Fabric	Name of the fabric.	detail
Remote-MAC	Media access control (MAC) address of the remotely connected FCoE device VN_Port interface.	detail
Tagging	Not used. Value is shown as untagged .	detail
Mode	Logical interface (LIF) mode of operation.	detail
H/W token	Unique identifier for the FCoE VLAN interface, internal to Junos OS.	detail

Sample Output

show fibre-channel interfaces

```

user@switch> show fibre-channel interfaces

```

Interface	Idx	Type	Native Fabric-id	NPIV	Config Mode	Oper Mode	State
fc-0/0/1.0	70	FC	200	YES	NP	NP	up
vlan.100	84	FCOE	200	YES	F	F	up

show fibre-channel interfaces detail

```
user@switch> show fibre-channel interfaces detail
```

```
Interface: fc-0/0/1.0, Index: 70, Type: FC, Native Fabric-id: 200
```

```
NPIV: YES, Config-Mode: NP, Oper-Mode: NP, State: up
```

```
WWN: 10:00:00:15:17:a9:98:64, FSM-State: up, Class ID: 1, BB_SC_N: 0
```

```
Tx B2B credits: 32
```

Fabric	Remote-MAC	Tagging	Mode	Oper state
proxy2	-	untagged	NP	up

```
Interface: vlan.100, Index: 84, Type: FCOE, Native Fabric-id: 200
```

```
NPIV: YES, Config-Mode: F, Oper-Mode: F, State: up
```

```
WWN: 10:00:00:30:48:b0:ee:d2, FSM-State: up
```

```
H/W token: 13
```

Fabric	Remote-MAC	Tagging	Mode	Oper state
proxy2	00:10:94:00:00:02	untagged	VF	up
proxy2	00:10:94:00:00:03	untagged	VF	up

show fibre-channel next-hops

Syntax	show fibre-channel next-hops
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel next-hop route information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show fibre-channel routes on page 5362 • show route forwarding-table family fibre-channel on page 5400
List of Sample Output	show fibre-channel next-hops on page 5360
Output Fields	Table 431 on page 5360 lists the output fields for the show fibre-channel next-hops command. Output fields are listed in the approximate order in which they appear.

Table 431: show fibre-channel next-hops Output Fields

Field Name	Field Description
Type	Type of next hop internal to Junos OS.
State	State of the NP_Port interface: <ul style="list-style-type: none"> • Active—The interface is online. • Deleted—The interface is deleted.
Interface	Name of the interface.
Mac-Address	Media access control (MAC) address of the interface.
Index	Next-hop index identifier.
Ref-count	Reference count internal to Junos OS.
Flags	Flags internal to Junos OS.

Sample Output

show fibre-channel next-hops

```

user@switch> show fibre-channel next-hops
Type  State  Interface  Mac-Address  Index  Ref-count  Flags
----  -
intf  Active fc-0/0/0.0      00:15:17:a9:98:64  674  1          kernel, self
ucast Active vlan.100      0e:fc:00:03:00:01  675  1          kernel, self
ucast Active vlan.100      0e:fc:00:03:00:02  676  1          kernel, self
ucast Active vlan.100      0e:fc:00:03:00:03  677  1          kernel, self
ucast Active vlan.100      0e:fc:00:03:00:04  678  1          kernel, self

```


show fibre-channel routes

Syntax	show fibre-channel routes <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel route information.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show fibre-channel next-hops on page 5360 show route forwarding-table family fibre-channel on page 5400
List of Sample Output	show fibre-channel routes on page 5362
Output Fields	Table 432 on page 5362 lists the output fields for the show fibre-channel routes command. Output fields are listed in the approximate order in which they appear.

Table 432: show fibre-channel routes Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Route-prefix	Route destination.
State	State of the NP_Port interface: <ul style="list-style-type: none"> Active—The interface is online. Deleted—The interface is deleted.
Interface	Name of the interface.
Mac-Address	Media access control (MAC) address of the interface.
Index	Next-hop index identifier.
Flags	Flags internal to Junos OS.

Sample Output

show fibre-channel routes

```

user@switch> show fibre-channel routes
Fabric: proxy2
Route-prefix      State   Interface      Mac-Address      Index  Flags
0x030000/24      Active fc-0/0/0.0     00:15:17:a9:98:64 674    kernel

```

0x030001/24	Active	vlan.100	0e:fc:00:03:00:01	675	kernel
0x030002/24	Active	vlan.100	0e:fc:00:03:00:02	676	kernel
0x030003/24	Active	vlan.100	0e:fc:00:03:00:03	677	kernel
0x030004/24	Active	vlan.100	0e:fc:00:03:00:04	678	kernel

show fibre-channel proxy fabric-state

Syntax	show fibre-channel proxy fabric-state <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display Fibre Channel (FC) proxy fabric state information.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Fibre Channel Interface Load Balancing on page 5269 • show fibre-channel proxy login-table on page 5368 • show fibre-channel proxy np-port on page 5371 • show fibre-channel proxy statistics on page 5374
List of Sample Output	show fibre-channel proxy fabric-state on page 5366 show fibre-channel proxy fabric-state fabric on page 5366
Output Fields	Table 433 on page 5364 lists the output fields for the show fibre-channel proxy fabric-state command. Output fields are listed in the approximate order in which they appear.

Table 433: show fibre-channel proxy fabric-state Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.
Proxy load balance algorithm	<p>Load-balancing algorithm used on the FCoE-FC gateway FC fabric:</p> <ul style="list-style-type: none"> • Simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner. • ENode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. On a link load rebalance, all sessions are logged out. When the sessions log in again, they are placed on active NP_Port interfaces in a balanced manner. • FLOGI-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner.

Table 433: show fibre-channel proxy fabric-state Output Fields (*continued*)

Field Name	Field Description
Fabric WWN verification	<p>Fabric worldwide name (WWN) verification check state on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Yes—Fabric WWN verification check is enabled. • No—Fabric WWN verification check is disabled.
Auto load rebalance enabled	<p>Automated link load rebalancing configuration for the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • No—Automated load balancing is disabled (default state). • Yes—Automated load balancing is enabled.
Last rebalance start-time	<p>Time that the last link load rebalance began on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—The link load has never been rebalanced. • Timestamp value—Time the last link load rebalancing started.
Last rebalance end-time	<p>Time that the last link load rebalance ended on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—The link load has never been rebalanced. • Timestamp value—Time the last link load rebalancing ended.
Last rebalance trigger	<p>Event that triggered the last link load rebalance on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • None—The link load has never been rebalanced. • Config-CLI—Configure (enable) automated load balancing. • Request-CLI—Rebalance requested from the CLI using the request fibre-channel proxy load-rebalance fabric fabric-name operational command. • Preview-CLI—Rebalancing <i>dry run</i> requested from the CLI using the request fibre-channel proxy load-rebalance dry-run fabric fabric-name operational command. Indicates that the switch completed the dry run. A dry run simulates a link load rebalance and displays a list of sessions that might be affected if you request an actual rebalance. • Link-up—New FC link (NP_Port) up on the FCoE-FC gateway fabric, which causes a rebalance to distribute sessions to the new link. • Restore-complete—If the FC process on the switch restarts, the switch attempts to restore the session state that existed before the restart. When automated rebalance is enabled, restore-complete indicates that the sessions have been restored and rebalanced.
Last rebalance trigger-time	<p>Time that the last link load rebalance was triggered on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—Link load rebalancing has never been triggered. • Timestamp value—Time the last link load rebalancing was triggered.

Table 433: show fibre-channel proxy fabric-state Output Fields (*continued*)

Field Name	Field Description
Last rebalance trigger-result	<p>Result of the last trigger event on the FCoE-FC gateway fabric:</p> <ul style="list-style-type: none"> • Never—Link load rebalancing has never been triggered. • Not-configured—Automated rebalancing is not configured on the FCoE-FC gateway fabric. • Not-required—Last rebalance trigger did not require rebalancing the link load (the link load was already balanced across the active NP_Port links). • In-progress—Link load rebalancing is in progress and has not finished yet. • Restore-in-progress—The switch is recovering from an FC process restart and is in the process of restoring the sessions to the active NP_Port links. • Success—Link load rebalancing was successful. • Logged-out-all—All sessions have been logged out. • Preview-complete—The switch has finished simulating a dry run rebalancing request from the CLI (<code>request fibre-channel proxy load-rebalance dry-run fabric fabric-name</code> operational command) and reported the sessions that might be affected if you request an actual link load rebalance. • Fabric-deletion-in-progress—FCoE-FC gateway fabric is in the process of being deleted. <p>NOTE: A trigger event does not necessarily result in a rebalance action. Link load rebalancing only occurs if the NP_Port interface session load is not balanced at the time of the trigger event.</p>

Sample Output

show fibre-channel proxy fabric-state

```

user@switch> show fibre-channel proxy fabric-state
Fabric: san_fab1, Fabric-id: 10
Proxy load balance algorithm: Simple, Fabric WVN verification: Yes
Auto load rebalance enabled : No
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Link-up
Last rebalance trigger-time  : Mon Sep 10 21:42:30 2012 usec: 814602
Last rebalance trigger-result: Not-configured

Fabric: san_fab2, Fabric-id: 20
Proxy load balance algorithm: ENode based, Fabric WVN verification: Yes
Auto load rebalance enabled : No
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Link-up
Last rebalance trigger-time  : Mon Sep 17 17:23:35 2012 usec: 619684
Last rebalance trigger-result: Not-configured

```

show fibre-channel proxy fabric-state fabric

```

user@switch> show fibre-channel proxy fabric-state fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
Proxy load balance algorithm: FLOGI based, Fabric WVN verification: No
Auto load rebalance enabled : Yes
Last rebalance start-time   : Never
Last rebalance end-time     : Never
Last rebalance trigger      : Config-CLI
Last rebalance trigger-time  : Fri Nov 2 08:56:16 2012 usec: 004487
Last rebalance trigger-result: Not-required

```


show fibre-channel proxy login-table

Syntax	<pre>show fibre-channel proxy login-table <brief detail> <fabric <i>fabric-name</i>> <interface <i>interface-name</i>></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel (FC) proxy fabric login table information.
Options	<p>brief detail—(Optional) Display the specified level of output.</p> <p>fabric <i>fabric-name</i>—(Optional) Display output only for the specified fabric.</p> <p>interface <i>interface-name</i>—(Optional) Display output only for the specified interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179 • show fibre-channel proxy fabric-state on page 5364 • show fibre-channel proxy np-port on page 5371 • show fibre-channel proxy statistics on page 5374
List of Sample Output	<p>show fibre-channel proxy login-table on page 5369</p> <p>show fibre-channel proxy login-table detail on page 5369</p>
Output Fields	Table 434 on page 5368 lists the output fields for the show fibre-channel proxy login-table command. Output fields are listed in the approximate order in which they appear.

Table 434: show fibre-channel proxy login-table Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Fabric-id	Fabric ID number.	All
F-Port	<p>One of the following two values:</p> <ul style="list-style-type: none"> • VF_Port interface connected to the Fibre Channel over Ethernet (FCoE) host, shown as the FCoE VLAN interface. • QFX Series FC port that is logged in to the FC switch, shown by a hyphen (-) to indicate that it is not the FCoE device VN_Port. 	All
FCID	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet (FCoE) forwarder (FCF) or the Fibre Channel switch.	All

Table 434: show fibre-channel proxy login-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Port-WWN	Unique worldwide name (WWN) of the VN_Port.	All
Node-WWN	Unique WWN of the node hosting the VN_Ports.	detail
NP-Port	NP_Port interface connected to the FCoE forwarder (FCF) or the Fibre Channel switch.	All
Class	FLOGI service class.	detail
Fabric port WWN	Unique WWN of the fabric port (VF_Port).	detail
Fabric WWN	Unique WWN of the fabric generated by the FCF.	detail

Sample Output

show fibre-channel proxy login-table

```

user@switch> show fibre-channel proxy login-table
Fabric: proxy2, Fabric-id: 200
F-Port          FCID      Port-WWN          NP-Port
-               0x030000  10:00:00:15:17:a9:98:64  fc-0/0/0.0
vlan.100        0x030001  20:00:10:94:00:01:00:01  fc-0/0/0.0
vlan.100        0x030002  20:00:10:94:00:01:00:05  fc-0/0/0.0
vlan.100        0x030003  20:00:10:94:00:02:00:01  fc-0/0/0.0
vlan.100        0x030004  20:00:10:94:00:02:00:05  fc-0/0/0.0

```

show fibre-channel proxy login-table detail

```

user@switch> show fibre-channel proxy login-table detail
Fabric: proxy2, Fabric-id: 200

FCID:          0x030000
F-Port:        -
NP-Port:       fc-0/0/0.0
Port WWN:      10:00:00:15:17:a9:98:64
Node WWN:      20:c8:11:22:33:44:55:66
Class:         3
Fabric port WWN: 10:00:00:15:17:a9:99:48
Fabric WWN:     00:0a:df:ff:0b:11:22:34

FCID:          0x030001
F-Port:        vlan.100
NP-Port:       fc-0/0/0.0
Port WWN:      20:00:10:94:00:01:00:01
Node WWN:      10:00:10:94:00:00:00:01
Class:         3
Fabric port WWN: 10:00:00:15:17:a9:99:48
Fabric WWN:     00:0a:df:ff:0b:11:22:34

FCID:          0x030002
F-Port:        vlan.100
NP-Port:       fc-0/0/0.0
Port WWN:      20:00:10:94:00:01:00:05

```

Node WWN: 10:00:10:94:00:00:00:01
Class: 3
Fabric port WWN: 10:00:00:15:17:a9:99:48
Fabric WWN: 00:0a:df:ff:0b:11:22:34

FCID: 0x030003
F-Port: vlan.100
NP-Port: fc-0/0/0.0
Port WWN: 20:00:10:94:00:02:00:01
Node WWN: 10:00:10:94:00:00:00:02
Class: 3
Fabric port WWN: 10:00:00:15:17:a9:99:48
Fabric WWN: 00:0a:df:ff:0b:11:22:34

FCID: 0x030004
F-Port: vlan.100
NP-Port: fc-0/0/0.0
Port WWN: 20:00:10:94:00:02:00:05
Node WWN: 10:00:10:94:00:00:00:02
Class: 3
Fabric port WWN: 10:00:00:15:17:a9:99:48
Fabric WWN: 00:0a:df:ff:0b:11:22:34

show fibre-channel proxy np-port

Syntax	show fibre-channel proxy np-port <brief detail> <fabric <i>fabric-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel gateway fabric proxy Node Port (NP_Port) information.
Options	brief detail —(Optional) Display the specified level of output. fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric. interface <i>interface-name</i> —(Optional) Display output only for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179 • Monitoring Fibre Channel Interface Load Balancing on page 5269 • show fibre-channel proxy fabric-state on page 5364 • show fibre-channel proxy login-table on page 5368 • show fibre-channel proxy statistics on page 5374
List of Sample Output	show fibre-channel proxy np-port on page 5372 show fibre-channel proxy np-port detail on page 5372
Output Fields	Table 435 on page 5371 lists the output fields for the show fibre-channel proxy np-port command. Output fields are listed in the approximate order in which they appear.

Table 435: show fibre-channel proxy np-port Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Fabric-id	Fabric ID number.	All
NP-Port	NP_Port interface connected to the FCoE forwarder (FCF) or the Fibre Channel switch.	All
State	FCID state of the NP_Port interface.	All
Sessions	Number of active sessions on the NP_Port interface. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.	All

Table 435: show fibre-channel proxy np-port Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured max login sessions	Configured maximum number of FIP login sessions permitted on the interface.	detail
Enodes	Number of ENodes with sessions on the NP_Port.	detail
LB state	Load-balancing state: <ul style="list-style-type: none"> On—Load balancing is on Off—Load balancing is off. 	All
LB weight	Load balance weight, which reflects the port speed: <ul style="list-style-type: none"> 2—Port speed is 2 Gbps. 4—Port speed is 4 Gbps. 8—Port speed is 8 Gbps. 	All
Ref-count	Reference count internal to Junos OS.	detail
Flags	Flags internal to Junos OS. NOTE: When an NP_Port interface reaches its configured maximum number of FIP sessions, the Flags field displays the flag MAX-LOGINS-REACHED .	detail

Sample Output

show fibre-channel proxy np-port

```

user@switch> show fibre-channel proxy np-port
Fabric: proxy1, Fabric-id: 10
NP-Port   State      Sessions  LB state  LB weight
fc-0/0/0.0 online      3         ON        4
fc-0/0/1.0 online      3         ON        4
fc-0/0/2.0 online      3         ON        4
root@junos1> show fibre-channel proxy np-port detail

```

show fibre-channel proxy np-port detail

```

user@switch> show fibre-channel proxy login-table detail
Fabric: proxy1, Fabric-id: 10

NP-Port:          fc-0/0/0.0
State:            online
Sessions:         3
Configured max login sessions: 130
Enodes:           1
LB state:         ON
LB weight:        4
Ref-count:        4
Flags:            UP LB

```



```
NP-Port:          fc-0/0/1.0
State:            online
Sessions:         3
Configured max login sessions: 130
Enodes            2
LB state:         ON
LB weight:        4
Ref-count:        4
Flags:            UP LB

NP-Port:          fc-0/0/2.0
State:            online
Sessions:         130
Configured max login sessions: 130
Enodes            17
LB state:         OFF
LB weight:        4
Ref-count:        131
Flags:            UP MAX-LOGINS-REACHED
```

show fibre-channel proxy statistics

Syntax	show fibre-channel proxy statistics <fabric <i>fabric-name</i>>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel proxy fabric statistics.
Options	fabric <i>fabric-name</i> —(Optional) Display output only for the specified fabric.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring an FCoE-FC Gateway Fibre Channel Fabric on page 5179 • show fibre-channel proxy fabric-state on page 5364 • show fibre-channel proxy login-table on page 5368 • show fibre-channel proxy np-port on page 5371 • clear fibre-channel proxy statistics on page 5281
List of Sample Output	show fibre-channel proxy statistics on page 5375
Output Fields	Table 436 on page 5374 lists the output fields for the show fibre-channel proxy statistics command. Output fields are listed in the approximate order in which they appear.

Table 436: show fibre-channel proxy statistics Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.

Table 436: show fibre-channel proxy statistics Output Fields (*continued*)

Field Name	Field Description
NP-Port Transmit Command Statistics	Transmitted command statistics for the NP_Port.
• Command	Type of command issued on the NP_Port: <ul style="list-style-type: none"> • FLOGI—Fabric login commands issued. • FDISC—Fabric discovery commands issued. • LOGO—Logout commands issued. • Others—Other commands issued.
• Tx	Number of times the command type was transmitted.
• Rx-ACC	Number of times the NP_Port transmitted a receive accept message for the command type.
• Rx-RJT	Number of times the NP_Port transmitted a receive reject message for the command type.
• Abort	Number of times the NP_Port transmitted an abort message for the command type.
NP-Port Receive Command Statistics	Received command statistics for the NP_Port.
• Command	The type of command received on the NP_Port: <ul style="list-style-type: none"> • LOGO—Logout commands issued. • Others—Other commands issued.
• Rx	Number of times the command type was received.
• Tx-ACC	Number of times the NP_Port received a transmit accept message for the command type.
• Tx-RJT	Number of times the NP_Port received a transmit reject message for the command type.
• Abort	Number of times the NP_Port received an abort message for the command type.

Sample Output

show fibre-channel proxy statistics

```
user@switch> show fibre-channel proxy statistics
Fabric: proxy1, Fabric-id: 10
```

```
NP-Port Transmit Command Statistics:
Command      Tx      Rx-ACC  Rx-RJT  Abort
FLOGI        3        3        0        0
FDISC        3        3        0        0
LOGO         0        0        0        0
```

Others	0	0	0	0
--------	---	---	---	---

NP-Port Receive Command Statistics:

Command	Rx	Tx-ACC	Tx-RJT	Abort
LOGO	0	0	0	0
Others	0	0	0	0

show fip snooping

Syntax	show fip snooping <brief detail>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display FIP snooping information.
Options	none —Display FIP snooping information. brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • Example: Configuring an FCoE Transit Switch • show fip snooping enode on page 5381 • show fip snooping fcf on page 5385 • show fip snooping interface on page 5388 • show fip snooping statistics on page 5391 • show fip snooping vlan on page 5394
List of Sample Output	show fip snooping on page 5379 show fip snooping brief (QFX Series) on page 5379 show fip snooping detail (QFX Series) on page 5379 show fip snooping detail on page 5380
Output Fields	Table 437 on page 5377 lists the output fields for the show fip snooping command. Output fields are listed in the approximate order in which they appear.

Table 437: show fip snooping Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Mode	(QFX Series only) Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports. 	All

Table 437: show fip snooping Output Fields (*continued*)

Field Name	Field Description	Level of Output
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
FCF or FCF-MAC	MAC address of the FCF.	All
Session Count or Active Sessions	Current number of virtual link sessions with VN_Ports.	All
VN_Port Count	(QFX Series only) Number of VN_Ports active on an ENode.	brief
Configured FKA-ADV	<p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
Running FKA-ADV	<p>Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
Beacon Period	(QFX Series only) Beacon period interval in milliseconds.	detail
VN2VN Mode	<p>(QFX Series only) Mode of VN2VN_Port snooping:</p> <ul style="list-style-type: none"> Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target. 	detail
ENode-MAC	MAC address of the connected FCoE node (ENode).	All
Interface	Interface connected to the ENode.	detail
VN-Port MAC	MAC address of a VN_Port on the ENode.	All

Table 437: show fip snooping Output Fields (*continued*)

Field Name	Field Description	Level of Output
FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail
Active VN_Ports	(QFX Series only) Number of VN_Ports active on an ENode.	detail
Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping

```

user@switch> show fip snooping
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
Enode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:00:00:05
VN-Port-MAC : 0E:FC:00:00:00:01

```

show fip snooping brief (QFX Series)

```

user@switch> show fip snooping brief
VLAN: vlan100,    Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
FCF: 30:10:94:01:00:00  Session Count: 2
Enode-MAC: 10:10:94:01:00:01
VN-Port-MAC: 0e:fc:00:01:0d:01
VN-Port-MAC: 0e:fc:00:01:0e:01
VLAN: vlan101,    Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
VN-Port-MAC: 0e:fc:00:01:0a:01 Session Count: 2
Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0

```

show fip snooping detail (QFX Series)

```

user@switch> show fip snooping detail
root@sw-pa02v> show fip snooping detail
VLAN: vlan100,    Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
FCF Information
FCF-MAC          : 30:10:94:01:00:00
Active Sessions  : 2
Configured FKA-ADV : 258
Running FKA-ADV   : 188
Enode Information
Enode-MAC: 10:10:94:01:00:01,      Interface: xe-0/0/10
Configured FKA-ADV : 258
Running FKA-ADV    : 230
Session Information

```

```
VN-Port MAC: 0e:fc:00:01:0d:01,   FKA-ADV : 230
VN-Port MAC: 0e:fc:00:01:0e:01,   FKA-ADV : 245

VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Multi-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,   Interface: xe-0/0/10
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 2
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0c:01
      Enode-MAC: 10:10:94:01:00:02,   Interface: xe-0/0/11
      Active VN_Ports : 0
```

show fip snooping detail

```
user@switch> show fip snooping detail
VLAN : fcoevlan1   FC-MAP : 0e:fc:00
  FCF Information
    FCF-MAC : 00:10:94:00:00:01
    Active Sessions : 2
    Configured FKA-ADV : 258
    Running FKA-ADV : 244
    Enode Information
      Enode-MAC : 00:10:94:00:00:02   Interface : xe-0/0/1
      Configured FKA-ADV : 258
      Running FKA-ADV : 248
      Session Information
        VN-Port MAC : 0E:FC:00:00:00:05   FKA-ADV : 264
        VN-Port MAC : 0E:FC:00:00:00:01   FKA-ADV : 260
```


show fip snooping enode

Syntax	show fip snooping enode <i>enode-mac</i> <brief detail> <vlan <i>vlan-name</i>>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display FIP snooping FCoE node (ENode) information.
Options	brief detail —(Optional) Display the specified level of output. <i>enode-mac</i> —Display information for the ENode specified by the MAC address. vlan <i>vlan-name</i> —(Optional) Display FIP snooping information for the ENode on only the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • Example: Configuring an FCoE Transit Switch • show fip snooping on page 5377 • show fip snooping fcf on page 5385 • show fip snooping interface on page 5388 • show fip snooping statistics on page 5391 • show fip snooping vlan on page 5394
List of Sample Output	show fip snooping enode on page 5383 show fip snooping enode brief (QFX Series) on page 5383 show fip snooping enode detail (QFX Series) on page 5383 show fip snooping enode detail on page 5383
Output Fields	Table 438 on page 5381 lists the output fields for the show fip snooping enode command. Output fields are listed in the approximate order in which they appear.

Table 438: show fip snooping enode Output Fields

Field Name	Field Description	Level of Output
ENode and ENode MAC	MAC address of the ENode.	All
VLAN	Name of the VLAN.	All
Interface	Interface connected to the ENode.	All

Table 438: show fip snooping enode Output Fields (*continued*)

Field Name	Field Description	Level of Output
Mode	(QFX Series only) Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports. 	All
VN_Port Count	(QFX Series only) Number of VN_Ports active on an ENode.	brief
Session Count	Current number of virtual link sessions with VN_Ports.	All
Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCoE forwarder (FCF) multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	detail
Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	detail
VN-Port or VN-Port-MAC	MAC address of a VN_Port on the ENode.	All
FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail
FCF or FCF-MAC	MAC address of the FCF to which the VN_Port is connected.	All
Beacon Period	(QFX Series only) Beacon period interval in milliseconds.	detail

Table 438: show fip snooping enode Output Fields (*continued*)

Field Name	Field Description	Level of Output
VN2VN Mode	(QFX Series only) Mode of VN2VN_Port snooping: <ul style="list-style-type: none"> Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target. 	detail
Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping enode

```

user@switch> show fip snooping enode 00:10:94:00:00:02
Enode : 00:10:94:00:00:02   VLAN : vlan1   Interface : xe-0/0/1
      VN-Port-MAC          FCF-MAC
      0E:FC:00:00:00:05    00:10:94:00:00:01
      0E:FC:00:00:00:01    00:10:94:00:00:01

```

show fip snooping enode brief (QFX Series)

```

user@switch> show fip snooping enode 10:10:94:01:00:02 brief
Enode: 10:10:94:01:00:02 ,   VLAN: vlan101,   Interface: xe-0/0/10
Mode: VN2VF Snooping      VN_Port Count: 1
      VN_Port Information
      VN_Port Mac: 0e:fc:00:01:0a:01      Session Count: 2

```

show fip snooping enode detail (QFX Series)

```

user@switch> show fip snooping enode 10:10:94:01:00:02 detail
Enode MAC: 10:10:94:01:00:02,   VLAN: vlan101,   Interface: xe-0/0/10
Mode: VN2VF Snooping      VN_Port Count: 1
Beacon_Period: 90000      VN2VN Mode: Multi-Point
      VN_Port Information
      VN_Port Mac: 0e:fc:00:01:0a:01      Session Count: 2
Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
Vlink far-end VN-Port-MAC: 0e:fc:00:01:0c:01

```

show fip snooping enode detail

```

user@switch> show fip snooping enode 00:10:94:00:00:02 detail
Enode MAC : 00:10:94:00:00:02   VLAN : vlan1   Interface : xe-0/0/1
Configured FKA-ADV : 258      Running FKA-ADV : 213
      Session Information
      VN-Port : 0E:FC:00:00:00:05   FKA-ADV : 229   FCF : 00:10:94:00:00:01
      VN-Port : 0E:FC:00:00:00:01   FKA-ADV : 225   FCF : 00:10:94:00:00:01

```


show fip snooping fcf

Syntax	show fip snooping fcf <i>fcf-mac</i> <brief detail> <vlan <i>vlan-name</i>>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display FIP snooping FCoE forwarder (FCF) information.
Options	brief detail —(Optional) Display the specified level of output. <i>fcf-mac</i> —Display information for the FCF specified by the MAC address. <i>vlan-name</i> —(Optional) Display FIP snooping information for the FCF on only the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • Example: Configuring an FCoE Transit Switch • show fip snooping on page 5377 • show fip snooping enode on page 5381 • show fip snooping interface on page 5388 • show fip snooping statistics on page 5391 • show fip snooping vlan on page 5394
List of Sample Output	show fip snooping fcf on page 5386 show fip snooping fcf detail on page 5386
Output Fields	Table 439 on page 5385 lists the output fields for the show fip snooping fcf command. Output fields are listed in the approximate order in which they appear.

Table 439: show fip snooping fcf Output Fields

Field Name	Field Description	Level of Output
FCF or FCF-MAC	MAC address of the FCoE forwarder.	All
VLAN	Name of the VLAN.	All
Session Count	Current number of virtual link sessions with VN_Ports.	None

Table 439: show fip snooping fcf Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.	detail
Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.	detail
ENode-MAC	MAC address of the connected ENode.	All
• Interface	Interface connected to the ENode.	detail
• Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.	detail
• VN-Port MAC	MAC address of a VN_Port on the ENode.	All
• FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail

Sample Output

show fip snooping fcf

```
user@switch> show fip snooping fcf 00:10:94:00:00:01
FCF : 00:10:94:00:00:01  VLAN : v1an1  Session Count : 2
  ENode-MAC : 00:10:94:00:00:02
    VN-Port-MAC : 0E:FC:00:00:00:05
    VN-Port-MAC : 0E:FC:00:00:00:01
```

show fip snooping fcf detail

```
user@switch> show fip snooping fcf 00:10:94:00:00:01 detail
FCF-MAC : 00:10:94:00:00:01  VLAN : v1an1
Configured FKA-ADV : 258      Running FKA-ADV : 222
  ENode Information
    ENode-MAC : 00:10:94:00:00:02 Interface: xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV : 226
    Session Information
      VN-Port MAC : 0E:FC:00:00:00:05  FKA-ADV : 242
      VN-Port MAC : 0E:FC:00:00:00:01  FKA-ADV : 238
```


show fip snooping interface

Syntax	show fip snooping interface <i>interface-name</i> <brief detail>
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display FIP snooping information for the specified interface.
Options	brief detail —(Optional) Display the specified level of output. interface-name —Display information for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • show fip snooping on page 5377 • show fip snooping enode on page 5381 • show fip snooping fcf on page 5385 • show fip snooping statistics on page 5391 • show fip snooping vlan on page 5394
List of Sample Output	show fip snooping interface on page 5390 show fip snooping interface detail on page 5390
Output Fields	Table 440 on page 5388 lists the output fields for the show fip snooping interface interface-name command. Output fields are listed in the approximate order in which they appear.

Table 440: show fip snooping interface Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
FCF or FCF-MAC	MAC address of the FCF.	All
Session Count or Active Sessions	Current number of virtual link sessions with VN_Ports.	All

Table 440: show fip snooping interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured FKA-ADV	<p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
Running FKA-ADV	<p>Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
ENode-MAC	MAC address of the connected FCoE node (ENode).	All
Interface	Interface connected to the ENode.	detail
Configured FKA-ADV	<p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
Running FKA-ADV	<p>Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>	detail
VN-Port MAC	MAC address of a VN_Port on the ENode.	All

Table 440: show fip snooping interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail

Sample Output

show fip snooping interface

```

user@switch> show fip snooping interface xe-0/0/9.0
VLAN: vlan_100,    FC-MAP: 0e:fc:00
FCF: 30:10:94:01:00:00    Session Count: 1
  Enode-MAC: 10:10:94:01:00:01
    VN-Port-MAC: 0e:fc:00:01:0a:01

```

show fip snooping interface detail

```

user@switch> show fip snooping interface xe-0/0/9.0 detail
VLAN: vlan_100, FC-MAP: 0e:fc:00
FCF Information
FCF-MAC          : 30:10:94:01:00:00
Active Sessions  : 1
Configured FKA-ADV : 368640000
Running FKA-ADV   : 0
  Enode Information
  Enode-MAC: 10:10:94:01:00:01,      Interface: xe-0/0/9
  Configured FKA-ADV : 368640000
  Running FKA-ADV    : 0
    Session Information
    VN-Port MAC: 0e:fc:00:01:0a:01,   FKA-ADV : 0

```

show fip snooping statistics

Syntax	show fip snooping statistics <vlan vlan-name>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display FIP snooping statistics.
Options	vlan vlan-name —(Optional) Display FIP snooping statistics for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring an FCoE Transit Switch • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • show fip snooping on page 5377 • show fip snooping enode on page 5381 • show fip snooping fcf on page 5385 • show fip snooping interface on page 5388 • show fip snooping vlan on page 5394
List of Sample Output	show fip snooping statistics (FIP Snooping) on page 5393 show fip snooping statistics (VN2VN_Port Snooping) on page 5393
Output Fields	Table 441 on page 5391 lists the output fields for the show fip snooping statistics command. Output fields are listed in the approximate order in which they appear.

Table 441: show fip snooping statistics Output Fields

Field Name	Field Description
VLAN	Name of the VLAN for which a set of statistics is displayed.
Mode	(QFX Series only) Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.
Number of MDS	Number of multicast discovery solicitation messages sent on the VLAN.
Number of UDS	Number of unicast discovery solicitation messages sent on the VLAN.

Table 441: show fip snooping statistics Output Fields (*continued*)

Field Name	Field Description
Number of FLOGI	Number of fabric logins on the VLAN.
Number of FDISC	Number of fabric discovery logins on the VLAN.
Number of LOGO	Number of fabric logouts on the VLAN.
Number of ENode-keep-alive	Number of ENode keepalive messages sent on the VLAN.
Number of VNPort-keep-alive	Number of VN_Port keepalive messages sent on the VLAN.
Number of MDA	Number of multicast discovery advertisement messages sent on the VLAN.
Number of UDA	Number of unicast discovery advertisement messages sent on the VLAN.
Number of FLOGI_ACC	Number of fabric logins accepted on the VLAN.
Number of FLOGI_RJT	Number of fabric logins rejected on the VLAN.
Number of FDISC_ACC	Number of fabric discoveries accepted on the VLAN.
Number of FDISC_RJT	Number of fabric discoveries rejected on the VLAN.
Number of LOGO_ACC	Number of fabric logouts accepted on the VLAN.
Number of LOGO_RJT	Number of fabric logouts rejected on the VLAN.
Number of CVL	Number of clear virtual links (CVL) actions on the VLAN.
Number of VN_Port Probes Req	(QFX Series only) Number of multicast N_Port_ID probes sent to the ALL-VN2VN-ENode-MACs multicast address on the VLAN.
Number of VN_Port Claim Notif	(QFX Series only) Number of multicast N_Port_ID claim notifications sent on the VLAN.
Number of VN_Port Beacons	(QFX Series only) Number of multicast beacons sent on the VLAN.
Number of VN_Port Probes Reply	(QFX Series only) Number of replies to N_Port_ID probes sent on the VLAN. Replies are unicast to the ENode MAC address of the probe requester.
Number of VN_Port Claim Reply	(QFX Series only) Number of replies to N_Port_ID claim notifications sent on the VLAN. Replies are unicast to the ENode MAC address of the claim notifier.

Sample Output

show fip snooping statistics (FIP Snooping)

```

user@switch> show fip snooping statistics
VLAN: fcoevlan1      Mode: VN2VF Snooping

Number of MDS:                2
Number of UDS:                2
Number of FLOGI:              2
Number of FDISC:              2
Number of LOGO:               0
Number of Enode-keep-alive: 200
Number of VNPort-keep-alive: 200

Number of MDA:                25
Number of UDA:                2
Number of FLOGI_ACC:          2
Number of FLOGI_RJT:          0
Number of FDISC_ACC:          2
Number of FDISC_RJT:          0
Number of LOGO_ACC:           0
Number of LOGO_RJT:           0
Number of CVL:                0

```

show fip snooping statistics (VN2VN_Port Snooping)

```

user@switch> show fip snooping statistics
VLAN: vlan101      Mode: VN2VN Snooping

Number of VN_Port Probes Req:    3
Number of VN_Port Claim Notif:  3
Number of VN_Port Beacons:      0

Number of VN_Port Probes Reply:  3
Number of VN_Port Claim Reply:   3
Number of FLOGI:                 0
Number of FLOGI_ACC:             0
Number of FLOGI_RJT:             0
Number of FDISC:                 0
Number of FDISC_ACC:             0
Number of FDISC_RJT:             0
Number of LOGO:                  0
Number of LOGO_ACC:              0
Number of LOGO_RJT:              0

```

show fip snooping vlan

Syntax	show fip snooping vlan <i>vlan-name</i> <brief detail>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display FIP snooping VLAN information.
Options	brief detail —(Optional) Display the specified level of output. <i>vlan-name</i> —Display information for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring VN2VF_Port FIP Snooping on an FCoE Transit Switch on page 5199 • Example: Configuring an FCoE Transit Switch • show fip snooping on page 5377 • show fip snooping enode on page 5381 • show fip snooping fcf on page 5385 • show fip snooping interface on page 5388 • show fip snooping statistics on page 5391
List of Sample Output	show fip snooping vlan on page 5396 show fip snooping vlan (QFX Series, VN2VF_Port FIP Snooping) on page 5396 show fip snooping vlan (QFX Series, VN2VN_Port FIP Snooping) on page 5396 show fip snooping vlan detail (QFX Series, VN2VN_Port FIP Snooping) on page 5397 show fip snooping vlan detail on page 5397
Output Fields	Table 442 on page 5394 lists the output fields for the show fip snooping vlan command. Output fields are listed in the approximate order in which they appear.

Table 442: show fip snooping vlan Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All

Table 442: show fip snooping vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Mode	(QFX Series only) Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports. 	All
VN_Port count	(QFX Series only) Number of VN_Ports active on an ENode when the mode is VN2VN_Port FIP snooping.	
FC-MAP	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
Beacon_Period	(QFX Series only) Beacon period interval in milliseconds.	detail
VN2VN Mode	(QFX Series only) Mode of VN2VN_Port snooping: <ul style="list-style-type: none"> Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target. 	detail
FCF or FCF-MAC	MAC address of the FCF.	All
Session Count or Active Sessions	Current number of virtual link sessions with VN_Ports.	All
Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.	detail
Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.	detail

Table 442: show fip snooping vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
ENode-MAC	MAC address of the connected ENode.	All
• Interface	Interface connected to the ENode.	detail
• Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.	detail
• VN-Port MAC	MAC address of a VN_Port on the ENode.	All
• FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail
• Active VN_Ports	(QFX Series only) Number of VN_Ports active on an ENode.	detail
• Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping vlan

```

user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
ENode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:00:00:05
VN-Port-MAC : 0E:FC:00:00:00:01

```

show fip snooping vlan (QFX Series, VN2VF_Port FIP Snooping)

```

user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1    Mode: VN2VF Snooping
FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
ENode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:00:00:05
VN-Port-MAC : 0E:FC:00:00:00:01

```

show fip snooping vlan (QFX Series, VN2VN_Port FIP Snooping)

```

user@switch> show fip snooping vlan vlan101

```



```

VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
  Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
    VN-Port-MAC: 0e:fd:00:00:0a:01 Session Count: 2
  Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0

```

show fip snooping vlan detail (QFX Series, VN2VN_Port FIP Snooping)

```

user@switch> show fip snooping vlan vlan101 detail
VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Multi-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/10
      Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
        Active Sessions : 2
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0c:01
      Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/11
        Active VN_Ports : 0

```

show fip snooping vlan detail

```

user@switch> show fip snooping vlan fcoevlan1 detail
VLAN : fcoevlan1 FC-MAP : 0e:fc:00
FCF Information
FCF-MAC : 00:10:94:00:00:01
Active Sessions : 2
Configured FKA-ADV : 258
Running FKA-ADV : 235
  Enode Information
    Enode-MAC : 00:10:94:00:00:02 Interface : xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV : 239
    Session Information
      VN-Port MAC : 0E:FC:00:00:00:05 FKA-ADV : 255
      VN-Port MAC : 0E:FC:00:00:00:01 FKA-ADV : 251

```

show fip vlan-discovery

Syntax	show fip vlan-discovery (enodes statistics)
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display FCoE VLAN information from the Fibre Channel switch or FCoE forwarder (FCF).
Options	enodes —Display VLAN discovery information for each ENode. statistics —Display VLAN discovery information statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear fip vlan-discovery statistics on page 5285 • Understanding FIP Functions on page 4984 • Understanding FIP Implementation on page 4988
List of Sample Output	show fip vlan-discovery enodes on page 5399 show fip vlan-discovery statistics (QFX3500) on page 5399 show fip vlan-discovery statistics (QFabric Systems) on page 5399
Output Fields	Table 443 on page 5398 lists the output fields for the show fip vlan-discovery command. Output fields are listed in the approximate order in which they appear.

Table 443: show fip vlan-discovery Output Fields

Field Name	Field Description	Level of Output
Enode-MAC	Media access control (MAC) address of the ENode.	enodes
Interface	Name of the interface.	enodes
Unsolicited notification count	Number of unsolicited VLAN discovery notifications.	All
Solicited notification count	Number of solicited VLAN discovery notifications.	statistics
Node Group Name	Displays the name of the Node group on QFabric systems.	statistics
Request count	Number of VLAN discovery requests sent by the ENode. This number should match the Solicited notification count number.	statistics
VLAN tags	Tags of the FIP-enabled VLANs.	enodes

Sample Output

show fip vlan-discovery enodes

```
user@switch> show fip vlan-discovery enodes
```

Enode-MAC	Interface	Unsolicited Notification Count	Vlan Tags
00:10:94:00:00:02	xe-0/0/9.0	0	400

show fip vlan-discovery statistics (QFX3500)

```
user@switch> show fip vlan-discovery statistics
```

```
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

show fip vlan-discovery statistics (QFabric Systems)

```
user@switch> show fip vlan-discovery statistics
```

```
NW-NG-0:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
BBAK0399:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
FCC001:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

show route forwarding-table family fibre-channel

Syntax	<pre>show route forwarding-table family fibre-channel <brief detail extensive> <all> <destination <i>destination-prefix</i>> <interface-name <i>interface-name</i>> <label <i>label</i>> <matching <i>ip-prefix</i>> <multicast> <summary> <table <i>routing-table-name</i>> <vlan <i>vlan-name</i>> <vpn <i>vpn-instance-name</i>></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Fibre Channel family forwarding table route information.
Options	<p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>all—Display all routing forwarding tables.</p> <p>destination <i>destination-prefix</i>—Destination prefix.</p> <p>interface-name <i>interface-name</i>—Name of the interface.</p> <p>label <i>label</i>—Display route entries for the specified label name.</p> <p>matching <i>ip-prefix</i>—Display route entries for the specified IP prefix or length.</p> <p>multicast—Display multicast routes.</p> <p>summary—Display route count instead of details.</p> <p>table <i>routing-table-name</i>—Name of the routing table.</p> <p>vlan <i>vlan-name</i>—Name of the VLAN.</p> <p>vpn <i>vpn-instance-name</i>—Name of the VPN instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show fibre-channel next-hops on page 5360• show fibre-channel routes on page 5362
List of Sample Output	show route forwarding-table family fibre-channel on page 5401
Output Fields	Table 444 on page 5401 lists the output fields for the show route forwarding-table family fibre-channel command. Output fields are listed in the approximate order in which they appear.

Table 444: show route forwarding-table family fibre-channel Output Fields

Field Name	Field Description
Routing table	Name of the routing table.
Destination	Route destination.
Type	Type of route internal to Junos OS.
RtRef	Route reference count internal to Junos OS.
Next hop Type	Type of next hop internal to Junos OS.
Index	Next-hop index identifier.
NhRef	Number of routes that refer to the next hop.
Netif	Interface used to reach the next hop.

Sample Output

show route forwarding-table family fibre-channel

```

user@switch> show route forwarding-table family fibre-channel
Routing table: default.fibre-channel
Fibre Channel:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0x30000/24       user  0          ucst  674  2 fc-0/0/0.0
0x30001/24       user  0          ucst  675  2 vlan.100
0x30002/24       user  0          ucst  676  2 vlan.100
0x30003/24       user  0          ucst  677  2 vlan.100
0x30004/24       user  0          ucst  678  2 vlan.100

```


CHAPTER 62

Troubleshooting

- [Troubleshooting Procedures on page 5403](#)

Troubleshooting Procedures

- [Troubleshooting Dropped FCoE Traffic on page 5403](#)
- [Troubleshooting Fibre Channel Interface Deletion on page 5406](#)
- [Troubleshooting Dropped FIP Traffic on page 5407](#)

Troubleshooting Dropped FCoE Traffic

Problem **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

Cause There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.
5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.
6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

Solution The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 5113 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

Related Documentation

- [show class-of-service congestion-notification on page 5931](#)
- [show class-of-service forwarding-class-set on page 5938](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

Troubleshooting Fibre Channel Interface Deletion

Problem **Description:** You deleted a Fibre Channel (FC) interface at the **[edit interfaces]** hierarchy level, but the commit check fails so the interface is not deleted.

Cause You must first delete the FC interface from the FC fabric on the QFX Series before you can delete the FC interface at the **[edit interfaces]** hierarchy level. You must perform both operations to delete a FC interface.

Solution First delete the interface from the FC fabric and then delete the interface from the QFX Series:

1. Delete the FC interface from the FC fabric to which it belongs:

```
[edit]
user@switch# delete fc-fabrics fabric-name interface interface-name
```

For example, to delete the FC interface **fc-0/0/3.0** from an FC fabric named **sanfab1**:

```
[edit]
user@switch# delete fc-fabrics sanfab1 interface fc-0/0/3.0
```

2. Delete the FC interface at the **[edit interfaces]** hierarchy level:

```
[edit]
user@switch: delete interfaces interface-name
```

For example, to delete the interface **fc-0/0/3.0** from the switch:

```
[edit]
user@switch: delete interfaces fc-0/0/3.0
```

Related Documentation

- [fc-fabrics on page 5232](#)
- [interface on page 5242](#)
- [interfaces on page 2584](#)
- [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

Troubleshooting Dropped FIP Traffic

Problem **Description:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) traffic such as FIP VLAN discovery and notification frames is dropped on the QFX Series.

Cause The interface on which the FIP traffic is dropped does not have a native VLAN configured. FIP VLAN discovery and notification messages are exchanged as untagged packets on the native VLAN. (After the FCoE session with the Fibre Channel switch is established, FCoE traffic uses the FCoE VLAN.)

Solution Check to ensure that every 10-Gigabit Ethernet interface that connects to an FCoE device includes a native VLAN. Configure a native VLAN on all 10-Gigabit Ethernet interfaces that connect to FCoE devices.



NOTE: Make sure that the native VLAN you are using on the QFX Series is the same native VLAN that the FCoE devices use for Ethernet traffic.

To configure a native VLAN on an interface:

1. Set the interface port mode to **tagged-access** if you have not already done so:

```
[edit]
```

```
user@switch# set interfaces interface unit unit family ethernet-switching port-mode
tagged-access
```

For example, to set the port mode to **tagged-access** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure a native VLAN on the interface:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching native-vlan-id
vlan-id
```

For example, to set the native VLAN ID to 1 for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
```

3. Configure the native VLAN if it does not already exist:

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

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

Related Documentation

- [interfaces on page 2584](#)
- [vlans on page 2145](#)
- [Understanding FIP Functions on page 4984](#)

PART 20

Traffic Management

- [Overview on page 5411](#)
- [Configuration on page 5605](#)
- [Administration on page 5915](#)
- [Troubleshooting on page 6065](#)

CHAPTER 63

Overview

- [Software Features Overview on page 5411](#)
- [CoS Overview on page 5435](#)
- [QFabric-Specific CoS Overview on page 5577](#)

Software Features Overview

- [Overview of Junos OS CoS for the QFX Series on page 5412](#)
- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 5414](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 5416](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 5418](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 5420](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Overview of Policers on page 5430](#)

Overview of Junos OS CoS for the QFX Series

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) enables you to divide traffic into classes and set various levels of throughput and packet loss when congestion occurs. You have greater control over packet loss because you can configure rules tailored to your needs.

You can configure CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP) or IEEE 802.1p code-point bits of packets leaving an interface, thus allowing you to tailor packets for the network requirements of the remote peers.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

In designing CoS applications, you must carefully consider your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because QFX Series implements CoS in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet forwarding and switching performance.



NOTE: CoS policies can be enabled or disabled on each switch interface. Also, each physical and logical interface on the switch can have associated custom CoS rules.

When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

This topic describes:

- [CoS Standards on page 5412](#)
- [How Junos CoS Works on page 5413](#)
- [Default CoS Behavior on page 5414](#)

CoS Standards

The following RFCs define the standards for the QFX Series CoS capabilities:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2698, *A Two Rate Three Color Marker*

QFX Series also supports the following data center bridging (DCB) standards to provide the CoS (and other characteristics) Fibre Channel requires for transmitting storage traffic over an Ethernet network:

- IEEE 802.1Qbb, priority-based flow control (PFC)
- IEEE 802.1Qaz, enhanced transmission selection (ETS)
- IEEE 802.1AB (LLDP) extension called Data Center Bridging Capability Exchange Protocol (DCBX)

How Junos CoS Works


Junos CoS works by examining traffic entering at the edge of your network. The switch classifies traffic into defined service groups to provide the special treatment of traffic across the network. For example, you can send voice traffic across certain links and data traffic across other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic to meet the policies of the targeted peer by rewriting the DSCP or IEEE 802.1 code-point bits.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

In [Figure 201 on page 5413](#), Switch A is receiving traffic. As each packet enters, Switch A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined on the switch. This definition allows Switch A to prioritize its resources for servicing the traffic streams it receives. Switch A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the defined traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packet according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can reclassify (rewrite) the CoS code-point bits of the packets before transmitting them.

Figure 201: Packet Flow Across the Network



Default CoS Behavior

If you do not configure CoS settings, the software performs some CoS functions to ensure that the system forwards traffic and protocol packets with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

Related Documentation

- [Overview of Policers on page 4664](#)
- [Understanding Junos CoS Components on page 5436](#)
- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)

Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)

Before you upgrade to Junos OS Release 11.3, you must deactivate the CoS configuration if the CoS configuration includes any of the following features:

- **excess-rate** option
- **strict-high** or **high** priority queues
- Any of the Junos OS Release 11.1 or 11.2 default multidestination forwarding classes



CAUTION: If your CoS configuration contains any of the features listed above and you attempt to upgrade from Junos OS Release 11.1 or 11.2 to a later version without first editing the configuration, the Junos OS might not restart.

Junos OS Release 11.3 and later for QFX Series no longer supports the **excess-rate** statement, the **strict** priority option, or the default multidestination forwarding classes used in Junos OS Release 11.1 and 11.2. In addition, Junos OS Release 11.3 introduces new restrictions on how to configure and use **strict-high** priority queues.

This topic does not describe how to perform the software upgrade procedure. It describes how to deactivate your CoS configuration, edit your CoS configuration, and reactivate your CoS configuration at the appropriate times.

Use the following procedure to upgrade safely from Junos OS Release 11.1 or 11.2 to a later release:

1. Deactivate the CoS configuration *before* you upgrade the software:

```
user@switch# deactivate class-of-service
```
2. Follow the upgrade procedure to Junos OS Release 11.3 or later software.
3. Make the following changes to the CoS configuration while the CoS configuration is still deactivated:

- Remove the **excess-rate** statement from the CoS configuration if you have used it at the **[edit class-of-service schedulers]** or **[edit class-of-service traffic-control-profiles]** hierarchy level.
 - Remove the **strict-high** and **strict** priority queue configurations if you have used them at the **[edit class-of-service schedulers]** hierarchy level.
 - Remove the default multidestination forwarding classes (**mcast-be**, **mcast-af**, **mcast-ef**, and **mcast-nc**) if you have used them at the **[edit class-of-service schedulers]**, **[edit class-of-service rewrite-rules]**, **[edit class-of-service classifiers]**, **[edit class-of-service scheduler-maps]**, or **[edit class-of-service forwarding-class-sets]** hierarchy level. Alternatively, you can change the mapping of the multidestination traffic to use the new default multidestination forwarding class (**mcast**).
4. If desired, configure **strict-high** priority queues in accordance with the Junos OS Release 11.3 or later configuration rules, and map multidestination traffic to the default multidestination forwarding class (**mcast**).
 5. Activate the CoS configuration:

```
user@switch# activate class-of-service
```
 6. Commit the CoS configuration:

```
user@switch# commit
```



NOTE: If you configured the **transmit-rate** option for any queues under the **[edit class-of-service schedulers]** hierarchy level, if the rate is configured as an exact rate in Mbps, we recommend that you reconfigure the **transmit-rate** option as a percentage. This is because the scheduler converts exact rates to percentages, and when the exact rate is below 1 Gbps, some granularity may be lost in the conversion. You can avoid this potential issue by specifying the **transmit-rate** option as a percentage.

Related Documentation

- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 132](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 5416](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 5418](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)

Overview of CoS Upgrade Requirements to Junos OS Release 12.2

Before you upgrade to Junos OS Release 12.2, you might need to edit the class-of-service (CoS) configuration, because the way the QFX Series handles lossless forwarding classes has changed in Junos OS Release 12.2.

By default, the **fcoe** and **no-loss** forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

In Junos OS Release 12.1 and earlier, explicitly setting the lossless **fcoe** and **no-loss** forwarding classes resulted in the same CoS behavior as using the default configuration. However, in Junos OS Release 12.2, the behavior when you explicitly configure the lossless forwarding classes differs from the behavior when you use the default forwarding classes.



NOTE: The default behavior differs from the explicit configuration behavior even if the explicit configuration is exactly the same as the default configuration.

- If you use the default forwarding class configuration for the lossless queues (the configuration does not include explicit setting of the **fcoe** or the **no-loss** forwarding classes), then the **fcoe** and **no-loss** queues behave as lossless queues.

If your CoS configuration does not explicitly configure the **fcoe** and **no-loss** forwarding classes, you can upgrade from Junos OS Release 12.1 to Junos OS Release 12.2, and the behavior of the two lossless queues remains the lossless.

- If your configuration includes statements that explicitly configure the **fcoe** or the **no-loss** forwarding class (using the **[set class-of-service forwarding-classes class class-name queue-num queue-number]** statement), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and behave as lossy (**best-effort**) queues.

If your CoS configuration explicitly configures the **fcoe** and **no-loss** forwarding classes, to retain the lossless behavior of those queues, you need to remove the explicit configuration for these two forwarding classes from the CoS configuration *before* you upgrade.

If you upgrade to Junos OS Release 12.2 and the **fcoe** and **no-loss** forwarding classes are explicitly configured, then those two queues continue to be used, but the traffic is treated as lossy traffic, not lossless traffic. To make the queues for these two forwarding classes lossless, you must delete the explicit forwarding class configuration.



CAUTION: If you explicitly configured the **fcoe** or the **no-loss** forwarding class and you upgrade to Junos OS Release 12.2, the system does not return an upgrade error or a commit error, or a generate a syslog message, to notify you that these forwarding classes are no longer lossless. Traffic mapped to

these forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

Before you upgrade, delete the **fcoe** and **no-loss** forwarding classes from the explicit configuration to preserve the lossless behavior of traffic mapped to these forwarding classes.

- To delete the explicit **fcoe** forwarding class configuration:

```
[edit]
user@switch# delete class-of-service forwarding-class class fcoe queue-num 3
user@switch# commit
```

- To delete the explicit **no-loss** forwarding class configuration:

```
[edit]
user@switch# delete class-of-service forwarding-class class no-loss queue-num 4
user@switch# commit
```



NOTE: If you try to delete these forwarding classes and they have not been explicitly configured on the system, the system returns the message warning: **statement not found**. This simply means that there is no explicit configuration to delete and does not change the lossless behavior of the **fcoe** and **no-loss** forwarding classes.

After you delete the explicit configuration for the **fcoe** and **no-loss** forwarding classes, traffic mapped to those forwarding classes retains its lossless behavior after the upgrade to Junos OS Release 12.2.

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 53](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 5418](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 5420](#)
- [Understanding CoS Forwarding Classes on page 5469](#)
- [Example: Configuring Forwarding Classes on page 5668](#)

Overview of CoS Upgrade Requirements to Junos OS Release 12.3 (QFX3500 and QFX3600 Switches) or to Junos OS Release 13.1 (QFabric Systems)

Before you upgrade to Junos OS Release 12.3 (QFX3500 and QFX3600 switches) or to Junos OS Release 13.1 (QFabric systems), you might need to edit the class-of-service (CoS) configuration, because the way the QFX Series handles lossless forwarding classes has changed from earlier Junos OS releases. (Throughout this document, changes introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.)

- [Support for Six Lossless Forwarding Classes on page 5418](#)
- [Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5419](#)
- [Strict-High Priority Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5420](#)

Support for Six Lossless Forwarding Classes

By default, the *fcoe* and *no-loss* forwarding classes are mapped to output queue 3 and output queue 4, respectively, and to IEEE 802.1p priority 3 (code point 011) and priority 4 (code point 100), respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport in the default configuration.

If you use the default CoS configuration, you do not need to edit the CoS configuration after upgrading to Junos OS Release 12.3 (QFX3500 and QFX3600 switches) or to Junos OS Release 13.1 (QFabric system) because the default CoS configuration is backward-compatible.

Junos OS Release 12.3 increases the support for lossless forwarding classes (priorities) from two forwarding classes to six forwarding classes. To support configuring lossless forwarding classes, Junos OS Release 12.3 introduces a new option to forwarding class configuration: the *no-loss* packet drop attribute.



NOTE: The new *no-loss* packet drop attribute and the previously existing *no-loss* default forwarding class have the same name, but they are not the same. You can use the *no-loss* packet drop attribute on any unicast forwarding class.

If you explicitly configure any lossless forwarding class (including explicitly configuring the default *fcoe* and *no-loss* forwarding classes), you *must* specify the *no-loss* packet drop attribute to obtain lossless behavior. If you do not explicitly configure the *fcoe* and *no-loss* forwarding classes, those forwarding classes remain lossless.

The addition of the *no-loss* packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the *fcoe* and *no-loss* forwarding classes.

If you used the default forwarding class configuration for the *fcoe* and *no-loss* forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything

to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the no-loss packet drop attribute.)

However, if you explicitly configured the fcoe or the no-loss forwarding class by including the **set forwarding-classes class forwarding-class-name queue-num queue-number** at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. In Junos OS Release 12.3 and later, you must include the *no-loss* packet drop attribute in any explicit forwarding class configuration to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the no-loss packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the no-loss drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



NOTE: The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes are lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding class retains its explicit configuration).



NOTE: To achieve lossless behavior for the traffic belonging to any forwarding class, you must also enable PFC on the IEEE 802.1p priority mapped to the forwarding class and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

Scheduling on QFabric System Node Device Fabric (fte) Ports

Junos OS Release 13.1 introduces the ability to configure scheduling on the fabric (fte) ports of QFabric system Node devices. In earlier Junos OS releases, Node device fabric port scheduling was done by default, with no user configuration.

In Junos OS Release 13.1, the default fabric port scheduler configuration is similar to the default scheduler configuration on access interfaces. Similar to the access port default configuration, the default fabric port scheduler supports the five default forwarding

classes (best-effort, fcoe, no-loss, network-control, and mcast). If you configure any new forwarding classes, you must configure scheduling on the fabric ports to allocate bandwidth to those forwarding classes, just as you must configure scheduling on the access ports for user-defined forwarding classes.

Strict-High Priority Scheduling on QFabric System Node Device Fabric (fte) Ports

If a fabric interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best effort, lossless, strict-high priority, and multidestination traffic.

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 53](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 5416](#)

Overview of CoS Changes Introduced in Junos OS Release 11.3

Junos OS Release 11.3 introduces many changes to class-of-service (CoS) functionality and to the CoS default values. This overview summarizes the changes, which other documents describe in detail.



NOTE: Some of the CoS changes are not backward compatible with Junos OS Releases 11.1 and 11.2. “[Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)](#)” on page 53 describes how to upgrade to Junos OS Release 11.3 if you have configured CoS on your QFX3500 switch.

This topic describes the following changes in CoS default values and behavior:

- [CoS Default Value Changes on page 5420](#)
- [Queue Priority Configuration Changes on page 5425](#)
- [Minimum Guaranteed Bandwidth \(Transmit Rate and Guaranteed Rate\) Changes on page 5426](#)
- [Excess Rate Statement Disabled on page 5426](#)
- [Queue Scheduling \(Low and Strict-High Priority Queues\) on page 5427](#)
- [Multidestination Traffic Changes on page 5427](#)

CoS Default Value Changes

The default values of the following CoS components have changed in Junos OS Release 11.3:

- [Default Forwarding Classes on page 5421](#)
- [Default IEEE 802.1p Unicast Classifiers on page 5422](#)
- [Default IEEE 802.1p Multidestination Classifiers on page 5423](#)
- [Default Scheduler on page 5424](#)

Default Forwarding Classes

In Junos OS Releases 11.1 and 11.2, there were eight default forwarding classes, four unicast default forwarding classes and four default multdestination (multicast, broadcast, and destination lookup fail) forwarding classes. [Table 445 on page 5421](#) shows the old default forwarding classes and default queue mapping:

Table 445: Junos OS Release 11.1 and 11.2 Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping
best-effort (be)	Unicast best-effort traffic	0
no-loss	Unicast guaranteed delivery for TCP no-loss traffic	2
fcoe	Unicast guaranteed delivery for FCoE traffic	3
network-control	Unicast network control traffic	7
multicast-best-effort (mcast-be)	Multidestination best-effort traffic	8
multicast-expedited-forwarding (mcast-ef)	Multidestination low-loss, low-latency traffic	9
multicast-assured-forwarding (mcast-af)	Multidestination assured forwarding traffic	10
multicast-network-control (mcast-nc)	Multidestination network control traffic	11

Junos OS Release 11.3 changes the default forwarding classes and queue mapping in the following ways:

- Instead of eight default forwarding classes, there are five default forwarding classes.
- The same four unicast default forwarding classes remain valid, but the default queue mapping of the no-loss forwarding class has changed from queue 2 to queue 4.
- There is now only one default multidestination forwarding class instead of four default multidestination forwarding classes. All multidestination traffic is assigned by default to the default multidestination forwarding class.



NOTE: The rest of the forwarding class characteristics remain the same as before. For example, the QFX Series still supports 12 forwarding classes and 12 output queues. You can still configure a total of eight unicast forwarding classes and four multidestination forwarding classes. The unicast queues are still queues 0 through 7 and the multidestination queues are still queues 8 through 11. Unicast traffic must be mapped to unicast queues, and multidestination traffic must be mapped to multidestination queues. The queue to which a forwarding class is mapped determines whether the forwarding class is unicast or multidestination.

[Table 446 on page 5422](#) shows the default forwarding classes and queue mapping in Junos OS 11.3 and later:

Table 446: Junos OS Release 11.3 Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping
best-effort (be)	Best-effort traffic class	0
fcoe	Guaranteed delivery for FCoE traffic	3
no-loss	Guaranteed delivery for TCP no-loss traffic	4
network-control (nc)	Network control traffic	7
mcast	Multicast traffic	8

Default IEEE 802.1p Unicast Classifiers

In Junos OS Release 11.1 and 11.2, there were default unicast classifiers only for best-effort and network-control traffic, as shown in [Table 447 on page 5422](#):

Table 447: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	best-effort	low
af11 (100)	best-effort	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

Junos OS Release 11.3 introduces new default classifiers for FCoE and no-loss traffic, replacing the best-effort classifiers mapped to IEEE 802.1p code points 011 and 100, respectively, as shown in [Table 448 on page 5422](#):

Table 448: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low

Table 448: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

Default IEEE 802.1p Multidestination Classifiers

In Junos OS Release 11.1 and 11.2, there were default multidestination classifiers for best-effort and network-control traffic, as shown in [Table 449 on page 5423](#):

Table 449: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast-be	low
be1 (001)	mcast-be	low
ef (010)	mcast-be	low
ef1 (011)	mcast-be	low
af11 (100)	mcast-be	low
af12 (101)	mcast-be	low
nc1 (110)	mcast-nc	low
nc2 (111)	mcast-nc	low

Junos OS Release 11.3 replaces the best-effort and network-control multidestination classifiers and maps all IEEE 802.1p code points to the new default multidestination forwarding class, as shown in [Table 450 on page 5423](#):

Table 450: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low

Table 450: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

Default Scheduler

In Junos OS Release 11.1 and 11.2, there were four default schedulers:

- Unicast best effort
- Unicast network control
- Multidestination best effort
- Multidestination network control

[Table 451 on page 5424](#) shows the default scheduler configuration in Junos OS Release 11.1 and 11.2:

Table 451: Junos OS Release 11.1 and 11.2 Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Rate (Extra Bandwidth Sharing)	Priority
Best-effort scheduler (queue 0)	75%	None	25%	Low
Network-control scheduler (queue 7)	5%	None	25%	Low
Best-effort multidestination scheduler (queue 8)	15%	None	25%	Low
Network-control multidestination scheduler (queue 11)	5%	None	25%	Low

Junos OS Release 11.3 replaces the four old classifiers with five new classifiers:

- Unicast best effort
- FCoE
- No loss
- Unicast network control
- Multidestination

There are now four different default unicast classifiers to provide default CoS for lossless queues (FCoE and no-loss traffic). Because there is only one default multidestination forwarding class in Junos OS Release 11.3, there is only one default multidestination classifier for all multidestination traffic. Also, the excess rate default value is removed from the scheduler because the **excess-rate** statement is no longer supported, as described elsewhere in this document. [Table 452 on page 5425](#) shows the default scheduler configuration in Junos OS Releases 11.3:

Table 452: Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority
Best-effort scheduler (queue 0)	5%	None	5%	Low
FCoE scheduler (queue 3)	35%	None	35%	Low
No-loss scheduler (queue 4)	35%	None	35%	Low
Network-control scheduler (queue 7)	5%	None	5%	Low
Multidestination scheduler (queue 8)	20%	None	20%	Low



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

Queue Priority Configuration Changes

In Junos OS Release 11.1 and 11.2, you could configure strict-high priority queues with a guaranteed minimum bandwidth and configure forwarding class sets (priority groups) with a mix of low priority and strict-high priority queues. In Junos OS Release 11.3 and later, these configurations are invalid, and several other changes have also been implemented:

- Priority configuration in Junos OS Release 11.1 and 11.2 provided three priority levels: **strict-high**, **high**, and **low**. In Junos OS Release 11.3, the **high** priority option has been removed. Only the **strict-high** and **low** priority options are valid in Release 11.3.
- Minimum guaranteed bandwidth (transmit rate) is not allowed on strict-high priority queues. Minimum guaranteed bandwidth (guaranteed rate) is not allowed on forwarding class sets that contain strict-high priority queues.
- You cannot configure a multidestination queue as a strict-high priority queue. You cannot configure a queue as a strict-high priority queue if it belongs to the multidestination forwarding class set.
- Only one forwarding class set can contain strict-high priority queues. If you want to configure a strict-high priority queue, you must also configure a separate forwarding class set for the strict-high priority queue. A forwarding class set cannot contain a mixture of low priority and strict-high priority queues.

The rest of the queue priority characteristics remain the same as before. For example, you can configure only one queue as a strict-high priority queue.



NOTE: If you have configured strict-high or high priority queues in Junos OS Release 11.1 or 11.2, the changes in Release 11.3 are not backward compatible. Please read [“Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)”](#) on page 53 before you upgrade to Release 11.3.

Minimum Guaranteed Bandwidth (Transmit Rate and Guaranteed Rate) Changes

The following restrictions have been placed on minimum guaranteed bandwidth configuration in Junos OS Release 11.3:

- You cannot configure a guaranteed minimum bandwidth (transmit rate) for strict-high priority queues.
- Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for forwarding class sets that include strict-high priority queues.
- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

Excess Rate Statement Disabled

The **excess-rate** statement has been disabled in Junos OS Release 11.3. Excess rate was used to specify the way extra bandwidth was shared among queues.

The **excess-rate** statement was used at the **[edit class-of-service schedulers]** hierarchy level for queue scheduling configuration and at the **[edit class-of-service traffic-control-profiles]** hierarchy level for forwarding class set scheduling configuration.

In Junos OS Release 11.3, extra bandwidth sharing among queues is proportional to the minimum guaranteed bandwidth (transmit rate) of the queue. Extra bandwidth sharing among forwarding class sets (priority groups) is proportional to the minimum guaranteed bandwidth (guaranteed rate) of the forwarding class set.



NOTE: If you have configured the `excess-rate` option in Junos OS Release 11.1 or 11.2, the changes in Release 11.3 are not backward compatible. Please read [“Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)”](#) on page 53 before you upgrade to Release 11.3.

Queue Scheduling (Low and Strict-High Priority Queues)

In Junos OS Release 11.1 and 11.2, if you configured a guaranteed minimum bandwidth (transmit rate) for low-priority queues, the low-priority queues received their guaranteed minimum bandwidth from the same bandwidth pool as the strict-high priority queue, using round-robin scheduling. Until the minimum bandwidth requirements of all queues were met, the strict-high priority queue and low-priority queues that had a guaranteed minimum bandwidth were treated equally. After the minimum bandwidth requirements of all queues were met, the strict-high priority queue received as much of the leftover bandwidth as it needed. This meant that the only way to ensure that a strict-high priority queue received all of the bandwidth it needed was not to configure a guaranteed minimum bandwidth for other queues.

In Junos OS Release 11.3 and later, queue scheduling has changed so that queues receive bandwidth in the following sequence:

1. The strict-high priority queue receives all of the bandwidth it needs before any other queue is served. The strict-high priority queue can take the full port bandwidth if necessary and can starve other queues on the port.
2. The guaranteed minimum bandwidth (transmit rate) of low-priority queues is served until the minimum is met or the queues are empty.
3. All other low-priority queues and needs that exceed the minimum bandwidth are served.

Multidestination Traffic Changes

The changes to the default forwarding classes and classifiers affects multidestination traffic handling in Junos OS Release 11.3:

- The number of default multidestination forwarding classes has been reduced from four default multidestination forwarding classes in Junos OS Release 11.1 and 11.2 to one default multidestination in Release 11.3 (see [Table 446 on page 5422](#)).
- The default classifier configuration for multidestination traffic has changed so that there is now one default classifier for all multidestination traffic (see [Table 450 on page 5423](#)).
- By default, all IEEE 802.1p code points map to the default multidestination forwarding class.

- The default scheduler for multidestination traffic has changed so that there is now one default scheduler for all multidestination traffic (see [Table 452 on page 5425](#)).
- You cannot configure multidestination queues as strict-high priority queues and you cannot include strict-high priority queues in a forwarding class set that contains multidestination queues.

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 53](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

Overview of CoS Changes Introduced in Junos OS Release 12.2

Junos OS Release 12.2 introduces some changes to class-of-service (CoS) functionality and to the CoS default values. This overview summarizes the changes, which other documents describe in detail.

This topic describes the following changes in CoS default values and behavior:

- [Lossless Forwarding Classes \(fcoe and no-loss\) on page 5428](#)
- [Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes on page 5429](#)
- [CoS for Layer 3 Physical Interfaces on page 5429](#)
- [DSCP IPv6 Classifiers and Rewrite Rules on page 5429](#)

Lossless Forwarding Classes (fcoe and no-loss)

The way the QFX Series handles lossless forwarding classes (the **fcoe** and **no-loss** forwarding classes) changes in Junos OS Release 12.2. In Junos OS Release 12.2 and in earlier releases, by default, the **fcoe** and **no-loss** forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

In earlier releases, explicitly setting the lossless **fcoe** and **no-loss** forwarding classes resulted in the same CoS behavior as using the default configuration. However, in Junos OS Release 12.2, the behavior when you explicitly configure the lossless forwarding classes differs from the behavior when you use the default forwarding classes.



NOTE: The default behavior differs from the explicit configuration behavior even if the explicit configuration is exactly the same as the default configuration.

If you use the default forwarding class settings for the lossless queues (the configuration does not include explicit setting of the **fcoe** or the **no-loss** forwarding classes), then the **fcoe** and **no-loss** queues behave as lossless queues. When you upgrade to Junos OS Release 12.2, traffic assigned to the **fcoe** and **no-loss** queues continues to be treated as lossless traffic.

If your configuration explicitly sets the **fcoe** or the **no-loss** forwarding class (**set class-of-service forwarding-classes class class-name queue-num queue-number**), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and behave as lossy (**best-effort**) queues. To retain lossless treatment of the **fcoe** and **no-loss** queues, delete the explicit lossless forwarding class configuration before you upgrade to Junos OS Release 12.2.



CAUTION: If you explicitly configured the **fcoe** or the **no-loss** forwarding class, and you upgrade to Junos OS Release 12.2, the system does not return an upgrade error or a commit error, or a generate a syslog message, to notify you that these forwarding classes are no longer lossless. Traffic mapped to these forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes

The default maximum transmission unit (MTU) the system uses for buffer headroom calculation is 2500 bytes for traffic classified into the **fcoe** forwarding class or the **no-loss** forwarding class.

In Junos OS Release 12.2, the default MTU used for buffer headroom calculation for the **fcoe** and **no-loss** forwarding classes remains 2500 bytes. However, if the buffer is filled, in Junos OS Release 12.2 you might experience commit failures.

CoS for Layer 3 Physical Interfaces

Before Junos OS Release 12.2, the QFX Series supported only Layer 2 CoS. Junos OS Release 12.2 introduces CoS support for Layer 3 traffic at the physical interface level.

If a physical Layer 3 interface has at least one logical interface configured on it, you can configure Layer 3 CoS for the physical interface. The CoS configured on the physical interface applies to all of the logical Layer 3 interfaces on that physical interface. The system does not support Layer 3 CoS configuration on individual Layer 3 logical interfaces.

DSCP IPv6 Classifiers and Rewrite Rules

Junos OS Release 12.2 introduces support for DSCP IPv6 classifiers and rewrite rules. The existing DSCP IP default classifier is now also the DSCP IPv6 default classifier.

You can configure and apply DSCP IPv6 classifiers and DSCP IPv6 rewrite rules to Layer 2 logical interfaces and to Layer 3 physical interfaces.



NOTE: DSCP IPv6 classifiers are not supported for multidestination (multicast, broadcast, and destination lookup fail) traffic.

**Related
Documentation**

- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 5416](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 5420](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 5430](#)
- [Policer Types on page 5431](#)
- [Policer Actions on page 5432](#)
- [Policer Colors on page 5432](#)
- [Filter-Specific Policers on page 5433](#)
- [Suggested Naming Convention for Policers on page 5433](#)
- [Policer Counters on page 5434](#)
- [Policer Algorithms on page 5434](#)
- [How Many Policers are Supported? on page 5434](#)
- [Policers can Limit Egress Firewall Filters on page 5434](#)

Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 171 on page 4664](#) illustrates this process.

Figure 202: Flow of Tricolor Marking Policer Operation



9017049

After you name and configure a policer, you use it by specifying it as an action in one or more firewall filters.

Policer Types

A switch supports three types of policers:

- Single-rate two-color marker—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A two-color policer is most useful for metering traffic at the port (physical interface) level.

- Single-rate three-color marker—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate three-color marker—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 365 on page 4666](#) for information about how metering results are applied for each of these policer types.

Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 365 on page 4666](#) describes the policer actions.

Table 453: Policer Actions

Policer	Marking	Implicit Action	Configurable Action
Single-rate two-color	Green (conforming)	Assign low loss priority	None
	Red (nonconforming)	None	Discard
Single-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the EBS)	Assign high loss priority	Discard
Two-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the PIR and PBS)	Assign high loss priority	Discard



NOTE: If you specify a policer in an egress firewall filter, the only supported action is discard.

Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 4658](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

Suggested Naming Convention for Policers

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named **srTCM1-ca**. The second two-rate, color-blind three-color configured would be named **trTCM2-cb**. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

Policer Counters

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

How Many Policers are Supported?

You can configure and commit the following numbers of policers on QFX3500 and QFX3600 standalone switches and QFabric Node devices:

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

Policers can Limit Egress Firewall Filters

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.

- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
 - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
 - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

Related Documentation

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 4669](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 4671](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 4670](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 4672](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 4753](#)

CoS Overview

- [Understanding Junos CoS Components on page 5436](#)
- [Understanding CoS Packet Flow on page 5440](#)
- [CoS Inputs and Outputs Overview on page 5442](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding Host Inbound Traffic Classification on page 5451](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5451](#)
- [Understanding CoS Code-Point Aliases on page 5453](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)
- [Understanding CoS Forwarding Classes on page 5469](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Understanding CoS Priority Group and Queue Guaranteed Rates \(Minimum Bandwidth\) on page 5497](#)
- [Understanding CoS Priority Group Shaping and Queue Shaping \(Maximum Bandwidth\) on page 5500](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5502](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)
- [Understanding CoS Buffer Configuration on page 5527](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding DCB Features and Requirements on page 5550](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5554](#)
- [Understanding DCBX on page 5564](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5573](#)

Understanding Junos CoS Components

This topic describes the Junos operating system (OS) class-of-service (CoS) components for the QFX Series:

- [Code-Point Aliases on page 5436](#)
- [Policers on page 5436](#)
- [Classifiers on page 5437](#)
- [Forwarding Classes on page 5437](#)
- [Forwarding Class Sets on page 5437](#)
- [Flow Control \(Ethernet PAUSE and Priority-Based Flow Control\) on page 5438](#)
- [Tail-Drop Profiles on page 5438](#)
- [Schedulers on page 5439](#)
- [Rewrite Rules on page 5439](#)

Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

Policers

Policers limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding

class, a different loss priority, or both. You define policers with filters that you can associate with input interfaces.

Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

You can create unicast classifiers for unicast traffic and multidestination classifiers for multicast, broadcast, and destination lookup fail traffic. You cannot assign unicast traffic and multidestination traffic to the same classifier.

You can apply unicast classifiers to one or more interfaces. Multidestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces.

Forwarding Classes

Forwarding classes group packets for transmission and CoS. You assign each packet to an output queue based on the packet's forwarding class. Forwarding classes affect the forwarding, scheduling, and rewrite marking policies applied to packets as they transit the switch.

The switch provides five default forwarding classes:

- fcoe—Fibre Channel over Ethernet traffic
- no-loss—Lossless traffic
- be—Best-effort traffic
- nc—Network control traffic
- mcast—Multicast traffic

The switch supports a total of 12 forwarding classes (8 unicast forwarding classes and 4 multicast forwarding classes), which provide flexibility in classifying traffic.

Forwarding Class Sets

You can group forwarding classes (output queues) into *forwarding class sets* in order to apply CoS to groups of traffic that require similar treatment. Forwarding class sets map traffic into priority groups to support enhanced transmission selection (ETS, described in IEEE 802.1Qaz).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. For example, you can configure different forwarding class sets to apply CoS to unicast groups of local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and configure another group for multicast traffic.

Within each forwarding class set, you can configure special CoS treatment for the traffic mapped to each individual queue. This provides the ability to configure CoS in a two-tier hierarchical manner. At the forwarding class set tier, you configure CoS for groups of traffic using a *traffic control profile*. At the queue tier, you configure CoS for individual output queues within a forwarding class set using a *scheduler* that you map to a queue (forwarding class) using a *scheduler map*.

Flow Control (Ethernet PAUSE and Priority-Based Flow Control)

Ethernet PAUSE (described in IEEE 802.3X) is a link-level flow control mechanism. During periods of network congestion, Ethernet PAUSE stops all traffic on a full-duplex Ethernet link for a period of time specified in the PAUSE message.

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is part of the IEEE data center bridging (DCB) specifications for creating a lossless Ethernet environment to transport loss-sensitive flows such as Fibre Channel over Ethernet (FCoE) traffic.

PFC is a link-level flow control mechanism similar to Ethernet PAUSE. However, Ethernet PAUSE stops all traffic on a link for a period of time. PFC decouples the pause function from the physical link and divides the traffic on the link into eight priorities (3-bit IEEE 802.1p code points). You can think of the eight priorities as eight “lanes” of traffic. You can apply pause selectively to the traffic on any priority without pausing the traffic on other priorities on the same link.

The granularity that PFC provides allows you to configure different levels of CoS for different types of traffic on the link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop methods of congestion management for IP traffic on the same link.



NOTE: If you transport FCoE traffic, you must enable PFC on the priority assigned to FCoE traffic (usually code point 011 on interfaces that carry FCoE traffic).

Tail-Drop Profiles

A tail-drop profile (drop profile) defines parameters that enable the network to drop packets during periods of congestion. A drop profile defines the conditions under which packets of different loss priorities drop, by determining the probability of dropping a packet for each loss priority when output queues become congested. Drop profiles essentially set a value for a level of queue fullness—when the queue fills to the level of the queue fullness value, packets drop.

You can associate different drop profiles with different loss priorities to set the probability of dropping packets. You can apply a drop profile for each loss priority to a forwarding

class (output queue) by applying a drop profile to a scheduler, and then mapping the scheduler to a forwarding class using a scheduler map. When the queue mapped to the forwarding class experiences congestion, the drop profile determines the level of packet drop for traffic of each loss priority in that queue.

Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. Typically you mark packets exceeding a particular service level with a high loss priority.

Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining the sequence in which different types of packets should be transmitted.

You can define the priority (**priority**), minimum bandwidth (**transmit-rate**), maximum bandwidth (**shaping-rate**), and tail-drop profiles to be applied to a particular queue for packet transmission. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream device to classify the packet into the appropriate service group. Rewriting (marking) outbound packets is useful when the switch is at the border of a network and must change the CoS values to meet the policies of the targeted peer.



NOTE: Ingress firewall filters can also rewrite forwarding class and loss priority values.

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding CoS Code-Point Aliases on page 5453](#)
- [Overview of Policers on page 4664](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Forwarding Classes on page 5469](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding DCB Features and Requirements on page 4965](#)

Understanding CoS Packet Flow

When a packet traverses a QFX Series Node device, the Node device provides the appropriate level of service to the packet using either default class of service (CoS) settings or CoS settings that you configure. On ingress ports, the Node device classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the Node device applies packet scheduling and (if you have configured them) rewrite rules to re-mark packets.

On the QFX Series, you can configure CoS on Layer 2 logical interfaces, and you can configure CoS on Layer 3 physical interfaces if you have defined at least one logical interface on the Layer 3 physical interface. You cannot configure CoS on Layer 2 physical interfaces and Layer 3 logical interfaces.

For Layer 2 traffic, either use the default CoS settings or configure CoS on each logical interface. You can apply different CoS settings to different Layer 2 logical interfaces.

For Layer 3 traffic, either use the default CoS settings or configure CoS on the physical interface (not on the logical unit). The QFX Series uses the CoS applied on the physical Layer 3 interface for all logical Layer 3 interfaces configured on the physical Layer 3 interface.

The QFX Series applies to CoS to packets as they flow through the system:

- An interface has one or more classifiers of different types applied to it (configure this at the **[edit class-of-service interfaces]** hierarchy level). The classifier types are based on the portion of the incoming packet that the classifier examines (IEEE 802.1p code point bits or DSCP code point bits).
- When a packet enters an ingress port, the classifier assigns the packet to a forwarding class and a loss priority based on the code point bits of the packet (configure this at the **[edit class-of-service classifiers]** hierarchy level).
- The QFX Series assigns each forwarding class to an output queue (configure this at the **[edit class-of-service forwarding-classes]** hierarchy level).
- Input (and output) policers meter traffic and can change the forwarding class and loss priority if a traffic flow exceeds its service level.
- A scheduler map is applied to each interface. When a packet exits an egress port, the scheduler map controls how it is treated (configure this at the **[edit class-of-service interfaces]** hierarchy level). A scheduler map assigns schedulers to forwarding classes (configure this at the **[edit class-of-service scheduler-maps]** hierarchy level).
- A scheduler defines how traffic is treated at the egress interface output queue (configure this at the **[edit class-of-service schedulers]** hierarchy level). You control the transmit rate, shaping rate, priority, and drop profile of each forwarding class by mapping schedulers to forwarding classes in scheduler maps, then applying scheduler maps to interfaces.

- A drop-profile defines how aggressively to drop packets that are mapped to a particular scheduler (configure this at the **[edit class-of-service drop-profiles]** hierarchy level).
- A rewrite rule takes effect as the packet leaves an interface that has a rewrite rule configured (configure this at the **[edit class-of-service rewrite-rules]** hierarchy level). The rewrite rule writes information to the packet (for example, a rewrite rule can re-mark the code point bits of outgoing traffic) according to the forwarding class and loss priority of the packet.

Figure 203 on page 5441 is a high-level flow diagram of how packets from various sources enter QFX Series interfaces, are classified at the ingress, and then scheduled (provided bandwidth) at the egress queues.

Figure 203: CoS Classifier, Queues, and Scheduler





Figure 204 on page 5441 shows the packet flow through the CoS components that you can configure.

Figure 204: Packet Flow Through Configurable CoS Components



The middle box ("Forwarding Class and Loss Priority") represents two values that you can use on ingress and egress interfaces. The system uses these values for classifying traffic on ingress interfaces and for rewrite rule re-marking on egress interfaces. Each outer box represents a process component. The components in the top row apply to incoming packets. The components in the bottom row apply to outgoing packets.

The solid-line arrows show the direction of packet flow from ingress to egress. The dotted-line arrows show inputs and outputs or show settings and actions based on those settings.

For example, the BA classifier sets the forwarding class and loss priority of incoming packets, so the forwarding class and loss priority are outputs of the classifier and the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packets based on those settings, so the arrow points toward the scheduler.

Related Documentation

- [Understanding CoS Classifiers on page 5455](#)
- [Overview of Policers on page 4664](#)
- [Understanding CoS Forwarding Classes on page 5469](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)
- [Understanding CoS Rewrite Rules on page 5549](#)

CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 454 on page 5442](#).

Table 454: CoS Mappings—Inputs and Outputs

CoS Mappings	Inputs	Outputs	Comments
classifiers	code-points	forwarding-class , loss-priority	The map sets the forwarding class and packet loss priority (PLP) for a specific set of code points.
drop-profile-map	loss-priority , protocol	drop-profile	The map sets the drop profile for a specific PLP and protocol type.
rewrite-rules	loss-priority , forwarding-class	code-points	The map sets the code points for a specific forwarding class and PLP.
rewrite-value (Fibre Channel Interfaces)	forwarding-class	code-point	The map sets the code point for the forwarding class specified in the fixed classifier attached to the native Fibre Channel (NP_Port) interface.

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Defining CoS Rewrite Rules on page 5805](#)

Understanding Default CoS Settings

If you do not configure CoS settings on the QFX Series, Junos OS performs some CoS functions to ensure that traffic and protocol packets are forwarded with minimum delay when the network experiences congestion. Some default mappings are automatically applied to each logical interface that you configure.

You can display default CoS settings by issuing the **show class-of-service** operational mode command.

This topic describes the default configurations for the following CoS components:

- [Default Forwarding Classes and Queue Mapping on page 5443](#)
- [Default Forwarding Class Sets \(Priority Groups\) on page 5444](#)
- [Default Code-Point Aliases on page 5444](#)
- [Default Classifiers on page 5446](#)
- [Default Rewrite Rules on page 5449](#)
- [Default Drop Profile on page 5449](#)
- [Default Schedulers on page 5449](#)
- [Default Scheduler Maps on page 5449](#)
- [Default Shared Buffer Configuration on page 5450](#)

Default Forwarding Classes and Queue Mapping

Table 455 on page 5443 shows the default mapping of the default forwarding classes to queues and packet drop attribute.

Table 455: Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
best-effort (be)	Best-effort traffic class (priority 0, IEEE 802.1p code point 000)	0	drop
fcoe	Guaranteed delivery for FCoE traffic (priority 3, IEEE 802.1p code point 011)	3	no-loss
no-loss	Guaranteed delivery for TCP no-loss traffic (priority 4, IEEE 802.1p code point 100)	4	no-loss

Table 455: Default Forwarding Classes and Queue Mapping (*continued*)

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
network-control (nc)	Network control traffic (priority 7, IEEE 802.1p code point 111)	7	drop
mcast	Multidestination traffic	8	drop

NOTE: You cannot configure multidestination forwarding classes as no-loss (lossless) traffic classes.

Default Forwarding Class Sets (Priority Groups)

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set.

Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth.

The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

Default Code-Point Aliases

Table 456 on page 5444 shows the default mapping of code-point aliases to IEEE code points.

Table 456: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1	110

Table 456: Default IEEE 802.1 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
nc2	111

[Table 457 on page 5445](#) shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

Table 457: Default DSCP and DCSP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000

Table 457: Default DSCP and DCSP IPv6 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
nc1	110000
nc2	111000

Default Classifiers

The QFX Series applies default unicast IEEE 802.1, unicast DSCP, and multdestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

[Table 458 on page 5446](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 458: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

[Table 459 on page 5447](#) shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

Table 459: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low
110	best-effort	low
111	best-effort	low

[Table 460 on page 5447](#) shows the default mapping of IEEE 802.1 code-point values to multdestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

Table 460: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

[Table 461 on page 5448](#) shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

Table 461: Default DSCP IP and IPv6 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low
af11 (001010)	best-effort	low
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low



NOTE: There are no default DSCP IP or IPv6 classifiers for multidestination traffic. DSCP IPv6 classifiers are not supported for multidestination traffic.

Default Rewrite Rules

There are no default rewrite rules. If you do not explicitly configure rewrite rules, the switch does not reclassify egress traffic.

Default Drop Profile

Table 462 on page 5449 shows the default drop profile configuration.

Table 462: Default Drop Profile

Fill Level	Drop Probability
100	100

Default Schedulers

Table 463 on page 5449 shows the default scheduler configuration.

Table 463: Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	low	5%
FCoE scheduler (queue 3)	35%	None	35%	low	35%
No-loss scheduler (queue 4)	35%	None	35%	low	35%
Network-control scheduler (queue 7)	5%	None	5%	low	5%
Multidestination scheduler (queue 8)	20%	None	20%	low	20%



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

Default Scheduler Maps

Table 464 on page 5449 shows the default mapping of forwarding classes to schedulers.

Table 464: Default Scheduler Maps

Forwarding Class	Scheduler
best-effort	Default BE scheduler

Table 464: Default Scheduler Maps (*continued*)

Forwarding Class	Scheduler
fcoe	Default FCoE scheduler
no-loss	No-loss scheduler
network-control	Default network-control scheduler
mcast-be	Default multidestination scheduler

Default Shared Buffer Configuration

Table [Table 465 on page 5450](#) and [Table 466 on page 5450](#) show the default shared buffer allocations:

Table 465: Default Ingress Shared Buffer Configuration

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

Table 466: Default Egress Shared Buffer Configuration

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

Related Documentation

- [Overview of Junos OS CoS for the QFX Series on page 5412](#)
- [Understanding Junos CoS Components on page 5436](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)
- [Understanding CoS Code-Point Aliases on page 5453](#)
- [Understanding CoS Forwarding Classes on page 5469](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

Understanding Host Inbound Traffic Classification

The destination address of traffic that enters the switch can be an external device such as another switch, a router, or a server, or the destination can be the host (the switch Routing Engine or CPU). When the destination is an external device, the DSCP and IEEE 802.1p code-point bits of incoming traffic are preserved as the traffic travels through the switch to the egress port. At the egress port, the code-point bits are either preserved when the packets are sent to the next hop or they are rewritten according to the rewrite rule attached to the egress interface.

When the destination of incoming traffic is the host, DSCP bits are preserved. However, IEEE 802.1p bits are not preserved. The IEEE 802.1p bits of traffic destined for the host are set to zero (0). This does not affect system behavior because the switch prioritizes traffic destined for the host based on the protocol type. For example, the switch gives a higher priority to BPDU traffic than to ping traffic.

Understanding Host Routing Engine Outbound Traffic Queues and Defaults

The host Routing Engine and CPU generate outbound traffic that is transmitted using different protocols. You cannot configure a classifier to map different types of outbound traffic that the host generates to forwarding classes (queues). The traffic that the host generates is assigned to forwarding classes by default as shown in [Table 467 on page 5452](#).

If you want to separate host outbound traffic from other traffic or if you want to assign that traffic to a particular queue, you can configure a single forwarding class for all traffic that the host generates. If you configure a forwarding class for outbound host traffic, that forwarding class is used globally for all traffic generated by the host. (That is, the host outbound traffic is mapped to the selected queue on all egress interfaces.) Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.

Whether you use the default host outbound traffic forwarding class configuration or configure a forwarding class for all host outbound traffic, the configuration applies to all Layer 2 and Layer 3 protocols and to all application-level traffic such as FTP and ping operations.

If you configure a queue for host outbound traffic, the queue must be properly configured on all interfaces.



NOTE: Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (`fabric`) and the egress Node device.

By default, traffic generated by the host is sent to the best effort queue (queue 0) or to the network control queue (queue 7). [Table 467 on page 5452](#) lists the default host traffic to output queue mapping.

Table 467: Routing Engine Protocol Default Queue Mapping

Routing Engine Protocol	Default Queue Mapping
Address Resolution Protocol (ARP) reply	Queue 0
ARP request	Queue 0
Border Gateway Protocol (BGP)	Queue 0
BGP TCP Retransmission	Queue 7
Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP)	Queue 3
File Transfer Protocol (FTP)	Queue 0
Internet Control Message Protocol (ICMP) reply	Queue 0
ICMP request	Queue 0
Internet Group Management Protocol (IGMP) query	Queue 7
IGMP report	Queue 0
Link Aggregation Control Protocol (LACP)	Queue 7
Open Shortest Path First (OSPF) hello	Queue 7
OSPF protocol data unit (PDU)	Queue 7
OSPF link state advertisements (LSAs)	Queue 7
Protocol Independent Multicast (PIM)	Queue 7
PIM hello	Queue 7
Simple Network Management Protocol (SNMP)	Queue 0
Secure Shell (SSH)	Queue 0
Telnet	Queue 0
Virtual Router Redundancy Protocol (VRRP)	Queue 7
VLAN Spanning Tree Protocol (VSTP)	Queue 7
xnm-clear-text	Queue 0
xnm-ssl	Queue 0

- Related Documentation**
- [Understanding CoS Forwarding Classes on page 5469](#)
 - [Changing the Host Outbound Traffic Default Queue Mapping on page 5793](#)
 - [Example: Configuring Forwarding Classes on page 5668](#)

Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs) and IEEE 802.1 bits to associate incoming packets with a particular CoS servicing level. You can assign a meaningful name or alias to the CoS values and use that alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure alias names for user-defined classifiers. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp or dscp-ipv6—Handles incoming IP and IPv6 packets.
- ieee-802.1—Handles Layer 2 CoS.

This topic covers:

- [Default Code-Point Aliases on page 5453](#)

Default Code-Point Aliases

[Table 468 on page 5453](#) shows the default mapping of code-point aliases to IEEE code points.

Table 468: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101

Table 468: Default IEEE 802.1 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
nc1	110
nc2	111

Table 469 on page 5454 shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

Table 469: Default DSCP and DSCP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000

Table 469: Default DSCP and DSCP IPv6 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
cs5	101000
nc1	110000
nc2	111000

- Related Documentation**
- [Understanding Junos CoS Components on page 5436](#)
 - [Defining CoS Code-Point Aliases on page 5783](#)

Understanding CoS Classifiers

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and a loss priority, and assign packets to output queues based on the associated forwarding class. There are three general types of classifiers:

- Behavior aggregate (BA) classifiers—DSCP and DSCP IPv6 classify IP and IPv6 traffic, and IEEE 802.1p classifiers classify all other traffic.
 - Fixed classifiers—Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
 - Multifield (MF) classifiers—MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.
- [Interfaces and Output Queues on page 5455](#)
 - [Behavior Aggregate Classifiers on page 5456](#)
 - [Fixed Classifiers on Ethernet Interfaces on page 5459](#)
 - [Fixed Classifiers on Native Fibre Channel Interfaces \(NP_Ports\) on page 5459](#)
 - [Multifield Classifiers on page 5459](#)
 - [Packet Classification for Routed VLAN Interfaces \(RVIs\) on page 5460](#)

Interfaces and Output Queues

On Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and link aggregation (LAG) interfaces, you can apply classifiers to Layer 2 logical interfaces and to Layer 3 physical interfaces if the Layer 3 physical interface has at least one defined logical interface. Classifiers applied to Layer 3 physical interfaces are used on all logical interfaces on that physical interface. “[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)” on page 5460 describes the interaction between classifiers and interfaces in greater detail.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification is performed. If the

two classification results conflict, the MF classification result overrides the BA classification result.

You cannot configure a fixed classifier and a BA classifier on the same interface.

You can configure both a DSCP or a DSCP IPv6 classifier and an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. All other traffic uses the IEEE classifier. You can configure only one DSCP classifier on a physical interface (either one DSCP classifier or one DSCP IPv6 classifier, but not both).

You can create unicast BA classifiers for unicast traffic and multicast BA classifiers for multdestination traffic, which includes multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot assign unicast traffic and multdestination traffic to the same BA classifier.

On each interface, the switch has separate output queues for unicast traffic and for multdestination traffic:

- The switch supports 12 output queues, with 8 queues dedicated to unicast traffic and 4 queues dedicated to multdestination traffic.
- Queues 0 through 7 are unicast traffic queues. You can apply only unicast BA classifiers to unicast queues. A unicast BA classifier should contain only forwarding classes that are mapped to unicast queues.
- Queues 8 through 11 are multdestination traffic queues. You can apply only multdestination BA classifiers to multdestination queues. A multdestination BA classifier should contain only forwarding classes that are mapped to multdestination queues.

You can apply unicast classifiers to one or more interfaces. Multdestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Use the DSCP multdestination classifier for both IP and IPv6 multdestination traffic. The DSCP IPv6 classifier is not supported for multdestination traffic.

Behavior Aggregate Classifiers

The behavior aggregate classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

The switch supports two types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ (IP and IPv6)
- IEEE 802.1p CoS bits

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore, core devices, which handle high traffic volumes, are normally configured to perform BA classification.

Unicast and multicast traffic cannot share the same classifier. You can map unicast traffic and multicast traffic to the same classifier CoS value, but the unicast traffic must

belong to a unicast classifier and the multicast traffic must belong to a multidestination classifier.

Default Behavior Aggregate Classification

Juniper Networks Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. [Table 470 on page 5457](#) lists different types of interfaces and the corresponding implicit default BA classifiers.

Table 470: Default BA Classification

Type of Interface	Default BA Classification
Layer 2 interface in trunk mode or in tagged-access mode	ieee8021p-default
Layer 3 interface	dscp-default
Layer 2 interface in access mode	ieee8021p-untrusted



NOTE: There are default BA classifiers for the best-effort, fcoe, no-loss, network-control, and mcast forwarding classes.

When you explicitly associate a unicast classifier with a logical interface, you are overriding the default unicast classifier with the explicit unicast classifier.



NOTE: You can apply only one classifier of each type, DSCP and IEEE 802.1p, to a Layer 2 interface. If both types of classifiers are present, DSCP classifiers take precedence over IEEE 802.1p classifiers.

Importing a Classifier

You can use any existing classifier, including the default classifiers, as the basis for defining a new classifier. You accomplish this using the **import** statement.

The imported classifier is used as a template and is not modified. The modifications you make become part of a new classifier (and a new template) identified by the name of the new classifier. Whenever you commit a configuration that assigns a new class-name and loss-priority value to a code-point alias or set of bits, it replaces that entry in the new classifier template. As a result, you must explicitly specify every CoS value in every designation that requires modification.

Multidestination Classifiers

Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. You can configure both a DSCP multidestination classifier and an IEEE multidestination classifier. IP and IPv6 traffic use the DSCP classifier, and all other traffic uses the IEEE classifier.

DSCP IPv6 multdestination classifiers are not supported, so IPv6 traffic uses the DSCP multdestination classifier.

The default multdestination classifier is the IEEE 802.1p multdestination classifier.

PFC Priorities

The eight IEEE 802.1p code points correspond to the eight priorities that priority-based flow control (PFC) uses to differentiate traffic classes for lossless transport. When you map a forwarding class (which maps to an output queue) to an IEEE 802.1p CoS value, the IEEE 802.1p CoS value identifies the priority.

Although you can map a priority to any output queue (by mapping the priority to a forwarding class), we recommend that the priority and the unicast forwarding class match in a one-to-one correspondence in which priority 0 is assigned to queue 0, priority 1 is assigned to queue 1, and so on, as shown in [Table 471 on page 5458](#). A one-to-one correspondence of queue and priority numbers makes it easier to configure and maintain the mapping of forwarding classes to priorities and queues.

Table 471: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping

IEEE 802.1p Code Point	PFC Priority	Unicast Output Queue	Forwarding Class and Packet Drop Attribute
000	0	0	best-effort (drop)
001	1	1	best-effort (drop)
010	2	2	best-effort (drop)
011	3	3	fcoe (no-loss)
100	4	4	no-loss (no-loss)
101	5	5	best-effort (drop)
110	6	6	network-control (drop)
111	7	7	network-control (drop)



NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires. We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.

Fixed Classifiers on Ethernet Interfaces

Fixed classifiers map all traffic on an interface to a forwarding class and a loss priority. (As opposed to BA classifiers, which map traffic into multiple different forwarding classes based on the CoS field value in the packet header.) The forwarding class determines the output queue. Incoming traffic of all IEEE 802.1p priorities is classified into the forwarding class specified in the fixed classifier. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

You cannot configure a fixed classifier and a BA classifier on the same interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.

To switch from a fixed classifier to a BA classifier or to switch from a BA classifier to a fixed classifier, deactivate the existing classifier attachment on the interface, and then attach the new classifier to the interface.



NOTE: If you configure a fixed classifier that classifies all incoming traffic into the `fcoe` forwarding class (or any forwarding class designed to handle FCoE traffic), you must ensure that all traffic that enters the interface is FCoE traffic and is tagged with the FCoE IEEE 802.1p code point (priority).

Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports)

Applying a fixed classifier to a native Fibre Channel (FC) interface (NP_Port) is a special case. By default, native FC interfaces classify incoming traffic from the FC SAN into the `fcoe` forwarding class and map the traffic to IEEE 802.1p priority 3 (code point 011). When you apply a fixed classifier to an FC interface, you also configure a priority rewrite value for the interface. The FC interface uses the priority rewrite value as the IEEE 802.1p tag value for all incoming packets instead of the default value of 3.

For example, if you specify a priority rewrite value of 5 (code point 101) for an FC interface, the interface tags all incoming traffic from the FC SAN with priority 5 and classifies the traffic into the forwarding class specified in the fixed classifier.



NOTE: The forwarding class specified in the fixed classifier on FC interfaces must be a lossless forwarding class.

Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) support in end-user applications. On a switch at the edge of a network, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

Packet Classification for Routed VLAN Interfaces (RVIs)

You cannot apply classifiers directly to routed VLAN interfaces (RVIs) because the members of RVIs are VLANs, not ports. However, you can apply classifiers to the VLAN port members of an RVI. You can also apply MF classifiers to RVIs.

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p or DSCP class of service (CoS) code point bits in the packet header. At egress interfaces, you can use rewrite rules to change (re-mark) the code point bits before the interface forwards the packets. At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p or DSCP CoS code point bits in the packet header. At egress interfaces, rewrite rules can change (re-mark) the code point bits before the interface forwards the packets.

You can apply classifiers and rewrite rules to interfaces to control the level of CoS applied to each packet as it traverses the system and the network. This topic describes:

- [Supported Classifier and Rewrite Rule Types on page 5461](#)
- [Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration on page 5462](#)
- [Default Classifiers on page 5463](#)
- [Default Rewrite Rules on page 5464](#)
- [Classifier Precedence on page 5464](#)
- [Classifier Behavior and Limitations on page 5465](#)
- [Rewrite Rule Precedence and Behavior on page 5466](#)
- [Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration on page 5466](#)

Supported Classifier and Rewrite Rule Types

Table 472 on page 5461 shows the types of classifiers and rewrite rules that the QFX Series supports:

Table 472: Supported Classifiers and Rewrite Rules

Classifier or Rewrite Rule Type	Description
Fixed classifier	Classifies all ingress traffic on a physical interface into one fixed forwarding class, regardless of the CoS bits in the packet header.
DSCP and DSCP IPv6 unicast classifiers	Classifies IP and IPv6 traffic into forwarding classes and assigns loss priorities to the traffic.
IEEE 802.1p unicast classifier	Classifies Ethernet traffic into forwarding classes and assigns loss priorities to the traffic.
DSCP multidestination classifier (also used for IPv6 multidestination traffic)	Classifies IP and IPv6 multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
IEEE 802.1p multidestination classifier	Classifies Ethernet multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
DSCP and DSCP IPv6 rewrite rules	Re-marks the DSCP code points of IP and IPv6 packets before forwarding the packets.
IEEE 802.1p rewrite rule	Re-marks the IEEE 802.1p code points of Ethernet packets before forwarding the packets.



NOTE: On native Fibre Channel (FC) interfaces (NP_Ports) only, you can specify a rewrite value to set the IEEE 802.1p code point of incoming FC traffic when the NP_Port encapsulates the FC packet in Ethernet before forwarding it to the FCoE network (see [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#)).

DSCP and IEEE classifiers are behavior aggregate (BA) classifiers. Multidestination classifiers apply to all interfaces; you cannot apply them to individual interfaces.

Classifying packets into forwarding classes assigns packets to the output queues associated with the forwarding classes. Classifying traffic into a forwarding class associates the CoS scheduling for the forwarding class with that traffic.



NOTE: In addition to BA classifiers and fixed classifiers, which classify traffic based on the CoS field in the packet header, you can use firewall filters to configure multifield (MF) classifiers. MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration

To apply a classifier to incoming traffic or a rewrite rule to outgoing traffic, you need to apply the classifier or rewrite rule to one or more interfaces. When you apply a classifier or rewrite rule to an interface, the interface uses the classifier to group incoming traffic into forwarding classes and uses the rewrite rule to re-mark the CoS code point value of each packet before it leaves the system.

Not all interfaces types support all types of CoS configuration. This section describes:

- [Interface Types That Support Classifier and Rewrite Rule Configuration on page 5462](#)
- [Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support on page 5462](#)
- [Routed VLAN Interfaces \(RVIs\) on page 5463](#)

Interface Types That Support Classifier and Rewrite Rule Configuration

You can apply classifiers to all Ethernet interfaces. For Layer 3 LAGs, configure BA or fixed classifiers on the LAG (ae) interface. The classifier configured on the LAG is valid on all of the LAG member interfaces.

You can apply fixed classifiers to native FC interfaces (NP_Ports). You cannot apply other types of classifiers or rewrite rules to native FC interfaces. You can rewrite the value of the IEEE 802.1p code point of incoming FC traffic when the interface encapsulates it in Ethernet before forwarding it to the FCoE network as described in [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#).

Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support

The QFX Series Ethernet ports can function as:

- Layer 2 physical interfaces (family ethernet-switching)
- Layer 2 logical interfaces (family ethernet-switching)
- Layer 3 physical interfaces (family inet/inet6)
- Layer 3 logical interfaces (family inet/inet6)

You can apply CoS classifiers and rewrite rules only to the following interfaces:

- Layer 2 logical interfaces
- Layer 3 physical interfaces if at least one logical Layer 3 interface is configured on the physical interface



NOTE: The CoS you configure on a Layer 3 physical interface is applied to all of the Layer 3 logical interfaces on that physical interface. This means that each Layer 3 interface uses the same classifiers and rewrite rules for all of the Layer 3 traffic on that interface.

You cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 473 on page 5463](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

Table 473: Ethernet Interface Support for Classifier and Rewrite Rule Configuration

CoS Classifiers and Rewrite Rules	Layer 2 Physical Interfaces	Layer 2 Logical Interfaces	Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined)	Layer 3 Logical Interfaces
Fixed classifier	No	Yes	Yes	No
DSCP classifier	No	Yes	Yes	No
DSCP IPv6 classifier	No	Yes	Yes	No
IEEE 802.1p classifier	No	Yes	Yes	No
DSCP rewrite rule	No	Yes	Yes	No
DSCP IPv6 rewrite rule	No	Yes	Yes	No
IEEE 802.1p rewrite rule	No	Yes	Yes	No



NOTE: IEEE 802.1p and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Routed VLAN Interfaces (RVIs)

You cannot apply classifiers and rewrite rules directly to routed VLAN interfaces (RVIs) because the members of RVIs are VLANs, not ports. However, you can apply classifiers and rewrite rules to the VLAN port members of an RVI. You can also apply MF classifiers to RVIs.

Default Classifiers

If you do not explicitly configure classifiers on an Ethernet interface, the QFX Series applies default classifiers (see [“Understanding Default CoS Settings” on page 5443](#)) so that the traffic receives basic CoS treatment. The factors that determine the default classifier applied to the interface include the interface type (Layer 2 or Layer 3), the port mode

(trunk, tagged-access, or access), and whether logical interfaces have been configured. The system applies a default classifier using the following rules:

- If the physical interface has at least one Layer 3 logical interface configured, it uses the default DSCP classifier.
- If the physical interface has a Layer 2 logical interface in trunk mode or tagged-access mode, it uses the default trusted classifier.
- If the physical interface has a Layer 2 logical interface in access mode, it uses the default untrusted classifier.
- If the physical interface has no logical interface configured, no default classifier is applied.
- The default multidestination classifier is the IEEE 802.1p multidestination classifier.

Default Rewrite Rules

No default rewrite rules are applied to interfaces. If you want to re-mark packets at the egress interface, you must explicitly configure a rewrite rule.

Classifier Precedence

You can apply multiple unicast classifiers (MF, fixed, IEEE 802.1p, or DSCP) to a physical or logical Ethernet interface to handle different types of traffic. When you apply more than one classifier to an interface, the system uses an order of precedence to determine which classifier to use on physical and logical interfaces:

- [Unicast Classifier Precedence on Physical Ethernet Interfaces on page 5464](#)
- [Unicast Classifier Precedence on Logical Ethernet Interfaces on page 5465](#)

Unicast Classifier Precedence on Physical Ethernet Interfaces

The precedence of unicast classifiers on physical interfaces, from the highest-priority classifier to the lowest-priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the physical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a physical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier.



NOTE: You cannot apply a fixed classifier and a DSCP or IEEE classifier to the same interface. If a DSCP classifier, an IEEE classifier, or both are on an interface, you cannot apply a fixed classifier to that interface unless you first delete the DSCP and IEEE classifiers. If a fixed classifier is on an interface, you cannot apply a DSCP classifier or an IEEE classifier unless you first delete the fixed classifier.

Unicast Classifier Precedence on Logical Ethernet Interfaces

The precedence of unicast classifiers on logical interfaces, from the highest priority classifier to the lowest priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the logical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a logical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier.

Classifier Behavior and Limitations

Consider the following behaviors and constraints when you apply classifiers to physical and logical Ethernet interfaces:

- You can configure only one DSCP classifier (IP or IPv6) on a physical interface. You cannot configure both types of DSCP classifier on one physical interface. Both IP and IPv6 traffic use whichever DSCP classifier is configured on the interface.
- When you configure a DSCP or a DSCP IPv6 classifier on a physical interface and the physical interface has at least one logical Layer 3 interface, all packets (IP, IPv6, and non-IP) use that classifier.
- An interface with both a DSCP classifier (IP or IPv6) and an IEEE 802.1p classifier uses the DSCP classifier for IP and IPv6 packets, and uses the IEEE classifier for all other packets.
- Fixed classifiers and BA classifiers (DSCP and IEEE classifiers) are not permitted simultaneously on an interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.
- When you configure an IEEE 802.1p classifier on a physical interface and a DSCP classifier is not explicitly configured on that interface, the interface uses the IEEE classifier for all types of packets. No default DSCP classifier is applied to the interface. (In this case, if you want a DSCP classifier on the interface, you must explicitly configure it.)

- The system does not apply a default classifier to a physical interface until you create a logical interface on that physical interface. If you configure a Layer 3 logical interface, the system uses the default DSCP classifier. If you configure a Layer 2 logical interface, the system uses the default IEEE 802.1p trusted classifier if the port is in trunk mode or tagged-access mode, or the default IEEE 802.1p untrusted classifier if the port is in access mode.
- MF classifiers configured on logical interfaces take precedence over BA and fixed classifiers. (Use firewall filters to configure MF classifiers.) When BA or fixed classifiers are present on an interface, you can still configure an MF classifier on that interface.

Rewrite Rule Precedence and Behavior

The following rules apply on both physical and logical Ethernet interfaces for rewrite rules:

- If you configure both one DSCP (or DSCP IPv6) rewrite rule and one IEEE 802.1p rewrite rule on an interface, both rewrite rules take effect. Traffic with IP and IPv6 headers use the DSCP rewrite rule, and traffic with a VLAN tag uses the IEEE rewrite rule.
- If you do not explicitly configure a rewrite rule, there is no default rewrite rule, so the system does not apply any rewrite rule to the interface.
- You can apply a DSCP rewrite rule or a DSCP IPv6 rewrite rule to an interface, but you cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6.

Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration

You can apply classifiers and rewrite rules only on Layer 2 logical interfaces and Layer 3 physical interfaces (if the Layer 3 physical interface has at least one defined logical interface). This section focuses on BA classifiers, but the interaction between BA classifiers and interfaces described in this section also applies to fixed classifiers and rewrite rules.

There are two components to applying classifiers or rewrite rules to interfaces:

1. Setting the interface family (inet, inet6, or ethernet-switching; ethernet-switching is the default interface family) in the **[edit interfaces]** configuration hierarchy.
2. Applying a classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy.

These are separate operations that can be set and committed at different times. Because the type of classifier or rewrite rule you can apply to an interface depends on the interface family configuration, the system performs checks to ensure that the configuration is valid. The method the system uses to notify you of an invalid configuration depends on the **set** operation that causes the invalid configuration.

When applying the classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy causes an invalid configuration, the system rejects the configuration and returns a commit check error.

When setting the interface family in the **[edit interfaces]** configuration hierarchy causes an invalid configuration, the system creates a syslog error message. When you receive the error message, you need to remove the classifier or rewrite rule configuration from the logical interface and apply it to the physical interface, or remove the classifier or rewrite rule configuration from the physical interface and apply it to the logical interface. For classifiers, if you do not take action to correct the error, the system programs the default classifier for the interface family on the interface. (There are no default rewrite rules. If the commit check fails, no rewrite rule is applied to the interface.)

Two scenarios illustrate these situations:

- [Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error on page 5467](#)
- [Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error on page 5468](#)



NOTE: Both of these scenarios also apply to fixed classifiers and rewrite rules.

Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error

In Scenario 1, we set the interface family, and then specify an invalid classifier.

1. Set and commit the interface as a Layer 3 (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This commit operation succeeds.

2. Set and commit a DSCP classifier on the logical interface (this example uses a DSCP classifier named **dscp1**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This configuration is not valid, because it attempts to apply a classifier to a Layer 3 logical interface. Because the failure is caused by the class-of-service configuration and not by the interface configuration, the system rejects the commit operation and issues a commit error, not a syslog message.

Note that the commit operation succeeds if you apply the classifier to the physical Layer 3 interface as follows:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Because the logical unit is not specified, the classifier is applied to the physical Layer 3 interface in a valid configuration, and the commit check succeeds.

Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error

In Scenario 2, we set the classifier first, then set an invalid interface type.

1. Set and commit a DSCP classifier on a Layer 3 logical interface, assuming that the interface has no existing configuration:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This commit succeeds. Because no explicit configuration existed on the interface, it is by default a Layer 2 (**family ethernet-switching**) interface. Layer 2 logical interfaces support BA classifiers, so applying the classifier is a valid configuration.

2. Set and commit the interface as a Layer 3 interface (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This configuration is not valid because it attempts to change an interface from Layer 2 (**family ethernet-switching**) to Layer 3 (**family inet**) when a classifier has already been applied to a logical interface. Layer 3 logical interfaces do not support classifiers. Because the failure is caused by the interface configuration and not by the class-of-service configuration, the system does not issue a commit error, but instead issues a syslog message.

When the system issues the syslog message, it programs the default classifier for the interface type on the interface. In this scenario, the interface has been configured as a Layer 3 interface, so the system applies the default DSCP profile to the physical Layer 3 interface.

In this scenario, to install a configured DSCP classifier, you remove the misconfigured classifier from the Layer 3 logical interface and apply it to the Layer 3 physical interface. For example:

```
[edit]
user@switch# delete class-of-service interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
user@switch# set class-of-service interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)

- [Defining CoS Rewrite Rules on page 5805](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP_Ports\) on page 5801](#)

Understanding CoS Forwarding Classes

Forwarding classes group traffic and assign the traffic to output queues. Each forwarding class is mapped to an output queue. Classification identifies the output queue for each incoming packet by mapping the packet code point bits to forwarding classes. The forwarding class to queue mapping defines the output queue used for the packet.

A classifier must associate each packet with one of the following five default forwarding classes or with a user-configured forwarding class in order to assign an output queue to the packet:

- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.
- **no-loss**—Guaranteed delivery for TCP lossless traffic.
- **best-effort**—Provides best-effort delivery without a service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- **network-control**—Supports protocol control and is typically high priority.
- **mcast**—Provides no service profile for multdestination (multicast, broadcast, and destination lookup fail) packets.

The switch supports up to 12 forwarding classes, thus enabling flexible, differentiated, packet classification. For example, you can configure multiple classes of best-effort traffic such as **best-effort**, **best-effort1**, and **best-effort2**.

The switch supports up to 12 output queues: 8 output queues for unicast traffic (queues 0 through 7) and 4 output queues for multdestination traffic (queues 8 through 11). Forwarding classes mapped to unicast queues are associated with unicast traffic, and forwarding classes mapped to multdestination queues are associated with multdestination traffic. You cannot map unicast and multdestination traffic to the same queue. You cannot map a strict-high priority queue to a multdestination forwarding class (queues 8 through 11 do not support strict-high priority configuration).

- [Default Forwarding Classes on page 5469](#)
- [Forwarding Class Configuration Rules on page 5471](#)
- [Lossless Transport Support on page 5472](#)

Default Forwarding Classes

[Table 474 on page 5470](#) shows the four default forwarding classes defined for unicast traffic, and [Table 475 on page 5471](#) shows the four default forwarding classes defined for multicast traffic.

If desired, you can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. CoS configurations can be

quite complicated, so unless it is required by your scenario, we recommend that you not alter the default class names or queue number associations.

Table 474: Default Forwarding Classes for Unicast Packets

Forwarding Class Name	Default Queue Mapping	Comments
best-effort (be)	0	<p>The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>
fcoe	3	<p>The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end to end for packets in this service class. The software accepts excess traffic in this class, but in contrast to the assured forwarding class, the out-of-profile expedited-forwarding class packets can be forwarded out of sequence or dropped.</p> <p>NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires.</p> <p>We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.</p> <p>By default, this is a lossless forwarding class with a packet drop attribute of no-loss.</p>
no-loss	4	<p>The software offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.</p> <p>The software accepts excess traffic, but it applies a tail-drop profile to determine if the excess packets are dropped and not forwarded.</p> <p>Up to two drop probabilities (low and high) are defined for this service class.</p> <p>By default, this is a lossless forwarding class with a packet drop attribute of no-loss.</p>
network-control (nc)	7	<p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>

Table 475: Default Forwarding Classes for Multicast Packets

Forwarding Class Name	Default Queue Mapping	Comments
mcast	8	<p>The software does not apply any special CoS handling to the multidestination packets. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>



NOTE: Mirrored traffic is always sent to the queue that corresponds to the multidestination forwarding class. The switched copy of the mirrored traffic is forwarded with the priority determined by the behavior aggregate classification process.

Forwarding Class Configuration Rules

Take the following rules into account when you configure forwarding classes:

- [Queue Assignment Rules on page 5471](#)
- [Scheduling Rules on page 5471](#)
- [Rewrite Rules on page 5472](#)

Queue Assignment Rules

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit operation fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that appears in the default configuration is the forwarding class currently associated with that queue.
- Only unicast forwarding classes can be mapped to unicast queues (0 through 7), and only multidestination forwarding classes can be mapped to multidestination queues (8 through 11).
- Strict-high priority queues cannot be mapped to multidestination forwarding classes. (Strict-high priority traffic cannot be mapped to queues 8 through 11).
- If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

Scheduling Rules

When you define a forwarding class that is used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you

must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

Rewrite Rules

On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

Lossless Transport Support

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the QFX3500 or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.
- Changing any portion of a PFC configuration on a port blocks the entire port until the change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Changing the PFC configuration means any change to a congestion notification profile that is configured on a port (enabling or disabling PFC on a code point, changing the MRU or cable-length value, or specifying an output flow control queue). Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.



NOTE: Junos OS Release 12.2 introduces changes to the way the QFX Series handles lossless forwarding classes (the `fcoe` and `no-loss` forwarding classes).

In Junos OS Release 12.1, both explicitly configuring the `fcoe` and `no-loss` forwarding classes, and using the default configuration for these forwarding classes, resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the `fcoe` or the `no-loss` forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the `fcoe` or the `no-loss` forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the explicit `fcoe` and `no-loss` forwarding class configuration before you upgrade to Junos OS Release 12.2.

See “[Overview of CoS Changes Introduced in Junos OS Release 12.2](#)” on [page 5428](#) for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the `fcoe` and `no-loss` forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new `no-loss` packet drop attribute or the forwarding class is lossy.

Related Documentation

- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding Junos CoS Components on page 5436](#)
- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Defining CoS Forwarding Classes on page 5787](#)

Understanding CoS Forwarding Class Sets (Priority Groups)

A forwarding class set is the Junos OS configuration construct that equates to a priority group in enhanced transmission selection (ETS, described in IEEE 802.1Qaz). The switch implements ETS using a two-tier hierarchical scheduler.

A priority group is a group of queues (priorities). Mapping a forwarding class to a queue defines the traffic for that queue, so a priority equates to a queue (forwarding class). The queues in a priority group share the port bandwidth allocated to that priority group. The traffic for queues in one priority group usually share similar traffic-handling requirements.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. Only unicast forwarding classes can belong to unicast forwarding class sets. Only multicast forwarding classes can belong to the multicast forwarding class set.

If you configure a strict-high priority queue, you must observe the following rules when configuring forwarding class sets:

- You must create a separate forwarding class set for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for a forwarding class set that includes a strict-high priority queue. (You also cannot configure a guaranteed minimum bandwidth for a strict-high queue.)

You must use hierarchical scheduling to define CoS for output queues. The two-tier hierarchical scheduler defines bandwidth resources for the priority group, and then allocates those resources among the priorities that belong to the priority group.

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set. Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.

When you explicitly configure forwarding class sets and map them to an interface, any forwarding class that you do not map to a forwarding class set receives no guaranteed bandwidth on that interface. Forwarding classes that belong to the default forwarding class set might receive bandwidth if the other forwarding class sets are not using all of the port bandwidth. However, the amount of bandwidth forwarding classes that are not in explicitly configured forwarding class sets receive is not guaranteed. The bandwidth

for the default forwarding class depends on whether extra port bandwidth is available and therefore is not deterministic.

To guarantee bandwidth for forwarding classes in a predictable manner, be sure to map all forwarding classes that you expect to carry traffic on an interface to a forwarding class set and map the forwarding class set to the interface.

Related Documentation

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)

Understanding Default CoS Scheduling and Classification

If you do not configure hierarchical scheduling on an interface, the switch uses the default classifiers for ingress traffic and the default schedulers for egress traffic. Default scheduling and classification handle all traffic types (best-effort, FCoE, no-loss, network-control, and multdestination traffic).

Hierarchical scheduling groups egress queues (priorities, configured as forwarding classes) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.



NOTE: If you explicitly configure one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group).

The following sections describe:

- [Default Classification on page 5475](#)
- [Default Scheduling on page 5478](#)
- [Default DCBX Advertisement on page 5479](#)
- [Default Scheduling and Classification Summary on page 5479](#)

Default Classification

The default classifiers assign unicast and multicast best-effort and network-control ingress traffic to forwarding classes and loss priorities. The QFX Series applies default unicast IEEE 802.1, unicast DSCP, and multdestination classifiers to each interface that

does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

[Table 476 on page 5476](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 476: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

[Table 477 on page 5476](#) shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

Table 477: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low

Table 477: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier) (continued)

Code Point	Forwarding Class	Loss Priority
110	best-effort	low
111	best-effort	low

[Table 478 on page 5477](#) shows the default mapping of IEEE 802.1 code-point values to multdestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

Table 478: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

[Table 479 on page 5477](#) shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

Table 479: Default DSCP IP and IPv6 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low
af11 (001010)	best-effort	low
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low

Table 479: Default DSCP IP and IPv6 Unicast Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low



NOTE: There are no default DSCP IP or IPv6 multdestination classifiers for multdestination traffic. DSCP IPv6 multdestination classifiers are not supported for multdestination traffic.

Default Scheduling

The default schedulers allocate egress bandwidth resources to unicast and multicast best-effort and network-control egress traffic as shown in [Table 480 on page 5478](#):

Table 480: Default Scheduler Configuration

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	low	5%

Table 480: Default Scheduler Configuration (*continued*)

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
FCoE scheduler (queue 3)	35%	None	35%	low	35%
No-loss scheduler (queue 4)	35%	None	35%	low	35%
Network-control scheduler (queue 7)	5%	None	5%	low	5%
Multidestination scheduler (queue 8)	20%	None	20%	low	20%



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

Default Scheduling and Classification Summary

If you do not configure hierarchical scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Example: Configuring Queue Schedulers on page 5674](#)

Understanding CoS Hierarchical Port Scheduling (ETS)

Scheduling defines the class-of-service (CoS) properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the queue priority, and the drop profiles associated with the queue.

Hierarchical port scheduling is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues. Hierarchical scheduling includes the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz).

The two tiers used in hierarchical scheduling are priorities and priority groups, as shown in [Table 481 on page 5481](#).

Table 481: Hierarchical Scheduling Tiers

Junos OS Configuration Construct	Equivalent ETS Construct	Description
Forwarding class	Priority	<p>Think about priorities (forwarding classes) as output queues. You map forwarding classes to queues, so each forwarding class is in essence an output queue.</p> <p>When you use a classifier to map a forwarding class to an IEEE 802.1p code point, the code point identifies that traffic's priority for priority-based flow control (PFC). Thus the forwarding class, the queue mapped to the forwarding class, and the priority mapped to the forwarding class all identify the same traffic.</p>
Forwarding class set	Priority group	<p>Priority groups (forwarding class sets) are groups of priorities. Forwarding class membership in a forwarding class set defines the priority group to which each priority belongs.</p> <p>You can configure up to three unicast priority groups and one multicast forwarding class set.</p>



NOTE: If you explicitly configure one or more priority groups on an interface, any priority that is not assigned to a priority group on that interface is assigned to an automatically generated default priority group and receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class that you want to forward traffic on that interface must belong to a forwarding class set.

This topic describes:

- [Hierarchical Scheduling and ETS on page 5482](#)
- [ETS Advertisement in DCBX on page 5483](#)
- [Hierarchical Scheduling Process on page 5483](#)

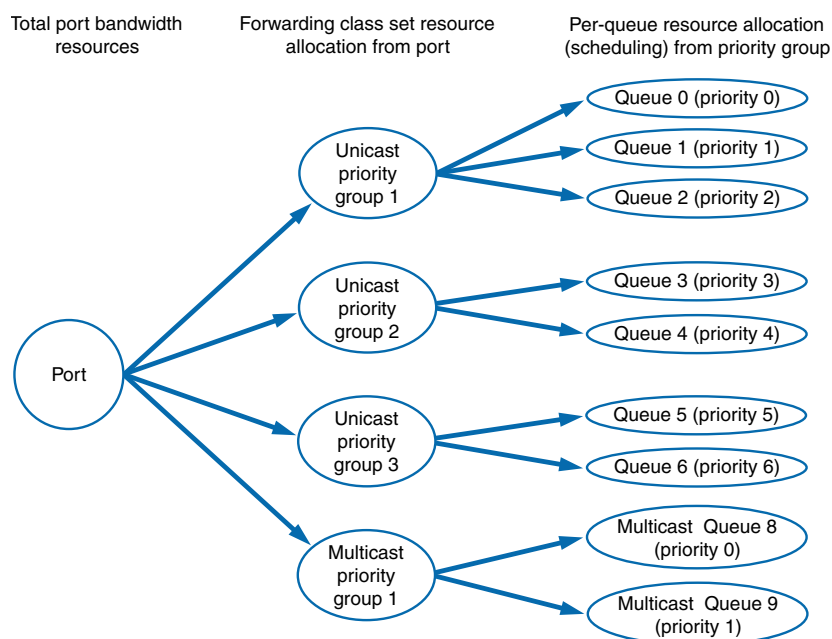
- [Strict-High Priority Queues and Hierarchical Scheduling on page 5484](#)
- [Default Hierarchical Scheduling on page 5485](#)

Hierarchical Scheduling and ETS

Two-tier hierarchical scheduling enables you to manage bandwidth efficiently by enabling you to define the CoS properties for each priority group and for each priority. One tier of the hierarchical scheduler allocates port bandwidth to a priority group. The other tier of the hierarchical scheduler determines the portion of the priority group bandwidth that a queue can use.

The CoS properties you configure for a priority group define the port bandwidth resources available to the queues in that priority group. The CoS properties you configure for each queue specify the portion or percentage of the total bandwidth configured for the priority group that is available to the queue. [Figure 205 on page 5482](#) shows the relationship of port resource allocation to priority groups and priority group resource allocation to queues (priorities).

Figure 205: Hierarchical Scheduling Tiers



g040722

If a queue is not using its allocated bandwidth, ETS shares the unused bandwidth among the other queues in the priority group. If link bandwidth is available or if a priority group on a link is not using its allocated bandwidth, ETS shares the unused bandwidth with other priority groups on the link. In this way ETS improves link bandwidth utilization and provides each queue with the maximum bandwidth. Priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.



NOTE: The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows.

ETS Advertisement in DCBX

When you configure hierarchical scheduling on a port, Data Center Bridging Capability Exchange Protocol (DCBX) advertises:

- Each priority group
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

When you configure hierarchical scheduling on a port, any priority that is not part of an explicitly configured priority group is assigned to the automatically generated default priority group and receives no bandwidth. The default priority group is transparent. It does not appear in the configuration.

Hierarchical Scheduling Process

Hierarchical scheduling consists of multiple configuration steps that create the priorities and the priority groups, schedule their resources, and assign them to interfaces. The steps below correspond to the six blocks in the packet flow diagram shown in [Figure 206 on page 5484](#):

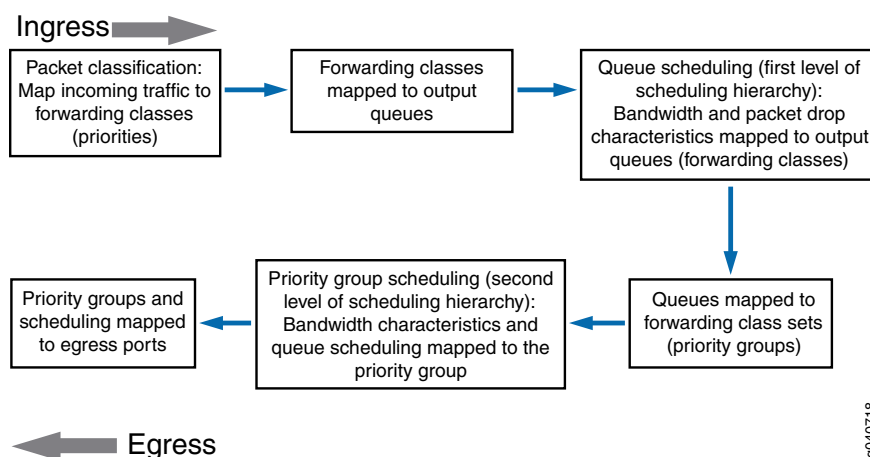
1. Packet classification:
 - Classify incoming traffic into priorities. This consists of either using the default classifiers or configuring classifiers to map IEEE 802.1p code points and loss priorities to the forwarding classes.
 - Apply the classifiers to ingress interfaces. This groups incoming traffic into forwarding classes (priorities) by mapping code points to forwarding classes and loss priorities on the specified interface.
2. Configure the output queues for the forwarding classes (priorities). This consists of either using the default forwarding classes and forwarding-class-to-queue mapping or creating your own forwarding classes and mapping them to queues.
3. Allocate resources to the forwarding classes:
 - Define resources for the priorities. This consists of configuring schedulers to set minimum guaranteed bandwidth, maximum bandwidth, drop profiles for Weighted Random Early Detection (WRED), and bandwidth priority to apply to a forwarding class. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.
 - Map resources to priorities. This consists of mapping forwarding classes to schedulers by using a scheduler map.
4. Configure priority groups. This consists of mapping forwarding classes (priorities) to forwarding class sets (priority groups) to define the priorities that belong to each priority group.
5. Define resources for the priority groups. This consists of configuring traffic control profiles to set minimum guaranteed bandwidth and maximum bandwidth for a priority group. Traffic control profiles also specify a scheduler map, which defines the resources

(schedulers) for the priorities in the priority group. Extra port bandwidth is shared among priority groups in proportion to the minimum guaranteed bandwidth of each priority group.

The traffic control profile bandwidth settings determine the port resources available to the priority group, and the schedulers specified in the scheduler map determine the amount of the priority group resources that each priority receives.

6. Apply the hierarchical scheduling to a port. This consists of attaching one or more priority groups to a port interface. For each priority group, you also attach a traffic control profile. Different priority groups on the same port can use different traffic control profiles.

Figure 206: Hierarchical Scheduling Packet Flow



Strict-High Priority Queues and Hierarchical Scheduling

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.



NOTE: On a QFabric system, if a fabric (fte) interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best effort, lossless, strict-high priority, and multidestination traffic.

Default Hierarchical Scheduling

If you do not explicitly configure hierarchical scheduling, the switch uses the default settings:

- The switch automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The switch assigns 100 percent of the port output bandwidth to the default forwarding class set. The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.
- Ingress traffic is classified based on the default classifier settings.
- The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings.

Related Documentation

- [Understanding CoS Packet Flow on page 5440](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [*Benefits of Configuring CoS Hierarchical Port Scheduling*](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Default CoS Scheduling and Classification on page 5475](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)

Understanding CoS Output Queue Schedulers

Output queue scheduling defines the class-of-service (CoS) properties of output queues (priorities). Queue scheduling works with priority group scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (priority group), and queue scheduling determines the portion of the priority group's bandwidth that a particular queue can use.

Scheduler maps associate queue schedulers with forwarding classes (output queues). You can associate each scheduler map with a traffic control profile, and then associate each traffic control profile with a forwarding class set (priority group) and a port interface. In conjunction with the priority group scheduling configured in the traffic control profile, queue scheduling configures the output queues, packet schedulers, and tail-drop processes that operate according to this mapping.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Output Queue Scheduling Components on page 5486](#)
- [Default Schedulers on page 5487](#)
- [Transmit Rate \(Minimum Guaranteed Bandwidth\) on page 5488](#)
- [Sharing Extra Bandwidth on page 5489](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 5489](#)
- [Scheduling Priority on page 5490](#)
- [Scheduler Drop-Profile Maps on page 5491](#)
- [Buffer Size on page 5491](#)
- [Scheduler Maps on page 5492](#)

Output Queue Scheduling Components

[Table 482 on page 5486](#) provides a quick reference to the scheduler components you can configure to determine the bandwidth properties of output queues, and [Table 483 on page 5487](#) provides a quick reference to some related scheduling configuration components.

Table 482: Output Queue Scheduler Components

Output Queue Scheduler Component	Description
Priority	Sets the scheduling priority applied to the queue.
Shaping rate	Sets the maximum bandwidth the queue can consume.

Table 482: Output Queue Scheduler Components (*continued*)

Output Queue Scheduler Component	Description
Transmit rate	Sets the minimum guaranteed bandwidth for the queue. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.
Drop profile	Sets the probability of dropping packets as the queue fills up.
Loss priority	Sets the traffic loss priority to which a drop profile applies.
Drop profile map	Maps a drop profile to a loss priority.
Buffer size	Sets the size of the queue buffer.

Table 483: Other Scheduling Components

Other Scheduling Components	Description
Scheduler map	Maps schedulers to queues (forwarding classes, also called priorities).
Forwarding class	Maps traffic to a queue (priority).
Traffic control profile	Sets scheduling for the forwarding class set (priority group) and associates a scheduler map with the forwarding class set to apply queue scheduling to the forwarding classes in the forwarding class set. Extra port bandwidth is shared among forwarding class sets in proportion to the minimum guaranteed bandwidth of each forwarding class set.
Forwarding class set	Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes.

Default Schedulers

Each forwarding class requires an associated scheduler. The default configuration uses only four forwarding classes, unicast best-effort (queue 0), unicast network-control (queue 3), multicast best-effort (queue 8), and multicast network control (queue 11). You can use the default schedulers or you can define new schedulers for these four forwarding classes. For any other forwarding class, you must explicitly configure a scheduler.

Table 484 on page 5487 shows the default schedulers.

Table 484: Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	Low	5%

Table 484: Default Schedulers (*continued*)

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
FCoE scheduler (queue 3)	35%	None	35%	Low	35%
No-loss scheduler (queue 4)	35%	None	35%	Low	35%
Network-control scheduler (queue 7)	5%	None	5%	Low	5%
Multidestination scheduler (queue 8)	20%	None	20%	Low	20%



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

There are no resources assigned by default to queues 1, 2, 5, 6, 9, 10, and 11. You can define new unicast forwarding classes to assign to queues 1, 2, 5, and 6, and new multidestination forwarding classes to assign to queues 9, 10, and 11, which do not have default mappings to a forwarding class. You can assign resources to these queues manually using schedulers. In addition, you can change the default forwarding-class-to-queue mappings and the default schedulers.

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues in the forwarding class set. When a forwarding class does not fully use the allocated minimum transmission bandwidth (transmit rate), other forwarding classes in the forwarding class set can use the remaining bandwidth if those forwarding classes receive more traffic than their allocated bandwidth.

Transmit Rate (Minimum Guaranteed Bandwidth)

The transmit rate determines the minimum guaranteed bandwidth for each forwarding class. It also determines how much excess (extra) bandwidth each low-priority queue can share; each queue shares extra bandwidth in proportion to its transmit rate. You specify the rate in bits per second as a fixed value such as 1 Mbps or as a percentage of the total forwarding class set minimum guaranteed bandwidth (the guaranteed rate set in the traffic control profile). Either the default scheduler or a scheduler you configure allocates a portion of the outgoing interface bandwidth to each forwarding class.



NOTE: For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

The allocated bandwidth can exceed the configured minimum rate if additional bandwidth is available from other queues in the forwarding class set. In case of congestion, the configured transmit rate is guaranteed for the queue. This property enables you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.



NOTE: Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

Sharing Extra Bandwidth

Extra bandwidth is available to low-priority queues when the minimum guaranteed bandwidth of the queues does not use the full amount of forwarding class set bandwidth. This extra bandwidth is shared among the forwarding classes in the set based on the minimum guaranteed bandwidth of each queue.

For example, in a forwarding class set, Queue A has a transmit rate of 1 Gbps, Queue B has a transmit rate of 1 Gbps, and Queue C has a transmit rate of 2 Gbps. After servicing the minimum guaranteed bandwidth of these queues, the forwarding class set has an extra 2 Gbps of bandwidth available, and all three queues still have packets to forward. The queues receive the extra bandwidth in proportion to their transmit rates, so Queue A receives an extra 500 Mbps, Queue B receives an extra 500 Mbps, and Queue C receives an extra 1 Gbps.

Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth each forwarding class can consume. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total forwarding class set maximum bandwidth (the shaping rate set in the traffic control profile).

The maximum bandwidth for a queue depends on the total bandwidth available to the forwarding class set to which the queue belongs and how much bandwidth the other queues in the forwarding class set consume.



NOTE: In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. A strict-high priority queue (or several queues with higher priorities than the starved queue) can consume all of the port bandwidth and prevent another queue from transmitting packets. To prevent a queue from being starved for bandwidth, you can configure a shaping rate on the queue or queues to prevent them from consuming all of the port bandwidth.

Scheduling Priority

Scheduling priority determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic receive better access to the outgoing interface. The priority setting in the scheduler determines the priority for the queue.

Two levels of scheduling priority are supported:

- **Low**—Low-priority queues transmit traffic based on the weighted round robin (WRR) algorithm. The scheduler first determines if an individual queue is within its defined bandwidth profile. The scheduler then regularly reevaluates whether each individual queue is within its defined bandwidth profile and compares the amount of data the queue transmits to the amount of bandwidth the scheduler allocates to the queue. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount. Out of profile queue data is transmitted only if bandwidth is available. Otherwise, it is buffered if buffer space is available. If no buffer space is available, the traffic may be dropped.
- **Strict-high**—You can configure only one queue as **strict-high** priority. The other 11 queues are **low** priority.

The **strict-high** priority queue receives preferential treatment over the low-priority queues. The **strict-high** priority queue receives all of its configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until the strict-high priority queue is empty. Carefully consider how much bandwidth you want to allocate to the **strict-high** priority queue to avoid starving the low-priority queues.

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.

- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- You cannot configure a minimum guaranteed bandwidth for a strict-high priority queue. (You cannot configure a transmit rate for a strict-high priority queue scheduler, and you cannot configure a guaranteed rate for a forwarding class set that has a strict-high priority queue.)

Junos OS performs priority queueing using the following steps:

1. Services the strict-high priority queue before any other queues are served
2. Services the minimum bandwidth (transmit rate) of low-priority queues until the minimum is met or the queues are empty
3. Services all other low-priority queues and needs that exceed the minimum bandwidth

Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. A drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type:

- PLP—Low, medium-high, high. You configure the PLP during classifier configuration. When you use a scheduler map to associate a forwarding class with a scheduler, you can use a drop-profile map to map different drop profiles to the forwarding class for different PLPs.
- Protocol type—Drop profiles match all protocol types.

Buffer Size

Most of the total system buffer space is divided into two buffer pools, shared buffers and dedicated buffers. Shared buffers are a global pool that the ports share dynamically as needed. Dedicated buffers are a reserved portion of the buffer pool that is distributed evenly to all of the ports. Each port receives an equal allocation of dedicated buffer space. The dedicated buffer allocation to ports is not configurable because it is reserved for the ports.

The queue buffers are allocated from the dedicated buffer pool assigned to the port. By default, ports divide their allocation of dedicated buffers among the egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (**transmit-rate**) for traffic. Only the queues included in the default scheduler receive dedicated buffers.

If you do not use the default configuration, you can explicitly configure the queue buffer size in either of two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is mapped to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining port dedicated buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit

scheduler means no dedicated buffer allocation for the queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



NOTE: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

For a complete discussion about queue buffer configuration in the context of ingress and egress port buffer configuration, see [“Understanding CoS Buffer Configuration” on page 5527](#).

Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map, associate the scheduler map with a traffic control profile, and then associate the traffic control profile with an interface and a forwarding class set.

You can associate up to four user-defined scheduler maps with traffic control profiles.

Related Documentation

- [Understanding Junos CoS Components on page 5436](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Buffer Configuration on page 5527](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5502](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Queue Scheduling Priority on page 5679](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)

Understanding CoS Priority Group Scheduling

Priority group scheduling defines the class-of-service (CoS) properties of a group of output queues (priorities). Priority group scheduling works with output queue scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (a priority group, called a forwarding class set in Junos OS configuration). Queue scheduling determines the portion of the priority group bandwidth that the particular queue can use.

You configure priority group scheduling in a traffic control profile and then associate the traffic control profile with a forwarding class set and an interface. You attach a scheduler map to the traffic control profile to specify the queue scheduling characteristics.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Priority Group Scheduling Components on page 5493](#)
- [Default Traffic Control Profile on page 5494](#)
- [Guaranteed Rate \(Minimum Guaranteed Bandwidth\) on page 5494](#)
- [Sharing Extra Bandwidth on page 5494](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 5495](#)
- [Scheduler Maps on page 5495](#)

Priority Group Scheduling Components

[Table 485 on page 5493](#) provides a quick reference to the traffic control profile components you can configure to determine the bandwidth properties of priority groups, and [Table 486 on page 5494](#) provides a quick reference to some related scheduling configuration components.

Table 485: Priority Group Scheduler Components

Traffic Control Profile Component	Description
Guaranteed rate	Sets the minimum guaranteed port bandwidth for the priority group. Extra port bandwidth is shared among priority groups in proportion to the guaranteed rate of each priority group on the port.
Shaping rate	Sets the maximum port bandwidth the priority group can consume.
Scheduler map	Maps schedulers to queues (forwarding classes, also called priorities). This determines the portion of the priority group bandwidth that a queue receives.

Table 486: Other Scheduling Components

Other Scheduling Components	Description
Forwarding class	Maps traffic to a queue (priority).
Forwarding class set	Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes.
Scheduler	Sets the bandwidth and scheduling priority of individual queues (forwarding classes).

Default Traffic Control Profile

There is no default traffic control profile.

Guaranteed Rate (Minimum Guaranteed Bandwidth)

The guaranteed rate determines the minimum guaranteed bandwidth for each priority group. It also determines how much excess (extra) port bandwidth the priority group can share; each priority group shares extra port bandwidth in proportion to its guaranteed rate. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total port bandwidth.

The minimum transmission bandwidth can exceed the configured rate if additional bandwidth is available from other priority groups on the port. In case of congestion, the configured guaranteed rate is guaranteed for the priority group. This property enables you to ensure that each priority group receives the amount of bandwidth appropriate to its level of service.



NOTE: Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

You cannot configure a guaranteed rate for forwarding class sets that include strict-high priority queues.

Sharing Extra Bandwidth

Extra bandwidth is available to priority groups when the priority groups do not use the full amount of available port bandwidth. This extra port bandwidth is shared among the priority groups based on the minimum guaranteed bandwidth of each priority group.

For example, Port A has three priority groups: fc-set-1, fc-set-2, and fc-set-3. Fc-set-1 has a guaranteed rate of 2 Gbps, fc-set-2 has a guaranteed rate of 2 Gbps, and fc-set-3 has a guaranteed rate of 4 Gbps. After servicing the minimum guaranteed bandwidth of these priority groups, the port has an extra 2 Gbps of available bandwidth, and all three priority groups have still have packets to forward. The priority groups receive the extra bandwidth in proportion to their guaranteed rates, so fc-set-1 receives an extra 500 Mbps, fc-set-2 receives an extra 500 Mbps, and fc-set-3 receives an extra 1 Gbps.

Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth the priority group can consume. You specify the rate in bits per second as a fixed value such as 5 Mbps or as a percentage of the total port bandwidth.

The maximum bandwidth for a priority group depends on the total bandwidth available on the port and how much bandwidth the other priority groups on the port consume.

Scheduler Maps

A scheduler map maps schedulers to queues. When you associate a scheduler map with a traffic control profile, then associate the traffic control profile with an interface and a forwarding class set, the scheduling defined by the scheduler map determines the portion of the priority group resources that each individual queue can use.

You can associate up to four user-defined scheduler maps with traffic control profiles.

Related Documentation

- [Understanding Junos CoS Components on page 5436](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5502](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)

Understanding CoS Traffic Control Profiles

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth that you assign to the forwarding class set in the traffic control profile.

This two-tier hierarchical scheduling architecture provides flexibility in allocating resources among queues and:

- Assigns a portion of port bandwidth to a priority group. You define the port resources for the priority group in a traffic control profile.
- Allocates priority group bandwidth among the queues that belong to the priority group. A scheduler map attached to the traffic control profile defines the amount of the priority group's resources that each queue can use.

Attaching a priority group and traffic control profile to a port defines the hierarchical scheduling properties of the group and the queues that belong to the group.

The ability to create priority groups supports enhanced transmission selection (ETS, described in IEEE 802.1Qaz). When a priority group does not use its allocated port bandwidth, ETS shares the excess port bandwidth among other priority groups on the port in proportion to their guaranteed minimum bandwidth (guaranteed rate). This utilizes the port bandwidth better than scheduling schemes that require setting strict priorities that reserve bandwidth for all groups whether it is needed or not. ETS allows traffic groups that need extra bandwidth to use it if the bandwidth is available, while preserving the ability to specify the minimum guaranteed bandwidth for traffic groups.

Traffic control profiles define the following CoS properties for priority groups:

- Minimum guaranteed bandwidth—Also known as the committed information rate (CIR). This is the minimum amount of port bandwidth the priority group receives. Priorities in the priority group receive their minimum guaranteed bandwidth as a portion of the priority group's minimum guaranteed bandwidth. The **guaranteed-rate** statement defines the minimum guaranteed bandwidth.



NOTE: You cannot apply a traffic control profile with a minimum guaranteed bandwidth to a priority group that includes strict-high priority queues.

- Shared excess (extra) bandwidth—When the priority groups on a port do not consume the full amount of bandwidth allocated to them or there is unallocated link bandwidth available, priority groups can contend for that extra bandwidth if they need it. Priorities in the priority group contend for extra bandwidth as a portion of the priority group's extra bandwidth. The amount of extra bandwidth for which a priority group can contend is proportional to the priority group's guaranteed minimum bandwidth (guaranteed rate).

- **Maximum bandwidth**—Also known as peak information rate (PIR). This is the maximum amount of port bandwidth the priority group receives. Priorities in the priority group receive their maximum bandwidth as a portion of the priority group's maximum bandwidth. The **shaping-rate** statement defines the maximum bandwidth.
- **Queue scheduling**—Each traffic control profile includes a scheduler map. The scheduler map maps priorities (forwarding classes) to schedulers to define the scheduling characteristics of the individual priorities in the priority group. The resources scheduled for each priority represent portions of the resources that the traffic control profile schedules for the entire priority group, not portions of the total link bandwidth. The **scheduler-maps** statement defines the mapping of forwarding classes to schedulers.

Related Documentation

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)

Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum Bandwidth)

You can set a guaranteed minimum bandwidth for individual forwarding classes (queues) and for groups of forwarding classes called forwarding class sets (priority groups). Setting a minimum guaranteed bandwidth ensures that priority groups and queues receive the bandwidth required to support the expected traffic.

This topic covers:

- [Guaranteeing Bandwidth Using Hierarchical Scheduling on page 5497](#)
- [Priority Group Guaranteed Rate \(Minimum Bandwidth\) on page 5499](#)
- [Queue Transmit Rate \(Minimum Bandwidth\) on page 5499](#)

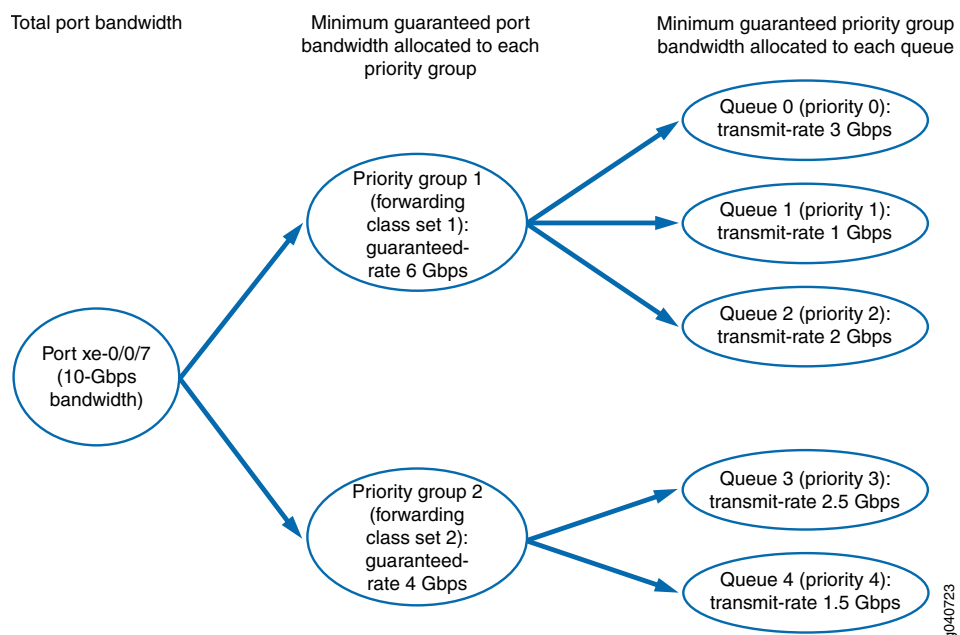
Guaranteeing Bandwidth Using Hierarchical Scheduling

The **guaranteed-rate** value for the priority group defines the minimum amount of bandwidth allocated to a forwarding class set on a port, whereas the **transmit-rate** value of the queue defines the minimum amount of bandwidth allocated to a particular queue in a priority group. The queue bandwidth is a portion of the priority group bandwidth.



NOTE: You cannot configure a minimum guaranteed bandwidth (transmit rate) for a forwarding class that is mapped to a strict-high priority queue, and you cannot configure a minimum guaranteed bandwidth (guaranteed rate) for a priority group that includes strict-high priority queues.

[Figure 207 on page 5498](#) shows how the total port bandwidth is allocated to priority groups (forwarding class sets) based on the guaranteed rate of each priority group. It also shows how the guaranteed bandwidth of each priority group is allocated to the queues in the priority group based on the transmit rate of each queue.

Figure 207: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling

The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue transmit rates cannot exceed the total guaranteed rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.



NOTE: You must set both the priority group **guaranteed-rate** value and the queue **transmit-rate** value in order to configure the minimum bandwidth for individual queues. If you set the **transmit-rate** value but do not set the **guaranteed-rate** value, the configuration fails.

You can set the **guaranteed-rate** value for a priority group without setting the **transmit-rate** value for individual queues in the priority group. However, queues that do not have a configured **transmit-rate** value can become starved for bandwidth if other higher-priority queues need the priority group's bandwidth. To avoid starving a queue, it is a good practice to configure a **transmit-rate** value for most queues.

If you configure the guaranteed rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Priority Group Guaranteed Rate (Minimum Bandwidth)

Setting a priority group **guaranteed-rate** enables you to reserve a portion of the port bandwidth for the forwarding classes (queues) in that forwarding class set. The minimum bandwidth (**guaranteed-rate**) that you configure for a priority group sets the minimum bandwidth available to all of the forwarding classes in the forwarding class set.

The combined **guaranteed-rate** value of all of the forwarding class sets associated with an interface cannot exceed the amount of bandwidth available on that interface.

You configure the priority group **guaranteed-rate** in the traffic control profile. You cannot apply a traffic control profile that has a guaranteed rate to a priority group that includes strict-high priority queues.

Queue Transmit Rate (Minimum Bandwidth)

Setting a queue **transmit-rate** enables you to reserve a portion of the priority group bandwidth for the individual queue. For example, a queue that handles Fibre Channel over Ethernet (FCoE) traffic might require a minimum rate of 4 Gbps to ensure the class of service that storage area network (SAN) traffic requires.

The priority group **guaranteed-rate** sets the aggregate minimum amount of bandwidth available to the queues that belong to the priority group. The cumulative total minimum bandwidth the queues consume cannot exceed the minimum bandwidth allocated to the priority group to which they belong. (The combined transmit rates of the queues in a priority group cannot exceed the priority group's guaranteed rate.)

You must configure the **guaranteed-rate** value of the priority group in order to set a **transmit-rate** value for individual queues that belong to the priority group. The reason is that if there is no guaranteed bandwidth for a priority group, there is no way to guarantee bandwidth for queues in that priority group.

You configure the queue **transmit-rate** in the scheduler configuration. You cannot configure a transmit rate for strict-high priority queues.

**Related
Documentation**

- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)

Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)

If the amount of traffic on an interface exceeds the maximum bandwidth of the interface, it leads to congestion. You can use priority group (forwarding class set) shaping and queue shaping to manage the excess traffic and avoid congestion.

The maximum bandwidth sets the most bandwidth a priority group or a queue can use after all of the priority group and queue minimum bandwidth requirements are met, even if more bandwidth is available.

This topic covers:

- [Priority Group Shaping on page 5500](#)
- [Queue Shaping on page 5500](#)
- [Shaping Maximum Bandwidth Using Hierarchical Scheduling on page 5501](#)

Priority Group Shaping

Priority group shaping enables you to shape the aggregate traffic of a forwarding class set on a port to a maximum rate that is less than the line or port rate. The maximum bandwidth (**shaping-rate**) that you configure for a priority group sets the maximum bandwidth available to all of the forwarding classes (queues) in the forwarding class set.

If a port has more than one priority group and the combined **shaping-rate** value of the priority groups is greater than the amount of port bandwidth available, the bandwidth is shared proportionally among the priority groups.

You configure the priority group **shaping-rate** in the traffic control profile.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-high priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues. Similarly, for any queue, you can configure queue shaping (**shaping-rate**) to set the maximum bandwidth for a particular queue.

The **shaping-rate** value of the priority group sets the aggregate maximum amount of bandwidth available to the queues that belong to the priority group. The cumulative total bandwidth the queues consume cannot exceed the maximum bandwidth of the priority group to which they belong on a port.

If a priority group has more than queue and the combined **shaping-rate** value of the queues is greater than the amount of bandwidth available to the priority group, the bandwidth is shared proportionally among the queues.

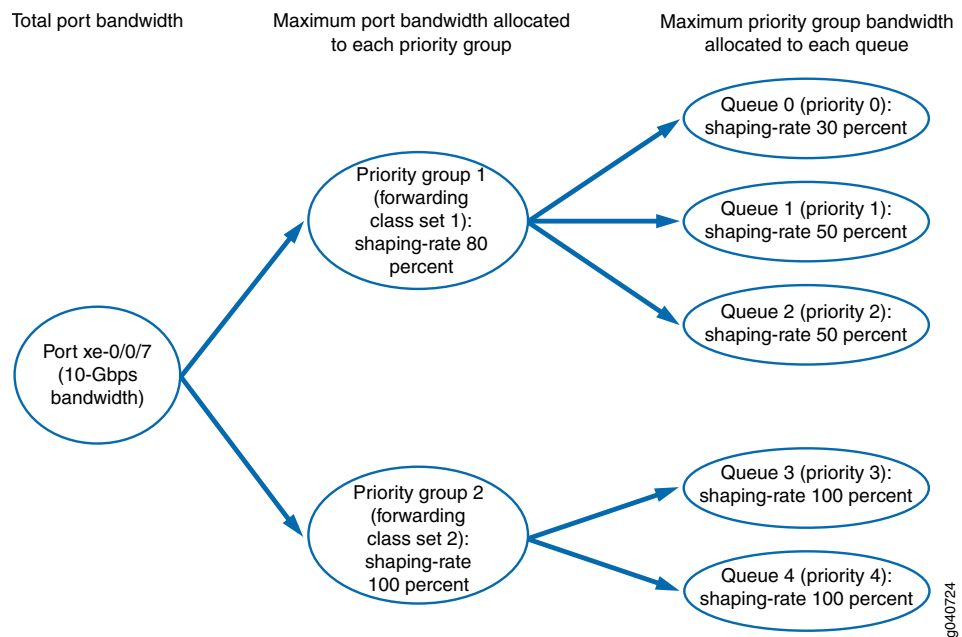
You configure the queue **shaping-rate** in the scheduler configuration.

Shaping Maximum Bandwidth Using Hierarchical Scheduling

Priority group shaping defines the maximum bandwidth allocated to a forwarding class set on a port, whereas queue shaping defines a limit on maximum bandwidth usage per queue. The queue bandwidth is a portion of the priority group bandwidth.

Figure 208 on page 5501 shows how the port bandwidth is allocated to priority groups (forwarding class sets) based on the shaping rate of each priority group, and how the bandwidth of each priority group is allocated to the queues in the priority group based on the shaping rate of each queue.

Figure 208: Setting Maximum Bandwidth Using Hierarchical Scheduling



Related Documentation

- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)

- [Defining CoS Queue Schedulers on page 5789](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)

Understanding CoS Scheduling Behavior and Configuration Considerations

Many factors affect scheduling configuration and bandwidth requirements, including:

- When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.
- When you define a forwarding class that will be used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:
 - Mapping a scheduler to the forwarding class in a scheduler map
 - Including the forwarding class in a forwarding class set
 - Associating the scheduler map with a traffic control profile
 - Attaching the traffic control profile to a forwarding class set and an interface
- On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.
- For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.
- Configuring the minimum guaranteed bandwidth (**transmit-rate**) for a queue (forwarding class) does not work unless you also configure the minimum guaranteed bandwidth (**guaranteed-rate**) for the priority group (forwarding class set) in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) If you configure transmit rates whose sum exceeds the guaranteed rate of the forwarding class set, the commit check fails and the system rejects the configuration.

- The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

- If you configure the **guaranteed-rate** of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.
- There are several factors to consider if you want to configure strict-high priority queues:
 - You cannot configure a minimum guaranteed bandwidth (**transmit-rate**) for a strict-high priority queue. You cannot configure a minimum guaranteed bandwidth (**guaranteed-rate**) for a forwarding class set that includes a strict-high priority queue.
 - You must create a separate forwarding class set for the strict-high priority queue.
 - Only one forwarding class set can contain strict-high priority queues.
 - Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
 - A strict-high priority queue cannot belong to a multidestination forwarding class set.
- In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. Failure of a queue to transmit packets for 12 consecutive seconds may be due to:
 - A strict-high priority queue consuming all of the port bandwidth
 - Several queues consuming all of the port bandwidth
 - Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
 - Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

If the cause is a strict-high priority queue consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the strict-high priority queue and prevent it from using all of the port bandwidth. To configure rate shaping, include the **shaping-rate (rate | percent percentage)** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level and apply the shaping rate to the strict-high priority scheduler.

If several queues consume all of the port bandwidth, you can use a scheduler to rate shape those queues and prevent them from using all of the port bandwidth.

- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.
- When you set the maximum bandwidth for a queue or for a priority group (**shaping-rate**) at 100 Kbps or lower, the traffic shaping behavior is accurate only within +/– 20 percent of the configured **shaping-rate**.
- Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the

ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



NOTE: Do not configure drop profiles for the **fcoe** and **no-loss** forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

- On an ingress port, do not configure classifiers that map the same IEEE 802.1p code point to both a multdestination traffic flow and a lossless unicast traffic flow (such as the default lossless **fcoe** or **no-loss** forwarding classes). Any code point used for multdestination traffic on a port should not be used to classify unicast traffic into a lossless forwarding class on the same port.

If a multdestination traffic flow and a lossless unicast traffic flow use the same code point on a port, the multdestination traffic is treated the same way as the lossless traffic. For example, if priority-based flow control (PFC) is applied to the lossless traffic, the multdestination traffic of the same code point is also paused. During periods of congestion, treating multdestination traffic the same as lossless unicast traffic can create ingress port congestion for the multdestination traffic and affect the multdestination traffic on all of the egress ports the multdestination traffic uses.

For example, the following configuration can cause ingress port congestion for the multdestination flow:

1. For unicast traffic, IEEE 802.1p code point 011 is classified into the **fcoe** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class fcoe loss-priority low code-points 011
```
2. For multdestination traffic, IEEE 802.1p code point 011 is classified into the **mcast** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast loss-priority low code-points 011
```
3. The unicast classifier that maps traffic with code point 011 to the **fcoe** forwarding class is mapped to interface **xe-0/0/1**:

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl
```
4. The multdestination classifier that maps traffic with code point 011 to the **mcast** forwarding class is mapped to all interfaces (multdestination traffic maps to all interfaces and cannot be mapped to individual interfaces):

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the same code point (**011**) maps unicast traffic to a lossless traffic flow and also maps multdestination traffic to a multdestination traffic flow, the multdestination traffic flow might experience ingress port congestion during periods of congestion.

To avoid ingress port congestion, do not map the code point used by the multdestination traffic to lossless unicast traffic. For example:

1. Instead of classifying code point **011** into the **fcoe** forwarding class, classify code point **011** into the **best-effort** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class best-effort loss-priority low code-points 011
```
2.

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast loss-priority low code-points 011
```
3.

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl
```
4.

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the code point **011** does not map unicast traffic to a lossless traffic flow, the multdestination traffic flow does not experience ingress port congestion during periods of congestion.

The best practice is to classify unicast traffic with IEEE 802.1p code points that are also used for multdestination traffic into best-effort forwarding classes.

Related Documentation

- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

Junos OS Release 12.3 increased support for lossless priorities from two lossless forwarding classes to up to six lossless forwarding classes on QFX3500 and QFX3600 switches. (Junos OS Release 13.1 introduced the same support on QFabric systems; throughout this document, features introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.) Each forwarding class is mapped to an IEEE 802.1p code point (priority).

Earlier Junos OS software releases supported two lossless forwarding classes, the default *fcoe* and *no-loss* forwarding classes, which are mapped by default to IEEE 802.1p priorities 3 (code point 011) and 4 (code point 100), respectively. Junos OS Release 12.3 also introduced a new output stanza in the congestion notification profile (CNP) to configure priority-based flow control (PFC) on output queues.

The default configuration is the same as the default configuration in Junos OS Release 12.2 and is backward-compatible. If you need only two (or fewer) lossless forwarding classes, use the default configuration. If you need more than two lossless forwarding classes, you can use the two default forwarding classes and configure additional lossless forwarding classes. If you do not want to use the default lossless forwarding classes, you can change them or use only the lossless forwarding classes that you explicitly configure.

- [Lossless Transport Features Introduced in Junos OS Release 12.3 on page 5506](#)
- [Default Lossless Priority Configuration on page 5507](#)
- [Configuring Lossless Priorities on page 5509](#)
- [Backward Compatibility with Junos OS Releases Earlier Than Release 12.3 on page 5522](#)
- [Configuration Rules and Recommendations on page 5523](#)

Lossless Transport Features Introduced in Junos OS Release 12.3

Support for lossless transport introduced in Junos OS Release 12.3 includes:

- Configuring up to six lossless forwarding classes.
- Configuring PFC pause on output queues to program the output queues that can respond to PFC pause messages received from the connected peer. The priorities you pause on output queues must match the priorities on which you enable PFC on the corresponding ingress interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101), then you must also enable pause on priorities 3 and 5 on the corresponding ingress interfaces. Configuring flow control on the output queues and enabling PFC on the corresponding input queues allows you to pause up to six priorities (forwarding classes).
- Controlling the headroom buffer on Ethernet interfaces by configuring the maximum receive unit (MRU) size for the traffic mapped to an IEEE 802.1p priority (configured per priority) and the length of the attached cable (configured per interface). The MRU size can range up to full jumbo packet size (9216 bytes).
- Remapping (rewriting) IEEE 802.1p priorities on native Fibre Channel (FC) interfaces when the system is acting as an FCoE-FC gateway. If the Ethernet (FCoE) network uses a different IEEE 802.1p priority than priority 3 (011) for FCoE traffic, then you can

use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that different priority (see [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#)).

Lossless transport still requires configuring previously existing features, including enabling PFC on the lossless priorities on ingress interfaces, and configuring classifiers to classify incoming traffic into lossless forwarding classes based on the IEEE 802.1p priority tag of the packet.



NOTE: If you expect a large amount of lossless traffic on your network and configure multiple lossless traffic classes, ensure that you reserve enough scheduling resources (bandwidth) and lossless headroom buffer space to support the lossless flows. ([“Understanding CoS Buffer Configuration” on page 5527](#) describes how to configure buffers and provides a recommended buffer configuration for networks with larger amounts of lossless traffic.)

Default Lossless Priority Configuration

If you do not explicitly configure forwarding classes, the system uses the default forwarding class configuration, which provides two default lossless forwarding classes (*fcoe* and *no-loss*). (If you change the forwarding class configuration on a QFX Series switch or on a QFabric Node device, the changes apply to all traffic on that device because forwarding classes are global to a particular device.)

If you do not explicitly configure classifiers, and you do not explicitly configure flow control to pause output queues (configured in the output stanza of the CNP), the default classifier and the default output queue pause configuration are applied to all Ethernet interfaces on QFX Series switches and Node devices. You can override the default classifier and the default output queue pause configuration on a per-interface basis by applying an explicit configuration to an Ethernet interface. The default configuration is used on all Ethernet interfaces that do not have an explicit configuration.



NOTE: If you do not configure flow control on output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause in the output stanza of the CNP.

In the default configuration, only queue 3 and queue 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza of the CNP. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza of the CNP.

The default configuration is the same as the default configuration in software releases earlier than Junos OS Release 12.3, and provides the same lossless behavior:

- There are two default lossless forwarding classes (the no-loss packet drop attribute is applied automatically):
fcoe—Mapped to output queue 3
no-loss—Mapped to output queue 4
- The default classifier maps the fcoe forwarding class to IEEE 802.1p priority 3 (011) and the no-loss forwarding class to IEEE 802.1p priority 4 (100)
- Priority-based flow control (PFC) is enabled on Ethernet interface output queues 3 and 4 when those queues carry lossless traffic (traffic that is mapped to the fcoe and no-loss forwarding classes, respectively). In Junos OS software releases earlier than Release 12.3, output queue flow control was not user-configurable.

On native FC interfaces (NP_Ports), default flow control is enabled on output queue 3 (IEEE 802.1p priority 3) for FCoE/FC traffic.

- PFC must be enabled explicitly on the lossless IEEE 802.1p priorities (code points) on ingress Ethernet interfaces; no default PFC configuration is applied at ingress interfaces. If you do not enable PFC on lossless priorities, those priorities might experience packet loss during periods of congestion. For example, if you want lossless FCoE traffic and you are using the default fcoe forwarding class, you use a CNP to enable PFC on priority 3 (code point 011), and apply that CNP to all ingress interfaces that carry FCoE traffic.
- On Ethernet ports, PFC buffer calculations use the following default values to determine the headroom buffer size:
Cable length—100 meters (approximately 328 feet)
MRU for priority 3 traffic—2500 bytes
MRU for priority 4 traffic—9216 bytes
Maximum transmission unit (MTU)—1522 (or the configured MTU value for the interface)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.

- DCBX is enabled on all interfaces in autonegotiation mode, and automatically exchanges FCoE application protocol type, length, and values (TLVs) on interfaces that carry FCoE traffic. However, if you explicitly configure DCBX protocol TLV exchange for any application, then you must explicitly configure protocol TLV exchange for every application for which you want DCBX to exchange TLVs, including FCoE.

The default CoS configuration is backward-compatible with the *default* CoS configuration of software releases before Junos OS Release 12.3. If you explicitly configure lossless transport, ensure that the input and output queues corresponding to the lossless forwarding classes are explicitly configured for PFC pause.



NOTE: If you *explicitly* configured the lossless fcoe or no-loss forwarding classes before upgrading from a release earlier than Junos OS Release 12.3, those forwarding classes are *not* lossless after the upgrade to Junos OS Release 12.3 or later. To regain lossless behavior, you can delete the explicit configuration and use the default lossless forwarding classes, or you can use the no-loss packet drop attribute introduced in Junos OS Release 12.3 to configure the forwarding classes for lossless behavior.

Table 487 on page 5509 summarizes the default unicast forwarding classes and their mapping to output queues, IEEE 802.1p priorities, and drop attributes.

Table 487: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute

Forwarding Class Name	Output Queue	Priority	Drop Attribute
best-effort	0	0	drop
fcoe	3	3	no-loss
no-loss	4	4	no-loss
network-control	7	7	drop

There is one default multdestination forwarding class named *mcast* for multicast, broadcast, and destination lookup fail (DLF) traffic that is mapped to output queue 8 with a drop attribute of drop. (Incoming multdestination traffic on all IEEE 802.1p priorities is mapped to the mcast forwarding class by default.)

Configuring Lossless Priorities

Configuring more than two lossless priorities (forwarding classes), or changing the default mapping of lossless forwarding classes to priorities and paused output queues, requires explicit configuration. Configuring lossless priorities includes:

- Configuring forwarding classes with the no-loss packet drop attribute
- Using a CNP to configure PFC on ingress interfaces and flow control (PFC) on egress interfaces
- Configuring a classifier to map IEEE 802.1p priorities (code points) to the correct forwarding classes (the forwarding classes for which you want lossless transport)

In addition, on Ethernet interfaces, DCBX must exchange the appropriate application protocol TLVs for the lossless traffic, and when the switch acts as an FCoE-FC gateway, you need to remap the FCoE priority on native FC interfaces if your network uses a priority other than 3 (IEEE code point 011) for FCoE traffic. This section describes:

- [Configuring Lossless Forwarding Classes \(Packet Drop Attribute\) on page 5510](#)
- [Congestion Notification Profiles \(PFC Configuration\) on page 5511](#)

- [Configuring DCBX \(Application Protocol TLV Exchange\) on page 5517](#)
- [Fate Sharing Among Traffic Classes on page 5518](#)
- [Transit Switch Configuration Versus FCoE-FC Gateway Configuration on page 5519](#)
- [Configuration Results and Commit Checks on page 5520](#)

Configuring Lossless Forwarding Classes (Packet Drop Attribute)

Junos OS Release 12.3 introduced the *no-loss* parameter for forwarding class configuration. (Although it uses the same name, this is not the no-loss default forwarding class. It is a packet drop attribute you can specify to configure any unicast forwarding class as a lossless forwarding class.)

You can configure up to six forwarding classes (depending on system architecture and the availability of system resources) as lossless forwarding classes by including the **no-loss** drop attribute at the **[edit class-of-service forwarding-classes class forwarding-class-name queue-num queue-number]** hierarchy level.

If you do not explicitly configure the fcoe or no-loss forwarding classes, they include the no-loss drop attribute by default. If you explicitly configure the fcoe or no-loss forwarding classes and you want to retain their lossless behavior, you must include the no-loss drop attribute in the configuration.



NOTE: All forwarding classes mapped to the same output queue must have the same packet drop attribute. (All forwarding classes mapped to the same output queue must be either lossy or lossless. You cannot map a lossy and a lossless forwarding class to the same queue.)

To avoid fate sharing (different flows receiving the same CoS treatment), use a one-to-one mapping of lossless forwarding classes to IEEE 802.1p code points (priorities) and queues. (Each forwarding class should be mapped to a different queue and classified into a different priority.) The classifier attached to the interface determines the forwarding class to priority mapping.

The fcoe and no-loss forwarding classes are special cases, because in the default configuration, they are configured for lossless behavior (providing that you also enable PFC on the priorities mapped to the fcoe and no-loss forwarding classes in the CNP input stanza).

[Table 488 on page 5511](#) summarizes the possible configurations of the fcoe and no-loss forwarding classes in Junos OS Release 12.3 and later, and the result of those configurations in terms of lossless traffic behavior. It is assumed that PFC, DCBX, and classifiers are properly configured.

Table 488: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3

Explicit (User-Configured) or Default Forwarding Class Configuration	Packet Drop Attribute	Result and Notes
Default	Default	The fcoe and no-loss forwarding classes are lossless. NOTE: Even if you explicitly configure other forwarding classes (lossy or lossless forwarding classes), the fcoe and no-loss forwarding classes remain lossless because they are not explicitly configured.
Explicit	Not specified in the explicit forwarding class configuration	The fcoe and no-loss forwarding classes are lossy because they do not include the no-loss drop attribute.
Explicit	No-loss	The fcoe and no-loss forwarding classes are lossless.
Explicit, configured in Junos OS Release 12.2 or earlier	Not specified (packet drop attribute was not available before Junos OS Release 12.3)	The fcoe and no-loss forwarding classes are lossy in Junos OS Release 12.3 and later because they do not include the no-loss drop attribute. NOTE: To retain lossless behavior, before you upgrade to Junos OS Release 12.3, delete the explicit configuration so that the system uses the default configuration. Alternatively, you can reconfigure the forwarding classes with the no-loss packet drop attribute after upgrading to Junos OS Release 12.3 or later.

For all other forwarding classes, you must explicitly configure lossless transport by specifying the no-loss packet drop attribute, because the default configuration for all other forwarding classes is lossy.

Congestion Notification Profiles (PFC Configuration)

Use CNPs to configure lossless PFC characteristics on input and output interfaces.

The input stanza of a CNP enables PFC on specified IEEE 802.1p priorities (code points) and fine-tunes headroom buffer settings by configuring the maximum receive unit (MRU) value and cable length on ingress interfaces.

The output stanza of a CNP enables PFC (flow control) on output queues for specified IEEE 802.1p priorities so that the queues can respond to PFC pause messages from the connected peer on the priority of your choice. (By default, output queues 3 and 4 respond to received PFC messages when those queues carry lossless traffic in the fcoe and no-loss forwarding classes, respectively.)

To achieve lossless transport, the priority paused at the ingress interfaces must match the priority paused at the egress interfaces for a given traffic flow. For example, if you configure ingress interfaces to pause traffic tagged with IEEE 802.1p priority 5 (code point 101) and priority 5 traffic is mapped to output queue 5, then you must also configure the corresponding output interfaces to pause priority 5 on queue 5. In addition, the forwarding class mapped to queue 5 must be configured as a lossless forwarding class (using the no-loss drop attribute).



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
 3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

Configuring Input Interface Flow Control (PFC and Headroom Buffer Calculation)

On Ethernet interfaces, the input stanza of the CNP enables PFC on specified priorities so that the ingress interface can send a pause message to the connected peer during periods of congestion. Input CNPs also fine-tune the headroom buffers used for PFC support by allowing you to configure the MRU value and cable length (if you do not want to use the default configuration).

Headroom buffers support lossless transport by storing the traffic that arrives at an interface after the interface sends a PFC flow control message to pause incoming traffic.

Until the connected peer receives the flow control message and pauses traffic, the interface continues to receive traffic and must buffer it (and the traffic that is still on the wire after the peer pauses) to prevent packet loss.

The system uses the MRU and the length of the attached physical cable to calculate buffer headroom allocation. The default configuration values are:

- MRU for priority 3 traffic—2500 bytes
- MRU for priority 4 traffic—9216 bytes
- Cable length—100 meters (approximately 328 feet)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not explicitly configure an MRU value, the default MRU value is 2500 bytes.

You can fine-tune the MRU and the cable length to adjust the size of the headroom buffer on an interface. The QFX Series has a shared global buffer pool and dynamically allocates headroom buffer space to lossless queues as needed.

A lower MRU or a shorter cable length reduces the amount of headroom buffer required on an interface and leaves more headroom buffer space for other interfaces. A higher MRU or a longer cable length increases the amount of headroom buffer space required on an interface and leaves less headroom buffer space for other interfaces.

In many cases, you can better utilize the headroom buffers by reducing the MRU value (for example, an MRU of 2180 is sufficient for most FCoE networks) and by reducing the cable length value if the physical cable is less than 100 meters long.



NOTE: When you configure the headroom buffers by changing the MRU or the cable length, and commit the configuration, the system performs a commit check and rejects the configuration if sufficient headroom buffer space is not available.

However, the system does not perform a commit check but instead returns a syslog error if:

- The buffers are configured on a LAG interface.
- The default classifier is used on the interface (instead of a user-configured classifier).
- The interface has not been created yet.

Configuring Output Interface Flow Control (PFC)

On Ethernet interfaces, you can use the output stanza of the CNP to configure flow control on unicast output queues and enable PFC pause response on specified IEEE 802.1p

priorities. By default, output queues 3 and 4 are enabled for PFC pause on priorities 3 (IEEE 802.1p code point 011) and 4 (IEEE 802.1p code point 100). The default PFC pause response supports the default lossless forwarding class configuration, which maps the fcoe forwarding class to queue 3 and priority 3, and maps the no-loss forwarding class to queue 4 and priority 4.

Configuring PFC on output queues enables you to pause any priority on any unicast output queue on any Ethernet interface. Output flow control enables you to use more than two output queues to support lossless traffic flows (you can configure up to six lossless forwarding classes and map them to different output queues that are enabled for PFC pause). Output queue flow control also enables you to support multiple lossless forwarding classes (each mapped to a different priority and output queue) for one class of traffic.



NOTE: Output flow control only works when PFC is enabled in the CNP input stanza on the corresponding priorities on the interface.

For example, if the converged Ethernet network uses two different priorities for FCoE traffic (for example, priority 3 and priority 5), then you can classify those priorities into different lossless forwarding classes that are mapped to different output queues by:

1. Configuring two lossless forwarding classes for FCoE traffic, with each forwarding class mapped to a different output queue. For example, you could use the default fcoe forwarding class, which is mapped to queue 3, and you could configure a second lossless forwarding class called fcoe1 and map it to queue 5. The fcoe forwarding class is for priority 3 FCoE traffic (code point 011), and the fcoe1 forwarding class is for priority 5 (code point 101) FCoE traffic.
2. Configuring a classifier that maps each forwarding class to the desired IEEE 802.1p code point (priority). If FCoE traffic on both priorities uses one interface, the classifier must classify both forwarding classes to the correct priorities. If FCoE traffic of different priorities uses different interfaces, the classifier configuration on each interface must map the correct priority to the corresponding lossless forwarding class.
3. Applying the classifier to the interfaces that carry FCoE traffic. The classifier determines the mapping of forwarding classes to priorities on each interface.

To configure lossless transport for these forwarding classes, you also need to:

- Enable PFC on the two priorities (3 and 5 in this example) at the ingress interfaces in the CNP input stanza.
- Configure PFC on the output queues and priorities for the forwarding classes in the CNP output stanza so that the interface can respond to pause messages received from the connected peer.



NOTE: When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

- Configure DCBX to exchange application protocol TLVs on both FCoE priorities.



NOTE: If you do not configure flow control to pause output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer, and you must explicitly enable PFC on the corresponding priorities in the CNP input stanza to achieve lossless behavior.

If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause. For example, if you explicitly configure flow control on output queue 5, the default configuration is no longer valid, and only output queue 5 is enabled for PFC pause. Output queues 3 and 4 are no longer enabled for PFC pause, so traffic using those queues no longer responds to PFC pause messages even if the corresponding forwarding class is configured with the no-loss drop attribute. To retain the pause configuration on output queues 3 and 4 and configure flow control on queue 5, you need to explicitly configure flow control on queues 3, 4, and 5.

You cannot configure flow control to pause a multidestination output queue. You can configure flow control to pause only unicast output queues.

Output Interface Flow Control Profiles

Configuring the CNP output stanza creates an output flow control profile that tells egress ports the queues on which the Ethernet interface should respond to PFC pause messages. Although you can create an unlimited number of CNPs that contain input stanzas only, the number of CNPs that you can configure with output stanzas is limited:

- For QFX3500 and QFX3600 standalone switches that are not part of a QFabric system, you can configure up to two output interface flow control profiles. (You can configure up to two CNPs with output stanzas.)
- For QFabric systems, you can configure one output interface flow control profile per Node device. (You can configure one CNP with an output stanza per Node device.)

There are a total of four output flow control profiles.

The system has a default output flow control profile that is applied to all Ethernet interfaces when the CNP attached to the interface has only an input stanza and does

not include an output stanza. The default profile responds to PFC pause messages received on queue 3 (for priority 3, for the default fcqe forwarding class) and on queue 4 (for priority 4, for the default no-loss forwarding class), and is effective only if PFC is configured on those priorities in the CNP input stanza.

Additionally, the system has two internal output flow control profiles that it applies automatically to fabric (FTE) ports and to native FC interfaces (NP_Ports). When the QFX3500 switch or the QFX3600 switch is not part of a QFabric system, the profile normally used for FTE ports is available for user configuration and provides a second user-configurable profile. (That is why standalone QFX3500 and QFX3600 switches have two user-configurable output flow control profiles, but QFX3500 and QFX3600 switches that are part of a QFabric have only one user-configurable output flow control profile.)

Because one output CNP can configure PFC pause response on multiple output queues (priorities), one user-configurable output CNP is usually flexible enough to specify the desired PFC response on all programmed interfaces.



NOTE: Each port can use one output flow control profile. You cannot apply more than one profile to one port.

Output flow control profiles can be expressed in table format. For example, [Table 489 on page 5516](#) shows the default output flow control profile that pauses priorities 3 and 4 on queues 3 and 4 (remember that PFC must also be enabled on code points 3 and 4 in the CNP input stanza in order for PFC to work):

Table 489: Default Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	—
6 (110)	—
7 (111)	—

[Table 490 on page 5517](#) is an example of a user-configured output flow control profile. Using the example from the preceding section, the CNP output stanza configures flow control on output queue 5, and also explicitly configures output flow control on queues

3 and 4 for the fcoe and no-loss forwarding classes. (If you explicitly configure an output CNP, you must explicitly configure every output queue that you want to respond to PFC messages, because the user-configured profile overrides the default profile. If this example did not include queues 3 and 4, those queues would no longer respond to received PFC messages.)

Table 490: User-Configured Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	5
6 (110)	—
7 (111)	—

Remember that you must also enable PFC on code points 3, 4, and 5 in the CNP input stanza for this configuration to work. When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

Configuring DCBX (Application Protocol TLV Exchange)

For applications that require lossless transport, DCBX exchanges application protocol TLVs with the connected peer interface. By default, DCBX advertises FCoE application protocol TLVs on all interfaces that are enabled for DCBX, and by default, DCBX is enabled on all interfaces. DCBX advertises no other applications by default.

For each application (for example, iSCSI) that you want to configure for lossless transport, you must enable the interfaces which carry that application traffic to exchange DCBX protocol TLVs with the connected peer. The TLV exchange allows the peer interfaces to negotiate a compatible configuration to support the application.

If you configure DCBX to advertise any application, the default DCBX advertisement is overridden, and DCBX advertises only the configured applications. If you want an interface to advertise only the FCoE application, you do not have to configure DCBX application protocol TLV exchange; instead, you can use the default configuration.

If you want DCBX to advertise other applications, you must explicitly configure an application map and apply it to the interfaces on which you want to exchange protocol TLVs for those applications. If you want to exchange FCoE application protocol TLVs in

addition to other application protocol TLVs, you must also explicitly configure the FCoE application in the application map. “[Understanding DCBX Application Protocol TLV Exchange](#)” on page 5060 describes how application mapping works.



NOTE: Lossless transport also requires that you enable PFC on the correct priority (IEEE 802.1p code point) on the ingress interfaces using an input CNP. If the priority you pause at the ingress interfaces is not mapped to queue 3 or queue 4 (the two output queues that are enabled for PFC pause flow control by default), then you must also enable the output queues that correspond to paused input priorities to pause using the output stanza of the CNP.

Fate Sharing Among Traffic Classes

You can configure different lossless (or lossy) traffic flows to share fate—that is, to receive the same CoS treatment.

Fate sharing is not desirable for I/O convergence. Instead of independent control of the fate of each type of flow, different types of flows receive the same treatment. Fate sharing is particularly undesirable for lossless flows. If one lossless flow experiences congestion and must be paused, that affects flows that share fate with the congested flow even if the other flows are not experiencing congestion, and also can cause ingress port congestion. If your network requires that all 802.1p priorities be lossless, you can achieve that by allowing some fate sharing among the eight priorities by spreading them across up to six lossless forwarding classes.

If the number of lossless priorities is less than or equal to the number of configured lossless forwarding classes, then you can avoid fate sharing by configuring a one-to-one mapping of forwarding classes to IEEE 802.1p code points (priorities) and output queues. (Each forwarding class should be mapped to a different output queue and classified to a different priority.)

If you want to configure different traffic flows to share fate, the QFX Series supports two fate-sharing configurations: mapping one forwarding class to more than one IEEE 802.1p code point (priority), and mapping two forwarding classes to the same output queue:

1. If you map one lossless forwarding class to more than one priority, the traffic tagged with each of the priorities uses the same CoS properties associated (the CoS properties associated with the forwarding class). For example, configuring a forwarding class called `fc1`, mapping it to queue 1, and mapping it to code points 101 and 110 using a classifier named `classify1` results in the traffic tagged with priorities 101 and 110 sharing fate:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 110
```

In this case, if the traffic mapped to either priority experiences congestion, both priorities are paused because they are mapped to the same forwarding class and are therefore treated similarly.

2. If you map multiple lossless forwarding classes to the same output queue, the traffic mapped to the forwarding classes uses the same output queue. This increases the amount of traffic the queue needs to buffer and forward, and can create congestion that affects all of the traffic flows that are mapped to the queue. For example, configuring two forwarding classes called fc1 and fc2, mapping both forwarding classes to queue 1, and mapping the forwarding classes to code points 101 and 110 (respectively) using a classifier named classify1 results in the traffic tagged with priorities 101 and 110 sharing fate on the same output queue:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service forwarding-classes class fc2 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1p classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1p classify1 forwarding class fc2
loss-priority low code-points 110
```

in this case, even though the two forwarding classes use different IEEE 802.1p priorities, if one forwarding class experiences congestion, it affects the other forwarding class. The reason is that if the output queue is paused because of congestion on either forwarding class, all traffic that uses that queue is paused. Since both forwarding classes are mapped to the queue, the traffic mapped to both forwarding classes is paused.



NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

Transit Switch Configuration Versus FCoE-FC Gateway Configuration

On a transit switch (all Ethernet ports, no native FC ports) that forwards FCoE traffic (or other traffic that requires lossless transport across the Ethernet network), the configuration of classifiers, lossless forwarding classes, DCBX, and PFC on ingress and egress interfaces to support lossless transport is as described in this document.

When the QFX Series acts as an FCoE-FC gateway, the system uses native FC interfaces (NP_Ports) to connect to the FC switch (or FCoE forwarder) at the FC network edge. You cannot apply CNPs or DCBX to native FC interfaces, only to Ethernet interfaces.

On an FCoE-FC gateway, the Ethernet interface configuration of classifiers, DCBX, and PFC is the same as the Ethernet interface configuration on a transit switch. The configuration of lossless forwarding classes is also the same.

However, supporting lossless transport on native FC interfaces requires that you rewrite the IEEE 802.1p priority value *if* your network uses any priority other than 3 (IEEE code point 011) for FCoE traffic. If your network uses priority 3 for FCoE traffic, you can and should use the default configuration on native FC interfaces.

By default, native FC interfaces tag packets with priority 3 when they encapsulate the incoming FC packets in Ethernet. If your FCoE network uses a different priority than 3 for FCoE traffic, you need to rewrite the priority value to the value that your network uses on the FC interface, classify the FCoE traffic to the correct priority on the Ethernet interfaces, and enable PFC on the correct priority on the Ethernet interfaces, as described in [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#).

Configuration Results and Commit Checks

Different configurations of forwarding classes and their drop attributes, classifiers, CNPs (PFC flow control), and Ethernet PAUSE (IEEE 802.3X flow control) result in different system behaviors.

[Table 491 on page 5520](#) describes the results of the possible lossless transport configurations in each case. The assumption in the *Result* column is that the system's buffer headroom calculation resulted in a successful configuration.

However, if the system calculates that there is insufficient buffer space to support the configuration, a commit check prevents you from committing the configuration on an individual Ethernet interface. For LAG interfaces, the system does not issue a commit check error but instead issues a syslog message.



NOTE: After you configure lossless transport for a LAG interface, be sure to check the syslog messages to confirm that the commit was successful.

Table 491: Results of Lossless Priority Configuration

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
None (default classifier)	None	None	System default configuration. No flows are lossless. To achieve lossless behavior for the default fcoe and no-loss forwarding classes, you must configure a CNP to enable PFC on their IEEE 802.1p code points (011 and 100 respectively).
Classifier with no lossless forwarding classes	None	None	No lossless traffic flows are configured; all traffic is best effort.
Classifier with at least one lossless forwarding class	None	None	Because no CNP is attached to interfaces, PFC is not enabled on the code point of the lossless traffic and no headroom buffer is allocated to the lossless queue, so packets can drop during periods of congestion. This configuration does not achieve lossless behavior.

Table 491: Results of Lossless Priority Configuration (*continued*)

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
None (default classifier)	PFC enabled on the fcoe and no-loss forwarding class code points (priorities)	None	The default classifier classifies traffic into two lossless forwarding classes, fcoe and no-loss. The CNP enables PFC on the priorities mapped to both lossless forwarding classes, resulting in lossless behavior for traffic mapped to the fcoe and no-loss forwarding classes.
None (default classifier)	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length. The system does not calculate buffer headroom for individual output queues. Because Ethernet PAUSE is enabled on the link instead of PFC being enabled on the lossless priorities, the entire link is paused during periods of congestion. This configuration results in lossless behavior for all of the forwarding classes on the link, but because all traffic is paused, this can cause greater overall network congestion.
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	None	Headroom buffer allocated only to priorities that are mapped to the lossless forwarding classes and on which PFC is enabled. This configuration achieves lossless behavior for the lossless forwarding classes.
Classifier with no lossless forwarding classes	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	Flow control enabled on a <i>different</i> interface than the interface with the CNP	The system checks the available buffer space for both the PFC-enabled priorities and for the other link. If sufficient buffer space is available, the lossless forwarding classes configured with PFC on one interface and also all of the traffic on the link with Ethernet PAUSE enabled achieve lossless behavior.



NOTE: If you attempt to configure both PFC and Ethernet PAUSE on a link, the system returns a commit error. PFC and Ethernet PAUSE are mutually exclusive configurations on an interface.

Backward Compatibility with Junos OS Releases Earlier Than Release 12.3

The addition of the no-loss packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the fcoe and no-loss forwarding classes.

If you used the default forwarding class configuration for the fcoe and no-loss forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the no-loss packet drop attribute.)

However, if you explicitly configured the fcoe or the no-loss forwarding class by including the **set forwarding-classes class forwarding-class-name queue-num queue-number** statement at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. (They are lossy because explicit configuration in releases earlier than Junos OS Release 12.3 did not use the no-loss packet drop attribute.) In Junos OS Release 12.3 and later, you must include the no-loss packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the no-loss packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the no-loss drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



NOTE: The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes were lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding classes retains its explicit configuration).



NOTE: To achieve lossless behavior for the traffic belonging to any forwarding class, you must also use a CNP to enable PFC on the IEEE 802.1p priority mapped to the forwarding class and apply the CNP to the relevant interfaces, and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

Configuration Rules and Recommendations

Keep in mind the following configuration rules and recommendations when you configure lossless traffic flows:

- You can configure a maximum of six lossless forwarding classes (forwarding classes with the no-loss packet drop attribute).
- All forwarding classes that you map to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes must be lossless).
- You cannot configure flow control to pause a multidestination output queue. You can configure PFC flow control only to pause unicast output queues.
- Forwarding classes mapped to multidestination queues (queues 8 through 11) cannot have the no-loss packet drop attribute. (Multidestination forwarding classes cannot be configured as lossless forwarding classes.)
- Do not configure weighted random early detection (WRED) on lossless forwarding classes. (Do not associate a drop profile with a forwarding class that has the no-loss packet drop attribute.)

Related Documentation

- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS Buffer Configuration on page 5527](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)

- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

When the QFX Series acts as an FCoE-FC gateway, it connects an Ethernet network that carries Fibre Channel over Ethernet (FCoE) traffic to a Fibre Channel (FC) network. Ethernet interfaces connect to the FCoE network. Native FC interfaces (NP_Ports) connect to the FC network.

FCoE traffic typically uses IEEE 802.1p priority 3 (code point 011). The QFX Series default configuration maps priority 3 traffic to the FCoE forwarding class. If your FCoE network uses priority 3 for FCoE traffic, you do not need to remap priorities, because the default configuration maps priority 3 to the FCoE forwarding class. (But you do need to enable PFC on IEEE 802.1p code point 3 on the Ethernet interfaces to achieve lossless behavior.)

However, if the FCoE network uses a different IEEE 802.1p priority than priority 3 for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that priority (and classified to that priority on the FCoE Ethernet interfaces in the ingress classifier). You specify the lossless forwarding class used for the FCoE traffic by configuring a fixed classifier and applying it to the native FC (NP_Port) interface. All traffic received from the FC SAN on that NP_Port interface is classified into the forwarding class specified in the fixed classifier.

When native FC interfaces on the FCoE-FC gateway encapsulate incoming FC traffic in Ethernet to create FCoE frames, by default they assign IEEE 802.1p code point 011 to the FCoE traffic, forward the traffic internally to the gateway Ethernet interfaces, and then forward the traffic to the FCoE network. Setting a rewrite value for the IEEE 802.1p code point configures the gateway native FC interface to assign the rewrite value priority to the FCoE frames when the native FC interface forwards the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

You can configure one rewrite value per native FC interface. Each native FC interface can use a different rewrite value.

- [Priority Remapping Configuration on page 5524](#)
- [Configuration Rules on page 5525](#)
- [Fate Sharing on page 5526](#)

Priority Remapping Configuration

Native FC interfaces on an FCoE-FC gateway receive native FC traffic from the FC SAN and encapsulate it in Ethernet to create FCoE frames. Priority remapping enables you to map the encapsulated FC traffic (the FCoE traffic) to any IEEE 802.1p priority. (This is similar to the rewrite rules you can configure to remap forwarding classes to code points

on Ethernet egress interfaces, but the rewrite takes place at the ingress FC interface so that the QFX Series uses the correct priority for FCoE traffic on the converged Ethernet network.)

To support lossless traffic flows, you must configure the remapped priority correctly on the native FC interfaces and also on the Ethernet interfaces that connect to the FCoE network. Achieving lossless behavior for FCoE traffic when you remap the FCoE priority requires configuring:

- A lossless forwarding class for FCoE traffic (or using the default *fcoe* forwarding class)
- A behavior aggregate (BA) classifier on the FCoE Ethernet interfaces to map the FCoE forwarding class to the IEEE 802.1p code points (priority) used for FCoE traffic on the FCoE network (the ingress classifier priority for the forwarding class must be the same as the rewrite value priority)
- A fixed classifier on the FCoE-FC gateway FC interface that maps all traffic from the FC network into the lossless FCoE forwarding class (the forwarding class must be lossless)
- A priority rewrite value that remaps the IEEE 802.1p code point on the FCoE-FC gateway FC interface to the priority used for FCoE traffic on the FCoE network
- An input congestion notification profile (CNP) to enable priority-based flow control (PFC) on the FCoE code point (the code point used as the rewrite value) at the Ethernet ingress interfaces

The ingress and egress configurations must match to achieve lossless behavior. The priority and the forwarding class specified in the BA classifier and in the CNP on the Ethernet ingress interfaces must match the fixed classifier and rewrite value on the FC interfaces. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

For example, if you configure a lossless forwarding class named *my_fcoe_fc* and your Ethernet network uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic, then:

- The forwarding class configuration, the BA classifier, and the fixed classifier all specify *my_fcoe_fc* as the forwarding class
- The BA classifier, the input CNP, and the rewrite value all specify the IEEE 802.1p code point 101

Configuration Rules

The following configuration rules apply when you remap priorities on an FCoE-FC gateway:

- Each native FC interface (NP_Port) supports one IEEE 802.1p priority value. The interface rewrites the IEEE 802.1p code point of all incoming traffic on the interface to the rewrite value. (The FC interface uses either the default value of 3 or the rewrite value for all incoming traffic.)
- Ports in the same FCoE-FC gateway local fc-fabric must use the same rewrite value. For example, if ports fc-0/0/0 and fc-0/0/1 are in the same local FCoE-FC gateway

fabric, they must use the same rewrite value. If you attempt to commit a configuration that uses different IEEE 802.1p priority rewrite values, the system returns a commit error.

- Ports in different FCoE-FC gateway local fc-fabrics can use different rewrite values. An example scenario is:
 - Interfaces fc-0/0/0 and fc-0/0/1 are in FCoE-FC gateway fc-fabric *my_fc_fab1*.
 - Interfaces fc-0/0/4 and fc-0/0/5 are in FCoE-FC gateway fc-fabric *my_fc_fab2*.

In this scenario, interfaces fc-0/0/0 and fc-0/0/1 must use the same rewrite value because they belong to the same local FC fabric on the gateway. Interfaces fc-0/0/4 and fc-0/0/5 also must use the same rewrite value because they belong to the same local FC fabric. However, the rewrite value you use for interfaces fc-0/0/0 and fc-0/0/1 can be different than the rewrite value you use for interfaces fc-0/0/4 and fc-0/0/5 because the interfaces belong to different local FC fabrics.

- You can apply the rewrite value only to native FC interfaces; you cannot apply the rewrite value configuration to Ethernet interfaces.
- The forwarding class specified in the fixed classifier on the native FC interface must be a lossless forwarding class. You cannot apply a fixed classifier to a native FC interface unless the associated forwarding class is lossless. (The forwarding class must be one of the two default lossless forwarding classes, or you must explicitly configure the forwarding class with the *no-loss* drop attribute.)
- The lossless forwarding class and IEEE 802.1p priority configuration must match on the FCoE-FC gateway native FC interfaces and Ethernet interfaces:
 - The same IEEE 802.1p priority (code point) must be paused on the Ethernet ingress interfaces, classified to the lossless forwarding class used in the native FC interface fixed classifier, and set as the rewrite value on the native FC interfaces.
 - The same lossless forwarding class must be used in the fixed classifier on the native FC interfaces and in the classifier configuration on the Ethernet interfaces.

Fate Sharing

To ensure that congestion on one interface does not affect the fate of traffic on a native FC interface on which you remap priorities, avoid fate sharing (different traffic flows receiving the same CoS treatment) configurations.

You can avoid fate sharing by ensuring that the remapping priority (code point) on the native FC interface is classified only to the forwarding class used in the fixed classifier on all other interfaces. For example, if you configure a fixed classifier on an FC interface that classifies all of the traffic into lossless forwarding class *myfcoe1* and remaps the priority to priority 5 (IEEE 802.1p code point 101), then in all other classifier configurations on all other interfaces, priority 5 should always be classified to forwarding class *myfcoe1*. If you classify priority 6 on another interface to forwarding class *myfcoe1*, then congestion on priority 6 traffic affects priority 5 traffic unfairly.

Related Documentation

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP_Ports\) on page 5801](#)

Understanding CoS Buffer Configuration

Each QFX3500 and QFX3600 switch has 9 MB of Packet Forwarding Engine (PFE) wide common packet buffer memory that is used to store packets on interface queues. The buffer memory has separate ingress and egress accounting to make accept, drop, or pause decisions. Because the switch has a single pool of memory with separate ingress and egress accounting, the full amount of buffer memory is available from both the ingress and the egress perspective. Packets are accounted for as they enter and leave the switch, but there is no concept of a packet arriving at an ingress buffer and then being moved to an egress buffer.

The buffers are divided into two pools from both an ingress and an egress perspective:

1. *Shared buffers* are a global memory pool that the switch allocates dynamically to ports as needed, so the buffers are shared among the switch ports.
2. *Dedicated buffers* are a memory pool divided equally among the switch ports. Each port receives a minimum guaranteed amount of buffer space, dedicated to each port, not shared among ports.



NOTE: Lossless traffic is traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link enabled for Ethernet PAUSE (IEEE 802.3x).

The switch reserves nonconfigurable buffer space to ensure that ports and queues receive a minimum memory allocation. You can configure how the system uses the rest of the buffer space to optimize the allocation for your mix of network traffic. You can configure the percentage of available buffer space used as shared buffer space versus dedicated buffer space. You can also configure how shared buffer space is allocated to different types of traffic. You can optimize the buffer settings for the traffic on your network.

The default buffer configuration is designed for networks that have a balance of best-effort and lossless traffic.

The default class-of-service configuration provides two lossless forwarding classes (**fcoe** and **no-loss**), a best-effort unicast forwarding class, a network control traffic forwarding class, and one multidestination (multicast, broadcast, and destination lookup fail) forwarding class. Each default forwarding class maps to a different default output queue. The default configuration allocates the buffers in a manner that supports a moderate amount of lossless traffic while still providing the ability to absorb bursts in best-effort traffic transmission.

Changing the buffer settings changes the abilities of the buffers to absorb traffic bursts and handle lossless traffic. For example, networks with mostly best-effort traffic require allocating most of the shared buffer space to best-effort buffers. This provides deep, flexible buffers that can absorb traffic bursts with minimal packet loss, at the expense of buffer availability for lossless traffic.

Conversely, networks with mostly lossless traffic require allocating most of the shared buffer space to lossless headroom buffers. This prevents packet loss on lossless flows at the expense of absorbing bursty best-effort traffic efficiently.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

This topic describes the buffer architecture and settings:

- [Buffer Pools on page 5528](#)
- [Default Buffer Pool Values on page 5536](#)
- [Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios on page 5538](#)
- [Optimizing the Buffer Configuration on page 5542](#)

Buffer Pools

From both an ingress and an egress perspective, the 9-MB buffer is split into two main pools, a shared buffer pool and a dedicated buffer pool that ensures a minimum allocation to each port. You can configure the amount of buffer space allocated to each of the two pools. A portion of the buffer space is reserved so that there is always a minimum amount of shared and dedicated buffer space available to each port.

- **Shared buffer pool**—A global memory space that all of the ports on the switch share dynamically as they need buffers. The shared buffer pool is further partitioned into buffers for best-effort unicast, best-effort multdestination (broadcast, multicast, and destination lookup fail), and PFC (lossless) traffic types. You can allocate global shared memory space to buffer partitions to better support different mixes of network traffic. The larger the shared buffer pool, the better the switch can absorb traffic bursts because more shared memory is available for the traffic.
- **Dedicated buffer pool**—A reserved global memory space allocated equally to each port. The switch reserves a minimum dedicated buffer pool that is not user-configurable. You can divide the dedicated buffer allocation for a port among the port queues on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic.)

A larger dedicated buffer pool means a larger amount of dedicated buffer space for each port, so congestion on one port is less likely to affect traffic on another port because the traffic does not need to use as much shared buffer space. However, the larger the dedicated buffer pool, the less bursty traffic the switch can handle because there is less dynamic shared buffer memory.

You can configure the way the available unreserved portion of the buffer space is allocated to the global shared buffer pool and to the dedicated shared buffer pool by configuring the ingress and egress shared buffer percentages.

By default, 100 percent of the available unreserved buffer space is allocated to the shared buffer pool. If you change the percentage of space allocated to the shared buffer, the available buffer space that is not allocated to the shared buffer is allocated to the dedicated buffer. For example, if you configure the ingress shared buffer pool as 80 percent, the remaining 20 percent of the available buffer space is allocated to the dedicated buffer pool and divided equally across the ports.



NOTE: When 100 percent of the available (user-configurable) buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

You can separately configure ingress and egress shared buffer pool allocations. You can also partition the ingress and egress shared buffer pool to allocate percentages of the shared buffer pool to specific types of traffic. If you do not use the default configuration or one of the recommended configurations, pay particular attention to the ingress configuration of the lossless and lossless headroom buffers (these buffers handle PFC pause during periods of congestion) and to the egress configuration of the best-effort buffers to handle incast congestion (multiple synchronized sources sending data to the same receiver in parallel).

In addition to the shared buffer pool and the dedicated buffer pool, there is also a small ingress global headroom buffer pool that is reserved and is not configurable.

When contention for buffer space occurs, the switch uses an internal algorithm to ensure that the buffer pools are distributed fairly among competing flows. When traffic for a given flow exceeds the amount of dedicated port buffer reserved for that flow, the flow begins to consume memory from the dynamic shared buffer pool. Competing flows compete for shared buffer memory with other flows that also have exhausted their dedicated buffers. When there is no congestion, there are no competing flows.

- [Buffer Handling of Lossless Flows \(PFC\) Versus Ethernet PAUSE on page 5529](#)
- [Shared Buffer Pool and Partitions on page 5530](#)
- [Dedicated Port Buffer Pool and Buffer Allocation to Queues on page 5531](#)
- [Trade-off Between Shared Buffer Space and Dedicated Buffer Space on page 5535](#)
- [Order of Buffer Consumption on page 5536](#)

Buffer Handling of Lossless Flows (PFC) Versus Ethernet PAUSE

When we discuss lossless buffers in the following sections, we mean buffers that handle traffic on which you enable PFC to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x). The lossless ingress and egress shared buffers, and the ingress lossless headroom shared buffer, are used only for traffic on which you enable PFC.



NOTE: To support lossless flows, you must configure the appropriate data center bridging capabilities (PFC, DCBX, or ETS) and scheduling properties.

Shared Buffer Pool and Partitions

The shared buffer pool is a global memory space that all of the ports on the switch share dynamically as they need buffers. The switch uses the shared buffer pool to absorb traffic bursts after the dedicated buffer pool for a port is exhausted.

You can divide both the ingress shared buffer pool and the egress shared buffer pool into three partitions to allocate percentages of each buffer pool to different types of traffic. When you partition the ingress or egress shared buffer pool:

- If you explicitly configure one ingress shared buffer partition, you must explicitly configure all three ingress shared buffer partitions. (You either explicitly configure all three ingress partitions or you use the default setting for all three ingress partitions.)

If you explicitly configure one egress shared buffer partition, you must explicitly configure all three egress shared buffer partitions. (You either explicitly configure all three egress partitions or you use the default setting for all three egress partitions.)

The switch returns a commit error if you do not explicitly configure all three partitions when configuring the ingress or egress shared buffer partitions.

- The combined percentages of the three ingress shared buffer partitions must total exactly 100 percent.

The combined percentages of the three egress shared buffer partitions must total exactly 100 percent.

When you explicitly configure ingress or egress shared buffer partitions, the switch returns a commit error if the total percentage of the three partitions does not equal 100 percent.

- If you explicitly partition one set of shared buffers, you do not have to explicitly partition the other set of shared buffers. For example, you can explicitly configure the ingress shared buffer partitions and use the default egress shared buffer partitions. However, if you change the buffer partitions for the ingress buffer pool to match the expected types of traffic flows, you would probably also want to change the buffer partitions for the egress buffer pool to match those traffic flows.

You can configure the percentage of available unreserved buffer space allocated to the shared buffer pool. Space that you do not allocate to the shared buffer pool is added to the dedicated buffer pool and divided equally among the ports. The default configuration allocates 100 percent of the unreserved ingress and egress buffer space to the shared buffers.

Configuring the ingress and egress shared buffer pool partitions enables you to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic.

Ingress Shared Buffer Pool Partitions

You can configure three ingress buffer pool partitions:

- Lossless buffers—Shared buffer pool for all lossless ingress traffic. The recommended minimum value for lossless buffers is 5 percent.
- Lossless headroom buffers—Shared buffer pool for packets received while a pause is asserted. If PFC is enabled on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers for which the recommended value can be less than 5 percent.)
- Lossy buffers—Shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The recommended minimum value for best-effort buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and best-effort buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. If you explicitly configure an ingress shared buffer partition, you must explicitly configure all three ingress buffer partitions, even if the lossless headroom buffer partition has a value of 0 (zero) percent.

Egress Shared Buffer Pool Partitions

You can configure three egress buffer pool partitions:

- Lossless buffers—Shared buffer pool for all lossless egress queues. The recommended minimum value for lossless buffers is 5 percent.
- Lossy buffers—Shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The recommended minimum value for best-effort buffers is 5 percent.
- Multicast buffers—Shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The recommended minimum value for multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and should have a value of at least 5 percent. If you explicitly configure an egress shared buffer partition, you must explicitly configure all three egress buffer partitions, and each partition should have a value of at least 5 percent.

Dedicated Port Buffer Pool and Buffer Allocation to Queues

The global dedicated buffer pool is memory that is allocated equally to each port, so each port receives a guaranteed minimum amount of buffer space. Dedicated buffers

are not shared among ports. Each port receives an equal proportion of the dedicated buffer pool.

The amount of dedicated buffer space is not user-configurable and depends on the percentage of available nonreserved buffers allocated to the shared buffers. (The dedicated buffer space is equal to the minimum reserved port buffers plus the remainder of the available nonreserved buffers that are not allocated to the shared buffer pool.)

When traffic enters and exits the switch, the switch ports use their dedicated buffers to store packets. If the dedicated buffers are not sufficient to handle the traffic, the switch uses shared buffers. The only way to increase the dedicated buffer pool is to decrease the shared buffer pool from its default value of 100 percent of available unreserved buffers.



NOTE: If 100 percent of the available unreserved buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

The larger the shared buffer pool, the better the burst absorption across the ports. The larger the dedicated buffer pool, the larger the amount of dedicated buffer space for each port. The greater the dedicated buffer space, the less likely that congestion on one port can affect traffic on another port, because the traffic does not need to use as much shared buffer space.

Allocating Dedicated Port Buffers to Queues

You can divide the dedicated buffer allocation for an egress port among the port queues by including the **buffer-size** statement in the scheduler configuration. This enables you to control the egress port dedicated buffer allocation on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic, or to stop the port from reserving buffers for queues that do not carry traffic.) Egress dedicated port buffer allocation is a hierarchical structure that allocates a global dedicated buffer pool evenly among ports, and then divides the allocation for each port among the port queues.

By default, ports divide their allocation of dedicated buffers among their egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (the **transmit-rate** option) for traffic. Only the queues included in the default scheduler receive bandwidth and dedicated buffers, in the proportions shown in [Table 492 on page 5532](#):

Table 492: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler)

Forwarding Class	Queue	Minimum Guaranteed Bandwidth (transmit-rate)	Proportion of Reserved Dedicated Port Buffers
best-effort	0	5%	5%
fcoe	3	35%	35%

Table 492: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler) (*continued*)

Forwarding Class	Queue	Minimum Guaranteed Bandwidth (transmit-rate)	Proportion of Reserved Dedicated Port Buffers
no-loss	4	35%	35%
network-control	7	5%	5%
mcast	8	20%	20%

In the default configuration, no egress queues other than the ones shown in [Table 492 on page 5532](#) receive an allocation of dedicated port buffers.



NOTE: The switch uses hierarchical scheduling to control port and queue bandwidth allocation, as described in “[Understanding CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 5481 and shown in “[Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 5606. For egress queue buffer size configuration, when you attach a traffic control profile (includes the queue scheduler information) to a port, the dedicated egress buffers on the port are divided among the queues as configured in the scheduler.

If you do not want to use the default allocation of dedicated port buffers to queues, use the **buffer-size** option in the scheduler that is attached to the port to configure the queue allocation. You can configure the dedicated buffer allocation to queues in two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is attached to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining dedicated port buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler for a queue means no dedicated buffer allocation for that queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



NOTE: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

Configuring Dedicated Port Buffer Allocation to Queues

In a port configuration that includes multiple forwarding class sets, with multiple forwarding classes mapped to multiple schedulers, the allocation of port dedicated buffers to queues depends on the mix of queues with buffer sizes configured as explicit percentages and queues configured with (or defaulted to) the **remainder** option.

The best way to demonstrate how using the percentage and remainder options affects dedicated port buffer allocation to queues is by showing an example of queue buffer allocation, and then showing how the queue buffer allocation changes when you add another forwarding class (queue) to the port.

[Table 493 on page 5534](#) shows an initial configuration that includes four forwarding class sets, the five default forwarding classes (mapped to the five default queues for those forwarding classes), the **buffer-size** option configuration, and the resulting buffer allocation for each queue. [Table 494 on page 5535](#) shows the same configuration after we add another forwarding class (best-effort-2, mapped to queue 1) to the best-effort forwarding class set. Comparing the buffer allocations in each table shows you how adding another queue affects buffer allocation when you use remainders and explicit percentages to configure the buffer allocation for different queues.

Table 493: Egress Queue Dedicated Buffer Allocation (Example 1)

Forwarding Class Set (Priority Group)	Forwarding Class	Queue	Scheduler Buffer Size Configuration	Buffer Allocation per Queue (Percentage)
fc-set-be	best-effort	0	10%	10%
fc-set-lossless	fcoe	3	20%	20%
	no-loss	4	40%	40%
fc-set-strict-high	network-control	7	remainder	15%
fc-set-mcast	mcast	8	remainder	15%

In this first example, 70 percent of the egress port dedicated buffer pool is explicitly allocated to the best-effort, fcoe, and no-loss queues. The remaining 30 percent of the port dedicated buffer pool is split between the two queues that use the **remainder** option (network-control and mcast), so each queue receives 15 percent of the dedicated buffer pool.

Now we add another forwarding class (queue) to the best-effort priority group (fc-set-be) and configure it with a buffer size of *remainder* instead of configuring a specific percentage. Because a third queue now shares the remaining dedicated buffers, the queues that share the remainder receive fewer dedicated buffers, as shown in [Table 494 on page 5535](#). The queues with explicitly configured percentages receive the configured percentage of dedicated buffers.

Table 494: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2)

Priority Group (fc-set)	Forwarding Class	Queue	Scheduler Buffer Size Configuration	Buffer Allocation per Queue (Percentage)
fc-set-be	best-effort	0	10%	10%
	best-effort-2	1	remainder	10%
fc-set-lossless	fcoe	3	20%	20%
	no-loss	4	40%	40%
fc-set-strict-high	network-control	7	remainder	10%
fc-set-mcast	mcast	8	remainder	10%

The two tables show how the port divides the dedicated buffer space that remains after servicing the queues that have an explicitly configured percentage of dedicated buffer space.

Trade-off Between Shared Buffer Space and Dedicated Buffer Space

The trade-off between shared buffer space and dedicated buffer space is:

- Shared buffers provide better absorption of traffic bursts because there is a larger pool of dynamic buffers that ports can use as needed to handle the bursts. However, all flows that exhaust their dedicated buffer space compete for the shared buffer pool. A larger shared buffer pool means a smaller dedicated buffer pool, and therefore more competition for the shared buffer pool because more flows exhaust their dedicated buffer allocation. Too much shared buffer space results in no single flow receiving very much shared buffer space, to maintain fairness when many flows contend for that space.
- Dedicated buffers provide guaranteed buffer space to each port. The larger the dedicated buffer pool, the less likely that congestion on one port affects traffic on another port, because the traffic does not need to use as much shared buffer space. However, less shared buffer space means less ability to dynamically absorb traffic bursts.

For optimal burst absorption, the switch needs enough dedicated buffer space to avoid persistent competition for the shared buffer space. When fewer flows compete for the shared buffers, the flows that need shared buffer space to absorb bursts receive more of the shared buffer because fewer flows exhaust their dedicated buffer space.

The default configuration and all of the recommended configurations allocate 100 percent of the user-configurable memory space to the global shared buffer pool because the amount of space reserved for dedicated buffers provides enough space to avoid persistent competition for dynamic shared buffers. This results in fewer flows competing for the shared buffers, so the competing flows receive more of the buffer space.

Order of Buffer Consumption

The total buffer pool is divided into ingress and egress shared buffer pools and dedicated buffer pools. When traffic flows through the switch, the buffer space is used in a particular order that depends on the type of traffic.

On ingress, the order of buffer consumption is:

- Best-effort unicast traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Global headroom buffers (very small)
- Lossless unicast traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Lossless headroom buffers
 4. Global headroom buffers (very small)
- Multidestination traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Global headroom buffers (very small)

On egress, the order of buffer consumption is the same for unicast best-effort, lossless unicast, and multidestination traffic:

- Dedicated buffers
- Shared buffers

In all cases on all ports, the switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for the port or queue is exhausted. This reserves the maximum amount of dynamic shared buffer space to absorb traffic bursts.

Default Buffer Pool Values

You can view the default or configured ingress and egress buffer pool values in KB units using the **show class-of-service shared-buffer** operational command. You can view the configured shared buffer pool values in percent units using the **show configuration class-of-service shared-buffer** operational command.

This section provides the default total buffer, shared buffer, and dedicated buffer values for QFX3500 and QFX3600 switches.

- [Total Buffer Pool Size on page 5537](#)
- [Shared Buffer Pool Default Values on page 5537](#)
- [Dedicated Buffer Pool Default Values on page 5538](#)

Total Buffer Pool Size

The total buffer pool is common memory that has separate ingress and egress accounting, so the full buffer pool is available from both the ingress and egress perspective. The total buffer pool size is not configurable.

For ingress and egress buffer pools, the total buffer pool size is 9 MB (9360 KB).

Shared Buffer Pool Default Values

This section describes the default values in percent and in KB for the shared ingress and shared egress buffers.

- [Shared Ingress Buffer Default Values on page 5537](#)
- [Shared Egress Buffer Default Values on page 5537](#)

Shared Ingress Buffer Default Values

[Table 495 on page 5537](#) shows the default ingress shared buffer allocation values in KB units for QFX3500 and QFX3600 switches.

Table 495: Default Shared Ingress Buffer Values (KB)

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
7202 KB	648.18 KB	3240.9 KB	3312.92 KB

[Table 496 on page 5537](#) shows the default ingress shared buffer allocation values as percentages for QFX3500 and QFX3600. (If you change the default shared buffer allocation, you configure the change as a percentage.)

Table 496: Default Shared Ingress Buffer Values (Percentage)

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

Shared Egress Buffer Default Values

[Table 497 on page 5538](#) shows the default egress shared buffer allocation values in KB units.

Table 497: Default Shared Egress Buffer Values (KB)

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
6656 KB	3328 KB	2063.36 KB	1264.64 KB

[Table 498 on page 5538](#) shows the default egress shared buffer allocation values as percentages.

Table 498: Default Shared Egress Buffer Values (Percentage)

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

Dedicated Buffer Pool Default Values

The system reserves ingress and egress dedicated buffer pools that are divided equally among the switch ports. By default, the system allocates 100 percent of the available unreserved buffer space to the shared buffer pool. If you reduce the percentage of available unreserved buffer space allocated to the shared buffer pool, the remaining unreserved buffer space is added to the dedicated buffer pool allocation. You configure the amount of dedicated buffer pool space by reducing (or increasing) the percentage of buffer space allocated to the shared buffer pool. You do not directly configure the dedicated buffer pool allocation.

The default dedicated buffer pool values for QFX3500 and QFX3600 switches in KB units are:

- Ingress dedicated buffer—2158 KB
- Egress dedicated buffer—2704.0 KB

Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios

The way you configure the shared buffer pool depends on the mix of traffic on your network. This section provides shared buffer configuration recommendations for five basic network traffic scenarios:

- **Balanced traffic**—The network carries a balanced mix of unicast best-effort, lossless, and multicast traffic. (This is the default configuration.)
- **Best-effort unicast traffic**—The network carries mostly unicast best-effort traffic.
- **Best-effort traffic with Ethernet PAUSE (IEEE 802.3X) enabled**—The network carries mostly best-effort traffic with Ethernet PAUSE enabled on the links.
- **Best-effort multicast traffic**—The network carries mostly multicast best-effort traffic.
- **Lossless traffic**—The network carries mostly lossless traffic (traffic on which PFC is enabled).



NOTE: Lossless traffic is defined as traffic on which you enable PFC to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link on which you enable Ethernet PAUSE. Start with the recommended profiles for each network traffic scenario, and adjust them if necessary for your network traffic conditions.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete. This includes changing the default configuration to one of the recommended configurations.

Start with the recommended shared buffer configuration for your network traffic conditions:

- [Balanced Traffic \(Default Configuration\) on page 5539](#)
- [Best-Effort Unicast Traffic on page 5540](#)
- [Ethernet PAUSE Traffic on page 5540](#)
- [Best-Effort Multicast \(Multidestination\) Traffic on page 5541](#)
- [Lossless Traffic on page 5541](#)

Balanced Traffic (Default Configuration)

The default shared buffer configuration is optimized for networks that carry a balanced mix of best-effort unicast, lossless, and multidestination (multicast, broadcast, and destination lookup fail) traffic. The default class-of-service (CoS) configuration is also optimized for networks that carry a balanced mix of traffic.

We recommend that you use the default shared buffer configuration for networks that carry a balanced mix of traffic, especially if you are using the default CoS settings. [Table 499 on page 5539](#) shows the default ingress shared buffer allocations:

Table 499: Default Ingress Shared Buffer Configuration

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

[Table 500 on page 5539](#) shows the default egress shared buffer allocations:

Table 500: Default Egress Shared Buffer Configuration

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

Best-Effort Unicast Traffic

If your network carries mostly best-effort (lossy) unicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 501 on page 5540](#)) and egress shared buffer settings (see [Table 502 on page 5540](#)):

Table 501: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	5%	0%	95%

Table 502: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	75%	20%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic](#)” on page 5748 for an example that shows you how to configure the recommended buffer settings shown in [Table 501 on page 5540](#) and [Table 502 on page 5540](#).

Ethernet PAUSE Traffic

If your network carries mostly best-effort (lossy) traffic *and* enables Ethernet PAUSE on links, then the default shared buffer configuration allocates too much buffer space to the shared ingress buffer (Ethernet PAUSE traffic uses the dedicated buffers instead of shared buffers) and not enough space to the lossless-headroom buffers. We recommend that you use the following ingress shared buffer settings (see [Table 503 on page 5540](#)) and egress shared buffer settings (see [Table 504 on page 5541](#)):

Table 503: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
70%	5%	80%	15%

Table 504: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	75%	20%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled](#)” on page 5754 for an example that shows you how to configure the recommended buffer settings shown in [Table 501 on page 5540](#) and [Table 502 on page 5540](#).

Best-Effort Multicast (Multidestination) Traffic

If your network carries mostly best-effort (lossy) multicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 505 on page 5541](#)) and egress shared buffer settings (see [Table 506 on page 5541](#)):

Table 505: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best -Effort Multicast Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	5%	0%	95%

Table 506: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	20%	75%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic](#)” on page 5759 for an example that shows you how to configure the recommended buffer settings shown in [Table 505 on page 5541](#) and [Table 506 on page 5541](#).

Lossless Traffic

If your network carries mostly lossless traffic, then the default shared buffer configuration allocates too much buffer space to support best-effort traffic. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 507 on page 5542](#)) and egress shared buffer settings (see [Table 508 on page 5542](#)):

Table 507: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	15%	80%	5%

Table 508: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	90%	5%	5%

See [“Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic” on page 5764](#) for an example that shows you how to configure the recommended buffer settings shown in [Table 507 on page 5542](#) and [Table 508 on page 5542](#).

Optimizing the Buffer Configuration

Instead of using the default buffer configuration or one of the recommended buffer configurations, you can optimize the buffer configuration to best support the mix of traffic on your network:

- [Optimizing Buffer Configuration on page 5542](#)
- [General Rules and Considerations on page 5543](#)

Optimizing Buffer Configuration

If you feel that the recommended shared buffer configurations are not optimized enough for your network, adjust the settings gradually to fine-tune the shared buffer allocation. Use caution when adjusting the shared buffer configuration, not just when you fine-tune the ingress and egress buffer partitions, but also when you fine-tune the total ingress and egress shared buffer percentage. (Remember that if you allocate less than 100 percent of the available buffers to the shared buffers, the remaining buffers are added to the dedicated buffers). Tuning the buffers incorrectly can cause problems such as ingress port congestion.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

The relationship between the sizes of the ingress buffer pool and the egress buffer pool affects when and where packets are dropped. The buffer pool sizes include the shared buffers and the dedicated buffers. In general, if there are more ingress buffers than egress buffers, the switch can experience ingress port congestion because egress queues fill before ingress queues can empty.

Use the `show class-of-service shared-buffer` operational command to see the sizes in kilobytes (KB) of the dedicated and shared buffers and of the shared buffer partitions.

For best-effort traffic (unicast and multdestination), the combined ingress lossy shared buffer partition and ingress dedicated buffers must be *less than* the combined egress lossy and multicast shared buffer partitions plus the egress dedicated buffers. This prevents ingress port congestion by ensuring that egress best-effort buffers are deeper than ingress best-effort buffers, and ensures that if packets are dropped, they are dropped at the egress queues. (Packets dropping at the ingress prevents the egress schedulers from working properly.)

For lossless traffic (traffic on which you enable PFC), the combined ingress lossless shared buffer partition and a reasonable portion of the ingress headroom buffer partition, plus the dedicated buffers, must be *less than* the total egress lossless shared buffer partition and dedicated buffers. (A reasonable portion of the ingress headroom buffer is approximately 20 to 25 percent of the buffer space, but this varies depending on how much buffer headroom is required to support the lossless traffic.) When these conditions are met, if there is ingress port congestion, the ingress port congestion triggers PFC on the ingress port to prevent packet loss. If the total lossless ingress buffers exceed the total lossless egress buffers, packets could be dropped at the egress instead of PFC being applied at the ingress to prevent packet loss.



NOTE: If you commit a buffer configuration for which the switch does not have sufficient resources, the switch might log an error instead of returning a commit error. After you commit a buffer configuration, check the syslog messages to ensure that the new buffer configuration did not fail to commit.

If the buffer configuration commits but you receive a syslog message that indicates the configuration cannot be implemented, you can:

- Reconfigure the buffers or reconfigure other parameters (for example, the PFC configuration, which affects the need for lossless headroom buffers and lossless buffers—the more priorities you pause, the more lossless and lossless headroom buffer space you need), then attempt the commit operation again.
- Roll back the switch to the last successful configuration.

If you receive a syslog message that says the buffer configuration cannot be implemented, you must take corrective action. If you do not fix the configuration or roll back to a previous successful configuration, the system behavior is unpredictable.

General Rules and Considerations

Keep the following rules and considerations in mind when you configure the buffers:

- Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.
- If you configure the ingress or egress shared buffer percentages as less than 100 percent, the remaining percentage of buffer space is added to the dedicated buffer pool.
- The sum of all of the ingress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent except the lossless headroom buffer, which can have a value of 0 percent.
- The sum of all of the egress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent.
- Lossless and lossless headroom shared buffers serve traffic on which you enable PFC, and do not serve traffic subject to Ethernet PAUSE.
- The switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for a port or queue is exhausted.
- Too little dedicated buffer space results in too much competition for shared buffer space.
- Too much dedicated buffer space results in poorer burst absorption because there is less available shared buffer space.
- Always check the syslog messages after you commit a new buffer configuration.
- The optimal buffer configuration for your network depends on the types of traffic on the network. If your network carries less traffic of a certain type (for example, lossless traffic), then you can reduce the size of the buffers allocated to that type of traffic (for example, you can reduce the sizes of the lossless and lossless headroom buffers).

**Related
Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)

Understanding CoS Tail-Drop Profiles

When the number of packets queued is greater than the ability of the switch to empty an output queue, the queue requires a method for determining which packets to drop to relieve the congestion. Weighted random early detection (WRED) drop profiles define the drop probability of packets as the output queue fills. During periods of congestion, as the output queue fills, the switch drops incoming packets as determined by a drop profile until the output queue becomes less congested.

Depending on the drop probabilities, a drop profile can drop many packets long before the buffer becomes full, or it can drop only a few packets even if the buffer is almost full.

You configure drop profiles in the drop profile section of the class-of-service (CoS) configuration hierarchy. You apply drop profiles using a drop profile map in each scheduler configuration. For each scheduler, you can configure separate drop profiles for each combination of loss priority (low, medium-high, and high) and protocol.

Drop profiles define the meaning of each of the loss priorities by setting the values for when to drop packets and the probability that packets will drop.

You can configure a maximum of 32 drop profiles.



NOTE: You cannot apply drop profiles to multidestination (multicast) queues.

Do not apply drop profiles to lossless flows such as FCoE traffic, because the corresponding queues require lossless behavior. Use priority-based flow control (PFC) to prevent packet drop.

- [Drop Profile Parameters on page 5545](#)
- [Default Drop Profile on page 5546](#)
- [Packet Drop Method on page 5547](#)
- [Drop Profile Maps on page 5548](#)
- [Congestion Prevention on page 5548](#)

Drop Profile Parameters

Drop profiles specify two values:


- **Fill level**—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- **Drop probability**—The percentage value that corresponds to the likelihood that an individual packet is dropped.

You set two queue fill levels and two drop probabilities in each drop profile. The two fill levels and the two drop probabilities create two pairs of values. The first fill level and the first drop probability create one value pair and the second fill level and the second drop probability create the second value pair.

The first fill level value specifies the percentage of queue fullness at which packets begin to drop, known as the drop start point. Until the queue reaches this level of fullness, no packets are dropped. The second fill level value specifies the percentage of queue fullness at which all packets are dropped, known as the drop end point.

The first drop probability value is always 0 (zero). This pairs with the drop start point and specifies that until the queue fullness level reaches the first fill level, no packets drop. When the queue fullness exceeds the drop start point, packets begin to drop until the queue exceeds the second fill level, when all packets drop. The second drop probability value, known as the maximum drop rate, specifies the likelihood of dropping packets when the queue fullness reaches the drop end point. As the queue fills from the drop start point to the drop end point, packets drop in a smooth, linear pattern (called an interpolated graph) as shown in [Figure 209 on page 5546](#). After the drop end point, all packets drop.

Figure 209: Tail-Drop Profile Packet Drop



The thick line in [Figure 209 on page 5546](#) shows the packet drop characteristics for a sample tail-drop profile. At the drop start point, the queue reaches a fill level of 30 percent. At the drop end point, the queue fill level reaches 50 percent, and the maximum drop rate is 80 percent.

No packets drop until the queue fill level reaches the drop start point of 30 percent. When the queue reaches the 30 percent fill level, packets begin to drop. As the queue fills, the percentage of packets dropped increases in a linear fashion. When the queue fills to the drop end point of 50 percent, the rate of packet drop has increased to the maximum drop rate of 80 percent. When the queue fill level exceeds the drop end point of 50 percent, all of the packets drop until the queue fill level drops below 50 percent.

Default Drop Profile

If you do not configure default profiles and apply them to queue schedulers, the switch uses the default drop profile for lossy traffic classes. In the default drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent. As soon as packets arrive on a queue, the default profile might begin to drop packets.

Packet Drop Method

When a packet reaches the head of the queue, the switch generates a random number between 0 and 100. The switch plots the random number against the drop profile using the current fullness of the queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped.

To create the linear drop pattern from the drop start point to the drop end point, the drop probabilities are derived using a linear approximation with eight sections, or steps, from the minimum queue fill level to the maximum queue fill level. The fill levels are divided into the eight sections equally, starting at the minimum fill level and ending at the maximum fill level. As the queue fills, the percentage of dropped packets increases. The percentage of packets dropped is based on the maximum drop rate.

For example, the default drop profile (which specifies a maximum drop rate of 100 percent) has the following drop probabilities at each section, or step, in the eight-section linear drop pattern:

- First section—The minimum drop probability is 6.25 percent of the maximum drop rate. The maximum drop probability is 12.5 percent of the maximum drop rate.
- Second section—The minimum drop probability is 18.75 percent of the maximum drop rate. The maximum drop probability is 25 percent of the maximum drop rate.
- Third section—The minimum drop probability is 30.25 percent of the maximum drop rate. The maximum drop probability is 37.5 percent of the maximum drop rate.
- Fourth section—The minimum drop probability is 43.75 percent of the maximum drop rate. The maximum drop probability is 50 percent of the maximum drop rate.
- Fifth section—The minimum drop probability is 56.25 percent of the maximum drop rate. The maximum drop probability is 62 percent of the maximum drop rate.
- Sixth section—The minimum drop probability is 68.75 percent of the maximum drop rate. The maximum drop probability is 75.5 percent of the maximum drop rate.
- Seventh section—The minimum drop probability is 81.25 percent of the maximum drop rate. The maximum drop probability is 87.5 percent of the maximum drop rate.
- Eighth section—The minimum drop probability is 92.75 percent of the maximum drop rate. The maximum drop probability is 100 percent of the maximum drop rate.

Packets drop even when there is no congestion, because packet drops begin at the drop start point regardless of whether congestion exists on the port. The default drop profile example represents the worst-case scenario, because the drop start point fill level is 0 percent, so packet drop begins when the queue starts to receive packets.

You can specify when packets begin to drop by configuring a drop start point at a fill level greater than 0 percent. For example, if you configure a drop profile that has a drop start point of 30 percent, packets do not drop until the queue is 30 percent full.

The smaller the gap between the minimum drop rate (which is always 0) and the maximum drop rate, the smaller the gap between the minimum drop probability and the

maximum drop probability at each section (step) of the linear drop pattern. The default drop profile, which has the maximum gap between the minimum drop rate (0 percent) and the maximum drop rate (100 percent), has the highest gap between the minimum drop probability and the maximum drop probability at each step. Configuring a lower maximum drop rate for a drop profile reduces the gap between the minimum drop probability and the maximum drop probability.

Drop Profile Maps

Drop profile maps are part of scheduler configuration. A drop profile map maps a drop profile to a loss priority and a protocol. Specifying the drop profile map in a scheduler associates the drop profile with the queues (forwarding classes) that you map to the scheduler in a scheduler map.

You configure loss priority for a queue in the classifier section of the CoS configuration hierarchy, and the loss priority is applied to the queue at the ingress interface.

Each scheduler can have multiple drop profile maps, one for each combination of loss priority and protocol.

Congestion Prevention

Configuring drop profiles on output queues prevents them from impacting other queues on the egress ports. If you do not configure drop profiles and map them to output queues, output queues without drop profiles can impact output queues on other egress ports, even if those queues are not experiencing congestion.

For example, if an ingress port forwards traffic to more than one egress port, and at least one of the egress ports experiences congestion, that can cause ingress port congestion. Ingress port congestion (ingress buffer exceeds its resource allocation) can cause frames to drop at the ingress port instead of at the egress port. Ingress port frame drop affects all of the egress ports to which the congested ingress port forwards traffic, not just the congested egress port.



NOTE: Do not configure drop profiles for the `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Related Documentation

- [Understanding Junos CoS Components on page 5436](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)

Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. Rewrite rules set the value of the CoS bits within the header of the outgoing packet. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header. Rewrite rules must be assigned to an interface for rewrites to be activated.

You can apply (bind) one DSCP or DSCP IPv6 rewrite rule and one IEEE 802.1p rewrite rule to each interface as described in [“Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces” on page 5460](#).

You cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface; each interface supports only one DSCP rewrite rule. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6.



NOTE: There are no default rewrite rules.

You can look at behavior aggregate (BA) classifiers and rewrite rules as two sides of the same coin. A BA classifier reads the CoS bits of incoming packets and classifies the packets into forwarding classes, then the system applies the CoS configured for the forwarding class to those packets. Rewrite rules rewrite the CoS bits just before the packets leave the system so that the next switch can apply the appropriate level of CoS to the packets. When you apply a rewrite rule to an interface, the rewrite rule is the last CoS action performed on the packet before it is forwarded.

Rewrite rules alter CoS values in outgoing packets on the outbound interfaces of an edge switch to accommodate the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.



NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

QFX Series does not have default rewrite rules. If you want to apply a rewrite rule to outgoing packets, you must explicitly configure the rewrite rule.



NOTE: Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

**Related
Documentation**

- [Understanding Junos CoS Components on page 5436](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)
- [Defining CoS Rewrite Rules on page 5805](#)

Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series supports the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 5551](#)
- [ETS on page 5552](#)
- [DCBX on page 5552](#)

Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 5551](#)
- [Buffer Management on page 5551](#)
- [Physical Interfaces on page 5551](#)

PFC

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

Buffer Management

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

Physical Interfaces

The switch supports 10-Gbps, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps (or faster) Ethernet interfaces.

ETS

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).
- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.
- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict priority flows.

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

**Related
Documentation**

- [Overview of Fibre Channel on the QFX Series on page 4956](#)
- [Understanding FCoE on page 4968](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding DCBX on page 5051](#)
- [Understanding Fibre Channel Terminology on page 5074](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

The QFX Series supports two methods of flow control:

- IEEE 802.3X Ethernet PAUSE
- IEEE 802.1Qbb priority-based flow control (PFC)

Ethernet PAUSE and PFC are link-level flow control mechanisms.

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority is mapped to a 3-bit IEEE 802.1p CoS code point flag in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on specified types of traffic (for example, Fibre Channel over Ethernet traffic).



NOTE: Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error. Ethernet PAUSE and PFC are mutually exclusive configurations on an interface.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

- [Ethernet PAUSE on page 5554](#)
- [PFC on page 5558](#)
- [Lossless Transport Support Summary on page 5562](#)

Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface

responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

Symmetric flow control means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

Asymmetric flow control allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, the PFC configuration overrides Ethernet PAUSE flow control.) The QFX Series supports both symmetric and asymmetric flow control.

- [Symmetric Flow Control on page 5555](#)
- [Asymmetric Flow Control on page 5556](#)

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can

perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

Asymmetric Flow Control

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 393 on page 5066](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

Table 509: Asymmetric Ethernet PAUSE Flow Control Configuration

Receive (Rx) Buffer	Transmit (Tx) Buffer	Configured Flow Control State
On	Off	Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).
Off	On	Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.)
On	On	Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.
Off	Off	Ethernet PAUSE flow control is disabled.

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, the QFX Series supports autonegotiation of Ethernet PAUSE with the connected peer. (The QFX Series does not support autonegotiation on 10-Gigabit Ethernet interfaces.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM_DIR bits, as described in [Table 394 on page 5067](#):

Table 510: Flow Control State Advertised to the Connected Peer (Autonegotiation)

Rx Buffer State	Tx Buffer State	PAUSE Bit	ASM_DIR Bit	Description
Off	Off	0	0	The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.
On	On	1	0	The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).
On	Off	0	1	The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).
Off	On	1	1	The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.)

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control behavior (resolution) between them, as shown in [Table 395 on page 5068](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series and on the connected peer (also known as the link partner). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



NOTE: In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

Table 511: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces

Local Interface (QFX Series)		Peer Interface		Local Resolution	Peer Resolution
PAUSE Bit	ASM_DIR Bit	PAUSE Bit	ASM_DIR Bit		
0	0	Don't care	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	1	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive	Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive
1	0	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	0	1	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive
1	1	0	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	1	0	1	Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive
1	1	Don't care	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive



NOTE: For your convenience, [Table 395 on page 5068](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits

a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic with the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as Fibre Channel over Ethernet (FCoE), LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of link-level flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)
- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 396 on page 5070](#):

Table 512: Default PFC Priority to Queue and Forwarding Class Mapping

IEEE 802.1p Priority (Code Point)	Queue	Forwarding Class
0 (000)	0	best-effort
1 (001)	1	best-effort
2 (010)	2	best-effort
3 (011)	3	fcoe
4 (100)	4	no-loss
5 (101)	5	best-effort
6 (110)	6	network-control
7 (111)	7	network-control

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default forwarding class configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) and [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 5524](#).

You enable PFC on a priority by:

1. Specifying the IEEE 802.1p code point to pause in the input stanza of a CNP
2. Applying the CNP to the ingress interfaces on which you want to pause the traffic



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
 3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

Although unicast traffic and multdestination (multicast, broadcast, and destination lookup fail) traffic must use different classifiers, you can map a unicast queue (queue 0 through 7) and a multdestination queue (queue 8, 9, 10, or 11) to the same PFC priority so that both unicast and multicast traffic use that priority. Do not map multdestination traffic to lossless priorities. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



NOTE: You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone QFX3500 switches and QFX3600 switches, you can create two CNPs with an explicitly configured output stanza.

When a QFX3500 switch or a QFX3600 switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 5506 describes configuring output flow control.

Lossless Transport Support Summary

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



NOTE: Junos OS Release 12.2 introduced changes to the way the QFX Series handles lossless forwarding classes (including the default fcoe and no-loss forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the fcoe and no-loss forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the fcoe or the no-loss forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the fcoe or the no-loss forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit fcoe and no-loss forwarding class configuration before you upgrade to Junos OS Release 12.2.

See [“Overview of CoS Changes Introduced in Junos OS Release 12.2” on page 5428](#) for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the fcoe and no-loss forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new no-loss packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5506](#) provides detailed information about the explicit configuration of lossless priorities and about the default configuration of lossless priorities, including the input and output stanzas of the CNP.



NOTE: PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

Related Documentation

- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)

- [Understanding DCB Features and Requirements on page 4965](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)

Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.

This topic describes:

- [DCBX Basics on page 5564](#)
- [DCBX Modes and Support on page 5565](#)
- [DCBX Attribute Types on page 5568](#)
- [DCBX Application Protocol TLV Exchange on page 5569](#)
- [DCBX and PFC on page 5570](#)
- [DCBX and ETS on page 5571](#)

DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



NOTE: When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

DCBX Modes and Support

This section describes DCBX support on the QFX Series:

- [DCBX Modes \(Versions\) on page 5565](#)
- [Autonegotiation on page 5567](#)
- [CNA Support for DCBX Modes on page 5568](#)
- [Interface Support for DCBX on page 5568](#)

DCBX Modes (Versions)

The QFX Series supports the two most common DCBX modes:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



NOTE: The QFX Series does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The QFX Series drops LLDP frames that contain pre-CEE DCBX TLVs.

Table 392 on page 5053 summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

Table 513: Summary of Differences Between IEEE DCBX and DCBX Version 1.01

Characteristic	IEEE DCBX	DCBX Version 1.01
OUI	0x0080c2	0x001b21
Frame Format	Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs	Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.
Symmetric/asymmetric configuration with peer	Asymmetric or symmetric	Symmetric only
Differences in the show dcbx interface interface-name operational command	<ul style="list-style-type: none"> • Synchronization information is not shown because symmetric configuration is not required. • Operational state information is not shown because the operational states do not have to be symmetric. • TLV type is shown because unique TLVs are sent for each DCBX attribute. • ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs. 	<ul style="list-style-type: none"> • Synchronization information is shown because symmetric configuration is required. • Operational state information is shown because the operational states do have to be symmetric. • TLV type is not shown because one TLV is used for all attribute information. • Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/public/dcs2008/az-walker-dcb-capability-exchange-discovery-protocol-108-v101.pdf>



NOTE: As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



NOTE: On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

Autonegotiation

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on QFX3500 switches and QFabric systems:

- QFX3500 switch—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface

receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



NOTE: If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

CNA Support for DCBX Modes

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on QFX Series interfaces depends on the DCBX features that the CNAs in your network support.

Interface Support for DCBX

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

DCBX Attribute Types

DCBX has three attribute types:

- **Informational**—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- **Asymmetric**—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- **Symmetric**—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 5568](#)
- [Symmetric Attributes on page 5569](#)

Asymmetric Attributes

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the

interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing”, the configuration on the interface cannot be overridden by the peer interface configuration.

- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

Symmetric Attributes

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

DCBX Application Protocol TLV Exchange

DCBX advertises the switch’s capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 5569](#)
- [FCoE Application Protocol TLV Exchange on page 5569](#)
- [Disabling Application Protocol TLV Exchange on page 5570](#)

Application Protocol TLV Exchange

For all applications, DCBX advertises the application’s state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application’s TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

FCoE Application Protocol TLV Exchange

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)

- Does *not* have an application map



NOTE: If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.



NOTE: If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

Disabling Application Protocol TLV Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

DCBX and PFC

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

DCBX and ETS

This section describes:

- [Default DCBX ETS Advertisement on page 5571](#)
- [ETS Advertisement and Peer Configuration on page 5571](#)
- [ETS Recommendation TLV on page 5572](#)

Default DCBX ETS Advertisement

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

ETS Advertisement and Peer Configuration

DCBX does not control the switch's ETS (hierarchical scheduling) operational state. If the connected peer is configured as "willing," DCBX pushes the switch's ETS configuration to the switch's peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

ETS Recommendation TLV

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



NOTE: You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

Related Documentation

- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding DCB Features and Requirements on page 4965](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding FCoE on page 4968](#)
- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)

Understanding DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



NOTE: LLDP and DCBX are enabled by default on all interfaces.

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 5060](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case (see [“Application Maps” on page 5061](#)) and only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic.

This topic describes:

- [Applications on page 5573](#)
- [Application Maps on page 5574](#)
- [Classifying and Prioritizing Application Traffic on page 5575](#)
- [Enabling Interfaces to Exchange Application Protocol Information on page 5576](#)
- [Disabling DCBX Application Protocol Exchange on page 5576](#)

Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise, except FCoE if FCoE is the only application that you want the interface to advertise.



NOTE: If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

Classifying and Prioritizing Application Traffic

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



NOTE: You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

Related Documentation

- [Understanding DCBX on page 5051](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)

- [Example: Configuring Unicast Classifiers on page 5658](#)

[QFabric-Specific CoS Overview](#)

- [Understanding CoS Fabric Forwarding Class Sets on page 5578](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)

Understanding CoS Fabric Forwarding Class Sets

Fabric forwarding class sets (fabric fc-sets) are similar to the fc-sets (priority groups) you configure on Node devices. The major differences are:

1. Fabric fc-sets group traffic for transport across the QFX3008-I or QFX3600-I Interconnect device (the fabric). Node device fc-sets group traffic on a Node device for transport across that Node device.
2. Fabric fc-sets are global. They apply to the entire fabric. Node device fc-sets apply only to the Node device on which they are configured.
3. You can configure class of service (CoS) scheduling for Node device fc-sets, but you cannot configure CoS for fabric fc-sets.
4. Fabric fc-sets map to Interconnect device fabric output queues statically—you cannot configure the mapping of fabric fc-sets to fabric output queues. All traffic in a fabric fc-set maps to the same output queue.

Node device fc-sets include forwarding classes that map to Node device output queues, and you can configure the mapping of forwarding classes to output queues (or you can use the default mapping). Because output queues are mapped to forwarding classes, different classes of traffic in a Node device fc-set can be mapped to different output queues.

Node device fc-sets consist of forwarding classes containing traffic that requires similar CoS treatment. (Forwarding classes are default forwarding classes or user-defined forwarding classes.) You can configure CoS for each fc-set to determine how the traffic of its forwarding classes is scheduled on a Node device.

When traffic exits a Node device interface and enters an Interconnect device fabric interface, the Interconnect device uses the same forwarding classes to group traffic. The forwarding classes are mapped to global fabric fc-sets for transport across the fabric. Like fc-sets on a Node device, fabric fc-sets also contain traffic that requires similar CoS treatment. Unlike fc-sets on a Node device, you cannot configure CoS on fabric fc-sets.

Fabric fc-sets reside on the Interconnect device and are global to the QFabric system. Fabric fc-sets apply to all traffic that traverses the fabric. The mapping of forwarding classes to fabric fc-sets is global and applies to all forwarding classes with traffic that traverses the fabric from all connected Node devices. You can change the mapping of forwarding classes to fabric fc-sets. All mapping changes you make are global. For example, if you change the fabric fc-set to forwarding class mapping of the default best-effort forwarding class, then every Node device's best-effort forwarding class traffic that traverses the fabric is mapped to that fabric fc-set.

This topic describes:

- [Default Fabric Forwarding Class Sets on page 5579](#)
- [Fabric Forwarding Class Set Configuration and Implementation on page 5582](#)
- [Fabric Forwarding Class Set Scheduling \(CoS\) on page 5584](#)
- [Support for Flow Control and Lossless Transport Across the Fabric on page 5586](#)

- [Viewing Fabric Forwarding Class Set Information on page 5588](#)
- [Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences on page 5590](#)

Default Fabric Forwarding Class Sets

Interconnect devices have 12 default fabric fc-sets, including five visible default fabric fc-sets, four for unicast traffic and one for multidestination (multicast, broadcast, and destination lookup failure) traffic.

There are also seven hidden default fabric fc-sets. There are three hidden default fabric fc-sets for multidestination traffic that you can use if you want to map different multidestination forwarding classes to different multidestination fabric fc-sets. There are four hidden default fabric fc-sets for lossless traffic that you can use to map different lossless forwarding classes (priorities) to different lossless fabric fc-sets.

[Table 117 on page 1239](#) shows the default fabric fc-sets:

Table 514: Default Fabric Forwarding Class Sets

Fabric Forwarding Class Set Name	Characteristics
<code>fabric_fcset_be</code>	Transports best-effort unicast traffic across the fabric.
<code>fabric_fcset_strict_high</code>	Transports unicast traffic that has been configured with strict-high priority and in the network-control forwarding class across the fabric. This fabric fc-set receives as much bandwidth across the fabric as it needs to service the traffic in the group up to the entire fabric interface bandwidth. For this reason, exercise caution when mapping traffic to this fabric fc-set to avoid starving other traffic.
<code>fabric_fcset_noloss1</code>	Transports unicast traffic in the default fcoe forwarding class across the fabric.
<code>fabric_fcset_noloss2</code>	Transports unicast traffic in the default no-loss forwarding class across the fabric.
<code>fabric_fcset_noloss3</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
<code>fabric_fcset_noloss4</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
<code>fabric_fcset_noloss5</code>	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.

Table 514: Default Fabric Forwarding Class Sets (*continued*)

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_noloss6	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_multicast1	Transports multdestination traffic in the mcast forwarding class across the fabric. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast2	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.

The five default forwarding classes (**best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast**) are mapped to the fabric fc-sets by default as shown in [Table 118 on page 1240](#).

Table 515: Default Forwarding Class to Fabric Forwarding Class Set Mapping

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
best-effort	fabric_fcset_be	0	NA
network-control	fabric_fcset_strict_high	7	NA
fcoe	fabric_fcset_noloss1	1	9K
no-loss	fabric_fcset_noloss2	2	9K
mcast	fabric_fcset_multicast1	8	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss3	3	9k

Table 515: Default Forwarding Class to Fabric Forwarding Class Set Mapping (*continued*)

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss4	4	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss5	5	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss6	6	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast2	9	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast3	10	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast4	11	NA

The maximum fiber cable length between the QFabric system Node device and the QFabric system Interconnect device is 150 meters.



TIP: If you explicitly configure lossless forwarding classes, we recommend that you map each user-configured lossless forwarding class to an unused fabric fc-set (fabric_fcset_noloss3 through fabric_fcset_noloss6) on a one-to-one basis: one lossless forwarding class mapped to one lossless fabric fc-set.

The reason for one-to-one mapping is to avoid fate sharing of lossless flows. Because each fabric fc-set is mapped statically to an output queue, when you map more than one forwarding class to a fabric fc-set, all of the traffic in all of the forwarding classes that belong to the fabric fc-set uses the same output queue. If that output queue becomes congested due to congestion caused by one of the flows, the other flows are also affected. (They share fate because the flow that congests the output queue affects flows that are not experiencing congestion.)

However, it is important to understand that fabric_fcset_noloss1 and fabric_fcset_noloss2 have a scheduling weight of 35, while the other fabric fc-sets have a scheduling weight of 1. The scheduling weights mean that fabric_fcset_noloss1 and fabric_fcset_noloss2 receive most of the bandwidth available to lossless fabric fc-sets if the amount of traffic on fabric_fcset_noloss1 and fabric_fcset_noloss2 requires the bandwidth.

If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will consume most of that bandwidth, then you should place all lossless traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2`. If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will **<emphasis>not</emphasis>** consume most of that bandwidth, then you can map lossless forwarding classes in a one-to-one manner to lossless fabric fc-sets to avoid fate sharing.

If you want to map different multidestination forwarding classes to different multidestination fabric fc-sets, use one or more of the hidden multidestination fabric fc-sets.



NOTE: The global mapping of forwarding classes to fabric fc-sets is independent of the mapping of forwarding classes to Node device fc-sets. Global mapping of forwarding classes to fabric fc-sets occurs only on the Interconnect device. The Node device mapping of forwarding classes to fc-sets does not affect the global mapping of forwarding classes to fabric fc-sets on the Interconnect device, and vice versa.

When you define new forwarding classes on a Node device, you explicitly map those forwarding classes to Node device fc-sets. However, new (user-created) forwarding classes are mapped by default to fabric fc-sets. (You can override the default mapping if you want to configure the forwarding class to fabric fc-set mapping explicitly, as described in the next section.)

By default:

- All best-effort traffic forwarding classes that you create are mapped to the **`fabric_fcset_be`** fabric fc-set.
- All lossless traffic forwarding classes that you create are mapped to the **`fabric_fcset_noloss1`** or **`fabric_fcset_noloss2`** fabric fc-set.
- All multidestination traffic forwarding classes that you create are mapped to the **`fabric_fcset_multicast1`** fabric fc-set.
- All **strict-high** priority traffic and **network-control** forwarding classes that you create are mapped to the **`fabric_fcset_strict_high`** fabric fc-set.

Fabric Forwarding Class Set Configuration and Implementation

You can map forwarding classes to fabric fc-sets, but no other attributes of fabric fc-sets are user-configurable, including CoS. This section describes:

- [Mapping Forwarding Classes to Fabric Forwarding Class Sets on page 5583](#)
- [Fabric Forwarding Class Set Implementation on page 5583](#)

Mapping Forwarding Classes to Fabric Forwarding Class Sets

If you do not want to use the default mapping of forwarding classes to fabric fc-sets, you can map forwarding classes to fabric fc-sets the same way as you map forwarding classes to Node device fc-sets. To do this, use exactly the same statement that you use to map forwarding classes to fc-sets, but instead of specifying a Node device fc-set name, specify a fabric fc-set name.



NOTE: The global mapping of forwarding classes to fabric fc-sets does not affect the mapping of forwarding classes to Node device fc-sets. The global forwarding class mapping to fabric fc-sets pertains to the traffic only when it enters, traverses, and exits the fabric. The forwarding class mapping to fc-sets on a Node device is valid within that Node device.

Mapping forwarding classes to fabric fc-sets does not affect the scheduling configuration of the forwarding classes or fc-sets on Node devices. Fabric fc-set scheduling (which is not user-configurable) pertains to traffic only when it enters, traverses, and exits the Interconnect device fabric.

If you change the mapping of a forwarding class to a fabric fc-set, the new mapping is global and applies to all traffic in that forwarding class, regardless of which Node device forwards the traffic to the Interconnect device.

- To assign one or more forwarding classes to a fabric fc-set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric-forwarding-class-set-name class
forwarding-class-name
```

For example, to map a user-defined forwarding class named **best-effort-2** to the fabric fc-set **fabric_fcset_be**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric_fcset_be class best-effort-2
```



NOTE: Because fabric fc-set configuration is global, in this example all forwarding classes with the name **best-effort-2** on all of the Node devices attached to the fabric use the **fabric_fcset_be** fabric fc-set to transport traffic across the fabric.

Fabric Forwarding Class Set Implementation

The following rules apply to fabric fc-sets:

- You cannot create new fabric fc-sets. Only the twelve default fabric fc-sets are available.
- You cannot delete a default fabric fc-set.
- You cannot attach a fabric fc-set to a Node device interface. Fabric fc-sets are used only on the Interconnect device fabric, not on Node devices.

- You can map only multdestination forwarding classes to multdestination fabric fc-sets.
- You cannot map multdestination forwarding classes to unicast fabric fc-sets.
- You cannot map unicast forwarding classes to multdestination fabric fc-sets.
- You cannot configure CoS for fabric fc-sets. (However, default CoS scheduling properties are applied to traffic on the fabric, and the fabric interfaces use link layer flow control (LLFC) for flow control.)

Fabric Forwarding Class Set Scheduling (CoS)

Although fabric fc-set CoS is not user-configurable, CoS is applied to traffic on the fabric. (In addition, fabric interfaces use LLFC to ensure lossless transport for lossless traffic flows.) This section describes how the fabric applies CoS scheduling to traffic:

- [Class Groups for Fabric Forwarding Class Sets on page 5584](#)
- [Class Group Scheduling on page 5584](#)
- [QFabric System CoS on page 5586](#)

Class Groups for Fabric Forwarding Class Sets

To transport traffic across the fabric, the QFabric system organizes the fabric fc-sets into three classes called *class groups*. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric_fcset_strict_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes and in any forwarding classes you create on a Node device that consist of **strict-high** priority or **network-control** forwarding class traffic.
- **Unicast**—All traffic in the fabric fc-sets **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes and in any forwarding classes you create on a Node device that consist of best-effort or lossless traffic. If you use any of the hidden no loss fabric fc-sets (**fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric_fcset_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric_fcset_multicast2**, **fabric_fcset_multicast3**, or **fabric_fcset_multicast4**), that traffic is also classified as part of this class group.

Class Group Scheduling

You cannot configure CoS for class groups or for fabric fc-sets (that is, you cannot attach a traffic control profile to a fabric fc-set—you attach traffic control profiles to Node device fc-sets to apply scheduling to the traffic that belongs to the Node device fc-set). By default, the fabric uses weighted round-robin (WRR) scheduling in which each class group receives a portion of the total available fabric bandwidth based on its type of traffic, as shown in [Table 119 on page 1245](#):

Table 516: Class Group Scheduling Properties and Membership

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Strict-high priority	fabric_fcset_strict_high	<ul style="list-style-type: none"> All strict-high priority forwarding classes network-control 	Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.
Unicast	<ul style="list-style-type: none"> fabric_fcset_be fabric_fcset_noloss1 fabric_fcset_noloss2 <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_noloss3 fabric_fcset_noloss4 fabric_fcset_noloss5 fabric_fcset_noloss6 	<ul style="list-style-type: none"> best-effort fcoe no-loss <p>NOTE: No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p>	Traffic in the unicast class group receives an 80% weight in the WRR calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	fabric_fcset_multicast1 <p>Includes the hidden multidestination fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_multicast2 fabric_fcset_multicast3 fabric_fcset_multicast4 	<ul style="list-style-type: none"> mcast <p>NOTE: No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)

The fabric fc-sets within each class group are weighted equally and receive bandwidth using round-robin scheduling. For example:

- If the unicast class group has three member fabric fc-sets, **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**, then each of the three fabric fc-sets receives one-third of the bandwidth available to the unicast class group.
- If the multidestination class group has one member fc-set, **fabric_fcset_multicast1**, then that fc-set receives all of the bandwidth available to the multidestination class group.
- If the multidestination class group has two member fc-sets, **fabric_fcset_multicast1** and **fabric_fcset_multicast2**, then each of the two fabric fc-sets receives one-half of the bandwidth available to the multidestination class group.

QFabric System CoS

When traffic enters and exits the same Node device, CoS works the same as it works on a standalone QFX3500 switch.

However, when traffic enters a Node device, crosses the Interconnect device, and then exits a different Node device, CoS is applied differently:

1. Traffic entering the ingress Node device receives the CoS configured at the Node ingress (packet classification, congestion notification profile for PFC).
2. When traffic goes from the ingress Node device to the Interconnect device, the fabric fc-set CoS is applied as described in the discussion of fabric forwarding class set scheduling.
3. When traffic goes from the Interconnect device to the egress Node device, the egress Node device applies CoS at the egress port (egress queue scheduling, WRED, IEEE 802.1p or DSCP code-point rewrite).

Support for Flow Control and Lossless Transport Across the Fabric

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the fabric interfaces use LLFC to support lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class on the Node devices.
2. The lossless forwarding class must be mapped to a lossless fabric fc-set on the Interconnect device (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**).

When traffic meets the two configuration constraints, the fabric propagates the back pressure from the egress Node device across the fabric to the ingress Node device during periods of congestion. However, to achieve end-to-end lossless transport across the switch, you must also configure a congestion notification profile to enable PFC on the Node device ingress ports.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.

- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



NOTE: Lossless transport across the fabric also must meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is a 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- **MTU size**—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- **Total number of Node device fabric ports connected to the Interconnect device**—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multideestination traffic.

Table 120 on page 1248 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 120 on page 1248, the QFabric system supports six lossless priorities.



NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 517: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5



NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

Viewing Fabric Forwarding Class Set Information

You can display information about fabric fc-sets using the same CLI command you use to display information about Node device fc-sets:

```
user@switch> show class-of-service forwarding-class-set
Forwarding class set: fabric_fcset_be, Type: fabric-type, Forwarding class set
index: 1
  Forwarding class      Index
  best-effort           0

Forwarding class set: fabric_fcset_mcast1, Type: fabric-type, Forwarding class
set index: 5
  Forwarding class      Index
  mcast                 8
```


Forwarding class set: fabric_fcset_mcast2, Type: fabric-type, Forwarding class set index: 6

Forwarding class set: fabric_fcset_mcast3, Type: fabric-type, Forwarding class set index: 7

Forwarding class set: fabric_fcset_mcast4, Type: fabric-type, Forwarding class set index: 8

Forwarding class set: fabric_fcset_noloss1, Type: fabric-type, Forwarding class set index: 2

Forwarding class	Index
fcoe	1

Forwarding class set: fabric_fcset_noloss2, Type: fabric-type, Forwarding class set index: 3

Forwarding class	Index
no-loss	2

Forwarding class set: fabric_fcset_noloss3, Type: fabric-type, Forwarding class set index: 9

Forwarding class set: fabric_fcset_noloss4, Type: fabric-type, Forwarding class set index: 10

Forwarding class set: fabric_fcset_noloss5, Type: fabric-type, Forwarding class set index: 11

Forwarding class set: fabric_fcset_noloss6, Type: fabric-type, Forwarding class set index: 12

Forwarding class set: fabric_fcset_strict_high, Type: fabric-type, Forwarding class set index: 4

Forwarding class	Index
network-control	3

[Table 121 on page 1249](#) describes the meaning of the **show class-of-service forwarding-class-set** output fields when you display fabric fc-set information.

Table 518: show class-of-service forwarding-class-set Command Output Fields

Field Name	Field Description
Forwarding class set	Name of the fabric forwarding class set.
Type	Type of forwarding class set: <ul style="list-style-type: none"> Fabric-type—Fabric fc-set Normal-type—Node device fc-set
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of the forwarding class.

Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences

Table 122 on page 1250 summarizes the differences between fabric fc-sets and Node device fc-sets:

Table 519: Summary of Differences Between Fabric fc-sets and Node Device fc-sets

Characteristic	Fabric fc-set	Node device fc-set
Location	QFX3008-I or QFX3600-I Interconnect device (the fabric).	QFabric system Node device.
Global or Node-device specific	Global, valid for the entire fabric.	Local to the Node device on which the fc-set is configured.
Ability to create (define) a new fc-set	No. Use the 12 default fabric fc-sets provided.	Yes.
Ability to configure CoS	Default CoS settings only. CoS is not user-configurable.	User-configurable using traffic control profiles.
Ability to map forwarding classes to an fc-set	Yes. Mapping is global and applies to all forwarding classes across Interconnect device fabric (traffic from all connected Node devices).	Yes. Mapping is local to a Node device and applies only to the forwarding classes on the Node device.

Related Documentation

- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [show class-of-service forwarding-class-set on page 5938](#)

Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports

Beginning with Junos OS Release 13.1, you can configure two-tier hierarchical scheduling (enhanced transmission selection, IEEE 802.1Qaz) on the fabric (fte) ports of QFabric system Node devices. Configuring CoS on Node device fabric interfaces provides increased control over traffic scheduling and helps to ensure predictable bandwidth consumption.

You can configure CoS on the following QFabric system interface types:

- Node device access interfaces (xe interfaces)—Schedule traffic on the output queues of the 10-Gigabit Ethernet access ports using standard Node device CoS scheduling configuration components, as described elsewhere in the QFX Series documentation. You can configure different scheduling for different ports and output queues.
- Node device fabric interfaces (fte interfaces)—Schedule traffic on the output queues of the 40-Gbps fabric interfaces that connect a Node device to a QFX3008-I or a QFX3600-I Interconnect device using standard Node device CoS scheduling configuration components. You can configure different scheduling for different interfaces and output queues.

This topic describes:

- [Hierarchical Scheduling Architecture on QFabric System Node Devices on page 5591](#)
- [Default Scheduling on Node Device Fabric Interfaces on page 5592](#)
- [Configuring Scheduling on Node Device Fabric Interfaces on page 5593](#)

Hierarchical Scheduling Architecture on QFabric System Node Devices

CoS architecture on Node device access interfaces is the same as CoS architecture on standalone switch access interfaces. CoS architecture on Node device fabric interfaces is also the same as the CoS architecture on the access interfaces. You apply schedulers to queues (priorities), fc-sets (priority groups), and interfaces in the same hierarchical manner as described in [“Understanding CoS Hierarchical Port Scheduling \(ETS\)” on page 5481](#).

You configure scheduling on Node device fabric interfaces (fte interfaces) using the same statements and configuration constructs that you use to configure scheduling on Node device access interfaces (xe interfaces). For example, on Node device fabric interfaces you can:

- Define up to four fc-sets (three unicast, one multidestination)



NOTE: If the fabric interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best-effort, lossless, strict-high priority, and multidestination traffic.

- Map forwarding classes to fc-sets

- Configure scheduling for each forwarding class (scheduler)
- Configure scheduling for each fc-set (traffic control profile)

The differences in configuring CoS on Node device fabric interfaces compared to configuring CoS on Node device access interfaces are:

- You specify a Node device *fabric* interface instead of a Node device *access* interface when you apply CoS to an interface.
- You cannot attach classifiers, congestion notification profiles, or rewrite rules to fabric interfaces. Also, you cannot configure buffer settings on fabric interfaces. You can only attach fc-sets and traffic control profiles.

Default Scheduling on Node Device Fabric Interfaces

Default scheduling on Node device fabric interfaces is the same as default scheduling on Node device access interfaces. Only the default forwarding classes (best-effort, network-control, fcoe, no-loss, and multidestination) receive port bandwidth, based on the default minimum guaranteed bandwidth (transmit rate) scheduler settings for each default forwarding class.

To transport traffic on Node device fabric interfaces, the system organizes the default forwarding classes into three *class groups*. Class groups are not user-configurable. The three class groups are:

- **Unicast**—All traffic in the default forwarding classes **best-effort**, **network-control**, **fcoe**, and **no-loss** belong to this default class group.
- **Multidestination**—All traffic in the default forwarding class **mcast** belongs to this default class group.
- **Strict-high priority**—There is no default strict-high priority forwarding class, so there is no default strict-high priority class group and there is no default configuration for strict-high priority traffic.



NOTE: If you configure strict-high priority forwarding classes, you must also configure an fc-set (priority group) for strict-high priority traffic, map the strict-high priority forwarding classes to the strict-high priority fc-set, create a scheduler for the strict-high priority traffic and map it to the strict-high priority forwarding classes, create a traffic control profile for the strict-high priority traffic, and apply the strict-high priority fc-set and traffic control profile to the appropriate fabric interfaces.

The default forwarding classes receive port bandwidth based on their default transmit rate settings (weights). Forwarding classes that are not default forwarding classes receive no default bandwidth.

Default class group scheduling uses weighted round-robin (WRR) scheduling, in which each class group receives a portion of the total available fabric interface bandwidth based on the class group traffic type, as shown in [Table 520 on page 5593](#). Within each

class group, the scheduler bandwidth allocation for individual forwarding classes is based on the default transmit rate for each forwarding class.

Table 520: Class Group Default Scheduling Properties and Membership on Node Device Fabric Interfaces

Class Group	Forwarding Class Mapping and Bandwidth Allocation (Default Transmit Rate)	Class Group Scheduling Properties (Weight)
Unicast	<ul style="list-style-type: none"> • best-effort (5%) • fcoe (35%) • no-loss (35%) • network-control (5%) 	Traffic in the unicast class group receives an 80% weight in the weighted round-robin (WRR) calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	<ul style="list-style-type: none"> • mcast (20%) 	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)



NOTE: Strict-high priority traffic is served first, before any other traffic is served. Strict-high priority traffic receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.

Configuring Scheduling on Node Device Fabric Interfaces

If you do not want to use default CoS scheduling on Node device fabric interfaces, you can configure two-tier hierarchical scheduling (ETS) the same way that you configure ETS on Node device access interfaces.

- [Similarities Between Node Device Fabric Interface and Access Interface Scheduling on page 5593](#)
- [Differences Between Node Device Fabric Interface and Access Interface Scheduling on page 5594](#)

Similarities Between Node Device Fabric Interface and Access Interface Scheduling

Configuring scheduling on a Node device fabric interface is similar to configuring scheduling on an access interface in many ways. In both cases, you configure:

- Schedulers to specify the output scheduling for forwarding class traffic
- Scheduler maps to map schedulers to forwarding classes
- Forwarding classes (or use the default forwarding classes)

- Forwarding class sets (groups of forwarding classes that require similar CoS treatment)
- A separate fc-set for strict-high priority traffic (an fc-set cannot contain a mix of strict-high priority traffic and traffic with a different priority)
- Traffic control profiles to specify the output scheduling for fc-sets
- Traffic control profile and fc-set mapping to interfaces

On Node device fabric interfaces, you configure ETS in the same way, and ETS works the same way as on Node device access interfaces

In addition, strict-high priority queues are served first, and then the remaining port bandwidth is allocated to other traffic.

Differences Between Node Device Fabric Interface and Access Interface Scheduling

Configuring scheduling on a Node device fabric interface differs from configuring scheduling on an access interface in several ways. On fabric interfaces:

- You cannot attach classifiers.
- You cannot attach congestion notification profiles (flow control is applied automatically to lossless forwarding classes).
- You cannot attach rewrite rules.
- You cannot configure buffer settings.
- You specify a Node device fabric interface name instead of a Node device access interface name when you apply CoS to an interface.

Related Documentation

- [Understanding CoS Fabric Forwarding Class Sets on page 1238](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
- [Understanding Default CoS Settings on page 5443](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)

Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)

The default class-of-service (CoS) properties on the QFabric system Interconnect device interfaces are optimized to best utilize the fabric resources. You cannot configure CoS properties on QFabric System Interconnect device interfaces.

- [Hierarchical CoS Architecture Across a QFabric System Interconnect Device on page 5595](#)
- [Default CoS on Interconnect Device Fabric Interfaces on page 5597](#)

Hierarchical CoS Architecture Across a QFabric System Interconnect Device

Because Interconnect devices support traffic from multiple Node devices that have multiple CoS configurations, CoS on Interconnect device fabric interfaces differs from CoS on Node device access and fabric interfaces.

The hierarchical CoS scheduling structure on the Interconnect device interfaces consists of two tiers:

1. Fabric forwarding class sets—Similar to fc-sets on Node devices, fabric fc-sets group traffic for transport across the Interconnect device fabric. Fabric fc-sets are global and apply to all traffic that crosses the fabric from all Node devices. See [“Understanding CoS Fabric Forwarding Class Sets” on page 1238](#) for a detailed description of fabric fc-sets.
2. Class groups—Fabric fc-sets are grouped into class groups for transport across the Interconnect device.

Node devices and Interconnect devices each have a two-tier hierarchical CoS scheduling architecture. The architectures are slightly different, but each tier of the scheduling hierarchy performs analogous functions, as shown in [Table 521 on page 5595](#).

Table 521: Hierarchical Scheduler Architecture on Node Devices and Interconnect Devices

Bandwidth Pool	Bandwidth Configuration on Node Devices	Bandwidth Configuration on Interconnect Devices
Port—Entire amount of bandwidth available to traffic on a port.	Access (xe) or fabric (fte) interfaces	Fabric (fte) or Clos fabric (bfte) interfaces
Priority group—Group of traffic types that requires similar CoS treatment. Each priority group receives a portion of the total available port bandwidth.	Forwarding class set (fc-set)	Class group
Priority—Most granular tier of bandwidth allocation. Each priority receives a portion of the total available priority group bandwidth.	Forwarding class (mapped to output queue)	Fabric fc-set (mapped to output queue)

Fabric FC-Sets

Fabric fc-sets are groups of forwarding classes that receive similar CoS treatment across the Interconnect device. Fabric fc-sets are global to the QFabric system and apply to all traffic that traverses the fabric, from all connected Node devices. The CoS on a fabric fc-set applies to all the traffic that belongs to that fabric fc-set.

For example, a fabric fc-set that includes the **best-effort** forwarding class handles all of the **best-effort** traffic from all of the connected Node devices that traverses the Interconnect device fabric.

There are 12 default fabric fc-sets, including 5 visible fabric fc-sets and 7 hidden fabric fc-sets. The five visible fabric fc-sets have forwarding classes mapped to them by default. By default, the seven hidden fabric fc-sets do not carry traffic, but you can map forwarding classes to the hidden fabric fc-sets if you want to use them.

You can configure the forwarding class membership of each fabric fc-set. However, you cannot create new fabric fc-sets, and you cannot delete the 12 default fabric fc-sets.

Each fabric fc-set is mapped to an output queue. Each fabric interface has 12 output queues, one for each of the 12 fabric fc-sets. The traffic from all of the forwarding classes mapped to a fabric fc-set uses that fabric fc-set's output queue.

Fabric fc-sets are grouped into class groups for transport across the Interconnect device.

Class Groups for Fabric FC-Sets

To transport traffic across the fabric, the fabric organizes the fabric fc-sets into three classes called *class groups*. Class groups are not user-configurable. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric_fcset_strict_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes, and in any forwarding classes you create on a Node device that consist of **strict-high** priority traffic.
- **Unicast**—All traffic in the fabric fc-sets **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes, and the traffic in any forwarding classes you create on a Node device that consist of best-effort or lossless traffic. If you use any of the hidden no loss fabric fc-sets (**fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric_fcset_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric_fcset_multicast2**, **fabric_fcset_multicast3**, or **fabric_fcset_multicast4**), that traffic is part of this class group.

Default CoS on Interconnect Device Fabric Interfaces

The Interconnect device interfaces use the default CoS configuration as described in these sections:

- [Default Class Group Scheduling on page 5597](#)
- [Default Fabric FC-Set Scheduling on page 5598](#)
- [Default Class Group and Fabric FC-Set Scheduling Example on page 5600](#)
- [Default PFC and Lossless Transport Across the Interconnect Device on page 5602](#)

Default Class Group Scheduling

Default class group bandwidth scheduling is analogous to default fc-set (priority group) scheduling on a Node device. Default class group scheduling uses weighted round-robin (WRR) scheduling, in which each class group receives a portion of the total available fabric interface bandwidth, based on the class group's traffic type, as shown in [Table 119 on page 1245](#):

Table 522: Class Group Default Scheduling Properties and Membership

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Strict-high priority	fabric_fcset_strict_high	<ul style="list-style-type: none"> • All strict-high priority forwarding classes • network-control 	Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.
Unicast	<ul style="list-style-type: none"> • fabric_fcset_be • fabric_fcset_noloss1 • fabric_fcset_noloss2 <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> • fabric_fcset_noloss3 • fabric_fcset_noloss4 • fabric_fcset_noloss5 • fabric_fcset_noloss6 	<ul style="list-style-type: none"> • best-effort • fcoe • no-loss <p>NOTE: No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p>	Traffic in the unicast class group receives an 80% weight in the weighted round-robin (WRR) calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)

Table 522: Class Group Default Scheduling Properties and Membership (*continued*)

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Multidestination	fabric_fcset_multicast1 Includes the hidden multidestination fabric fc-sets if used: <ul style="list-style-type: none"> • fabric_fcset_multicast2 • fabric_fcset_multicast3 • fabric_fcset_multicast4 	<ul style="list-style-type: none"> • mcast <p>NOTE: No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)

Only the five visible fabric fc-sets have traffic mapped to them by default. The fabric fc-sets within each class group are weighted by their transmit rates (guaranteed minimum bandwidth), and they receive bandwidth from the class group's total bandwidth using weighted round-robin (WRR) scheduling.

Default Fabric FC-Set Scheduling

Default fabric fc-set bandwidth scheduling is analogous to default forwarding class (priority) scheduling on a Node device. Each fabric fc-set receives a guaranteed minimum percentage of the port bandwidth that the class group receives. The guaranteed minimum percentage is called the *transmit rate*.

Table 523 on page 5598 shows the default transmit rate for each of the default fabric fc-sets.

Table 523: Default Fabric FC-Set Scheduler Configuration

Default Fabric FC-Set	Transmit Rate (Percentage of Class Group Bandwidth)
fabric_fcset_strict_high	N/A
	Strict-high priority traffic is served first, before any other traffic is served. Strict-high priority traffic receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.
fabric_fcset_noloss1	35%
fabric_fcset_noloss2	35%
fabric_fcset_be	10%
fabric_fcset_multicast1	20%

Each fabric fc-set belongs to a class group. Each class group receives a portion of the total available port bandwidth. Each fabric fc-set in a class group receives a portion of

the total available class group bandwidth based on the transmit rate (weight) of the fabric fc-set.

Traffic in `fabric_fcset_strict_high` does not have a default transmit rate because `fabric_fcset_strict_high` receives all of the bandwidth needed to empty its queue before other queues are served. Traffic in the remaining fabric fc-sets receive bandwidth in a ratio proportional to the default transmit rate of each fabric fc-set.

Each of the following hidden fabric fc-sets receives a default scheduling weight of 1:

- `fabric_fcset_noloss3`
- `fabric_fcset_noloss4`
- `fabric_fcset_noloss5`
- `fabric_fcset_noloss6`
- `fabric_fcset_multicast2`
- `fabric_fcset_multicast3`
- `fabric_fcset_multicast4`

You must explicitly map forwarding classes to hidden fabric fc-sets if you want to use the hidden fabric fc-sets.



CAUTION: Bandwidth is allocated to fabric fc-sets based on scheduling weight. The scheduling weights of the visible (default) fabric fc-sets are the same as their transmit rates, so in the unicast class group, `fabric_fcset_noloss1` and `fabric_fcset_noloss2` each have a weight of 35 and `fabric_fcset_be` has a weight of 10. In the multdestination class group, the default `fabric_fcset_multicast1` has a weight of 20. The hidden multicast and noloss fabric fc-sets each have a scheduling weight of 1.

The scheduling weights mean that when the visible fabric fc-sets are fully utilizing their allocated bandwidth:

- The hidden noloss fc-sets (`fabric_fcset_noloss3`, `fabric_fcset_noloss4`, `fabric_fcset_noloss5`, and `fabric_fcset_noloss6`) receive bandwidth at a proportional rate of 1:35 compared to the default noloss fc-sets.
- The hidden multicast fc-sets (`fabric_fcset_multicast2`, `fabric_fcset_multicast3`, and `fabric_fcset_multicast4`) receive bandwidth at a proportional rate of 1:20 compared to the default multicast fc-sets.

If you map traffic to a hidden fabric fc-set, that fabric fc-set receives the proportional amount of class group bandwidth that corresponds to its scheduling weight (1). The amount of bandwidth allocated to a hidden fabric fc-set depends on how much bandwidth the other fc-sets in the same class group consume. When the visible fabric fc-sets fully utilize their bandwidth, hidden fabric fc-sets receive only their minimum weight in bandwidth.

(However, even a low scheduling weight results in a relatively large absolute bandwidth allocation because each fabric port is a 40-Gbps port.)

For example, if `fabric_fcset_noloss1` and `fabric_fcset_noloss2` each consume all of the 35 percent of bandwidth allocated to them, and `fabric_fcset_be` consumes all of the 10 percent of bandwidth allocated to it, then `fabric_fcset_noloss3`, `fabric_fcset_noloss4`, `fabric_fcset_noloss5`, and `fabric_fcset_noloss6` receive bandwidth at a rate of 1:80 compared to the visible noloss fabric fc-sets. (If the visible fabric fc-sets do not use all of their allocated bandwidth, then the hidden fabric fc-sets receive more bandwidth.)

Another example is if we map lossless traffic to `fabric_fcset_noloss3` and to `fabric_fcset_noloss4`. `Fabric_fcset_noloss1` uses 10 percent of its 35 percent allocation of unicast class group bandwidth. `Fabric_fcset_noloss2` uses 15 percent of its 35 percent allocation of unicast class group bandwidth. `Fabric_fcset_be` uses 5 percent of its allocated bandwidth. `Fabric_fcset_noloss3` and `fabric_fcset_noloss4` can use the remaining unicast class group bandwidth allocated to lossless traffic. However, if the traffic on `fabric_fcset_noloss1`, `fabric_fcset_noloss2`, or `fabric_fcset_be` increases, the bandwidth allocated to the hidden fabric fc-sets decreases.

Similarly, if you map traffic to a hidden multidestination fabric fc-set (`fabric_fcset_multicast2`, `fabric_fcset_multicast3`, `fabric_fcset_multicast4`), that multidestination fabric fc-set receives the proportional amount of class group bandwidth that corresponds to its scheduling weight (1). The amount of bandwidth allocated to a hidden multidestination fabric fc-set depends on how much bandwidth the other fc-sets in the multidestination class group consume. When `fabric_fcset_multicast1` (the visible fabric fc-set) fully utilizes its bandwidth, hidden fabric fc-sets receive only their minimum weight in bandwidth. For example, if `fabric_fcset_multicast1` uses its full bandwidth allocation, then the hidden multidestination fabric fc-sets receive bandwidth at a rate of 1:20 compared to `fabric_fcset_multicast1`.

Default Class Group and Fabric FC-Set Scheduling Example

The following example shows how default scheduling allocates the total port bandwidth among the class groups and their fabric fc-sets. In the example, traffic is mapped to each of the forwarding classes in the five visible fabric fc-sets, and the strict-high priority class group consumes an average of 10 percent of the 40-Gbps fabric interface bandwidth (4 gigabits), leaving 90 percent of the fabric interface bandwidth (36 gigabits) for the remaining class groups.

In this scenario, by default, the strict-high priority class group includes one fabric fc-set (`fabric_fcset_strict_high`), the unicast class group includes three fabric fc-sets (`fabric_fcset_be`, `fabric_fcset_noloss1`, and `fabric_fcset_noloss2`), and the multidestination class group includes one fabric fc-set (`fabric_fcset_multicast1`). Each individual fabric fc-set receives the following treatment:

- Strict-high priority class group (`fabric_fcset_strict_high`)—This group is assumed to average 10 percent (4 gigabits) for the purposes of this example. Because the strict-high priority class group is served first and receives all of the bandwidth it requires to empty its queue, in real networks the amount of required bandwidth fluctuates and affects the amount of bandwidth available to the other class groups.



TIP: To prevent strict-high priority traffic from using too much bandwidth, you can set a maximum bandwidth limit by configuring a scheduler shaping rate for the `fabric_fcset_strict_high` fabric fc-set.

- Unicast class group (`fabric_fcset_be`, `fabric_fcset_noloss1`, and `fabric_fcset_noloss2`)—Each of these fabric fc-sets receives a weighted portion of the 80 percent of the total port bandwidth available after the strict-high traffic has been served. The weight corresponds to the transmit rate of each fabric fc-set. The following calculations show the minimum port bandwidth allocated to each of the unicast class group fabric fc-sets:

- `fabric_fcset_be`

$10 / (35 + 35 + 10)\%$ of 80% of the available port bandwidth (12.5 percent of 80 percent of port bandwidth)

The 10 that is the numerator in $10 / (35 + 35 + 10)$ is the percentage of bandwidth allocated to the `fabric_fcset_be` by the transmit rate weight. The $(35 + 35 + 10)$ in the denominator sums the percentage of bandwidth (transmit rate weights) allocated to each of the three fabric fc-sets in the unicast class group.

The 80 percent represents 80 percent of the port bandwidth available after strict-high priority traffic is served (36 gigabits).

The resulting equation is:

$10 / (35 + 35 + 10)\% \times (0.8 \times 36 \text{ gigabits}) = \text{approximately } 3.6 \text{ gigabits}$

- `fabric_fcset_noloss1` and `fabric_fcset_noloss2`

The default minimum bandwidth for the two visible lossless fabric fc-sets is the same because both of these fabric fc-sets have the same transmit rate weight.

$35 / (35 + 35 + 10)\%$ of 80% of the port bandwidth (43.75 percent of 80 percent of port bandwidth)

The 35 that is the numerator in $35 / (35 + 35 + 10)$ is the percentage of bandwidth allocated to each of the noloss fabric fc-sets by the transmit rate weight. The $(35 + 35 + 10)$ in the denominator sums the percentage of bandwidth (transmit rate weights) allocated to each of the three fabric fc-sets in the unicast class group.

The 80 percent represents 80 percent of the port bandwidth available after strict-high priority traffic is served (36 gigabits).

The resulting equation is:

$35 / (35 + 35 + 10)\% \times (0.8 \times 36 \text{ gigabits}) = \text{approximately } 12.6 \text{ gigabits}$

- Multidestination class group (`fabric_fcset_multicast1`)—Because only one fabric fc-set is configured by default in the multidestination class group, it receives 100 percent of the 20 percent of the total port bandwidth available to the multidestination class group after the strict-high traffic has been served:

100 / (100)% of 20% of the available port bandwidth (100 percent of 20 percent of available port bandwidth)

The resulting equation is:

100 / 100% x (0.2 x 36 gigabits) = approximately 7.2 gigabits

Default PFC and Lossless Transport Across the Interconnect Device

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the Interconnect device fabric supports lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class (a forwarding class configured with the **no-loss** parameter or the default **fcoe** or **no-loss** forwarding class).
2. The lossless forwarding class must be mapped to one of the lossless fabric fc-sets (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**). If you do not explicitly map lossless forwarding classes to fabric fc-sets, lossless forwarding classes are mapped by default to lossless fabric fc-sets **fabric_fcset_noloss1** and **fabric_fcset_noloss2**.

When traffic meets these two constraints, the fabric propagates back-pressure from egress queues during periods of congestion. However, to achieve end-to-end lossless transport across the QFabric system, you must also configure a congestion notification profile to enable PFC on the Node device ingress interfaces. To achieve end-to-end lossless transport across the network, you must configure PFC on all of the devices in the lossless traffic path.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the default congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop, and the traffic does not receive lossless treatment.
- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.

- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop, and the traffic does not receive lossless treatment.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the lossy **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the lossy **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



NOTE: Lossless transport across the fabric must also meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- **MTU size**—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- **Total number of Node device fabric ports connected to the Interconnect device**—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multdestination traffic.

Table 120 on page 1248 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 120 on page 1248, the QFabric system supports six lossless priorities.



NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 524: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5



NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

Related Documentation

- [Understanding CoS Fabric Forwarding Class Sets on page 1238](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)

CHAPTER 64

Configuration

- [Configuration Examples on page 5605](#)
- [Configuration Tasks on page 5781](#)
- [Configuration Statements on page 5815](#)

Configuration Examples

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5627](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5635](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Queue Scheduling Priority on page 5679](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5770](#)

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Hierarchical port scheduling defines the class-of-service (CoS) properties of output queues, which are mapped to forwarding classes (forwarding classes are mapped to IEEE 802.1p priorities, so mapping queues to forwarding classes also maps queues to priorities). Hierarchical port scheduling enables you to group priorities that require similar CoS resources into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.

Hierarchical port scheduling is the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz). One major benefit of hierarchical port scheduling is greater port bandwidth utilization. If a priority group on a port does not use all of its allocated bandwidth, other priority groups on that port can use that bandwidth. Also, if a priority within a priority group does not use its allocated bandwidth, other priorities within that priority group can use that bandwidth.

Configuring hierarchical scheduling is a multistep procedure that includes:

- Mapping forwarding classes to queues
- Defining forwarding class sets (priority groups)
- Defining behavior aggregate classifiers
- Configuring priority-based flow control (PFC) for lossless priorities (queues)
- Applying classifiers and PFC configuration to ingress interfaces
- Defining drop profiles
- Defining schedulers
- Mapping forwarding classes to schedulers
- Defining traffic control profiles
- Assigning priority groups and traffic control profiles to egress ports

This example describes how to configure hierarchical scheduling:

- [Requirements on page 5607](#)
- [Overview on page 5607](#)

- [Configuration on page 5610](#)
- [Verification on page 5618](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Keep the following considerations in mind when you plan the port bandwidth allocation for priority groups and for individual priorities:

- How much traffic and what types of traffic you expect to traverse the system.
- How you want to divide different types of traffic into priorities (forwarding classes, also called queues) to apply different CoS treatment to the traffic. Dividing traffic into priorities includes:
 - Mapping the code points of ingress traffic to forwarding classes using behavior aggregate (BA) classifiers. This classifies incoming traffic into the appropriate forwarding class.
 - Mapping forwarding classes to output queues. This defines the output queue for each type of traffic.
 - Attaching the BA classifier to the desired ingress interfaces so that incoming traffic maps to the desired forwarding classes and queues.
- How you want to organize priorities into priority groups (forwarding class sets).

Traffic that requires similar treatment usually belongs in the same priority group. To do this, place forwarding classes that require similar bandwidth, loss, and other characteristics in the same forwarding class set. For example, you can map all types of best-effort traffic forwarding classes into one forwarding class set.

- How much of the port bandwidth you want to allocate to each priority group and to each of the priorities in each priority group. The following considerations apply to bandwidth allocation:
 - Estimate how much traffic you expect in each forwarding class (output queue) and how much traffic you expect in each forwarding class set (the aggregate amount of traffic in the forwarding classes that belong to the forwarding class set).
 - The combined minimum guaranteed bandwidth of the priorities (forwarding classes) in a priority group should not exceed the minimum guaranteed bandwidth of the priority group. The transmit rate scheduler parameter defines the minimum guaranteed bandwidth for forwarding classes. Scheduler maps associate schedulers with forwarding classes.
 - The combined minimum guaranteed bandwidth of the priority groups (forwarding class sets) on a port should not exceed the port's total bandwidth. Traffic control profiles define the minimum bandwidth for a forwarding class set. Associating a

scheduler map with a traffic control profile sets the scheduling for the individual forwarding classes in the forwarding class set.

This example creates hierarchical port scheduling by defining priority groups for best effort, guaranteed delivery, and high-performance computing (HPC) traffic. Each priority group includes priorities that need to receive similar CoS treatment. Each priority group and each priority within each priority group receive the CoS resources needed to service their flows. Lossless priorities use PFC to prevent packet loss when the network experiences congestion.

Topology

Table 525 on page 5608 shows the configuration components for this example.

Table 525: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology

Property	Settings
Hardware	QFX3500 switch
Mapping of forwarding classes (priorities) to queues	<p>best-effort to queue 0</p> <p>be to queue 1</p> <p>fcoe (Fibre Channel over Ethernet) to queue 3</p> <p>no-loss to queue 4</p> <p>hpc (high-performance computing) to queue 5</p> <p>network-control to queue 7</p> <p>NOTE: If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does <i>not</i> receive lossless treatment.</p> <p>In Junos OS Release 12.3 and later, you can include the <i>no-loss</i> packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.</p>
Forwarding class sets (priority groups)	<p>best-effort-pg: contains forwarding classes best-effort, be, and network control</p> <p>guar-delivery-pg: contains forwarding classes fcoe and no-loss</p> <p>hpc-pg: contains forwarding class hpc</p>
Behavior aggregate classifier (maps forwarding classes and loss priorities to incoming packets by IEEE 802.1 code point)	<p>Name—hsclassifier1</p> <p>Code point mapping:</p> <ul style="list-style-type: none"> • 000 to forwarding class best-effort and loss priority low • 001 to forwarding class be and loss priority high • 011 to forwarding class fcoe and loss priority low • 100 to forwarding class no-loss and loss priority low • 101 to forwarding class hpc and loss priority low • 110 to forwarding class network-control and loss priority low

Table 525: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology (*continued*)

Property	Settings
PFC	Congestion notification profile name— gd-cnp PFC enabled on code points: 011 (fcoe priority), 010 (no-loss priority)
Drop profiles	dp-be-low : drop start point 25 , drop end point 50 , maximum drop rate 80 dp-be-high : drop start point 10 , drop end point 40 , maximum drop rate 100 dp-hpc : drop start point 75 , drop end point 90 , maximum drop rate 75 dp-nc : drop start point 80 , drop end point 100 , maximum drop rate 100
Queue schedulers	be-sched : minimum bandwidth 3g , maximum bandwidth 100% , priority low , drop profiles dp-be-low and dp-be-high fcoe-sched : minimum bandwidth 2.5g , maximum bandwidth 100% , priority low hpc-sched : minimum bandwidth 2g , maximum bandwidth 100% , priority low , drop profile dp-hpc nc-sched : minimum bandwidth 500m , maximum bandwidth 100% , priority low , drop profile dp-nc nl-sched : minimum bandwidth 2g , maximum bandwidth 100% , priority low
Forwarding class-to-scheduler mapping	Scheduler map be-map : Forwarding class best-effort , scheduler be-sched Forwarding class be , scheduler be-sched Forwarding class network-control , scheduler nc-sched Scheduler map gd-map : Forwarding class fcoe , scheduler fcoe-sched Forwarding class no-loss , scheduler nl-sched Scheduler map hpc-map : Forwarding class hpc , scheduler hpc-sched
Traffic control profiles	be-tcp : scheduler map be-map , minimum bandwidth 3.5g , maximum bandwidth 100% gd-tcp : scheduler map gd-map , minimum bandwidth 4.5g , maximum bandwidth 100% hpc-tcp : scheduler map hpc-map , minimum bandwidth 2g , maximum bandwidth 100%
Interfaces	This example configures hierarchical port scheduling on interfaces xe-0/0/20 and xe-0/0/21 . Because traffic is bidirectional, you apply the ingress and egress configuration components to both interfaces: <ul style="list-style-type: none"> • Classifier Name—hsclassifier1 • Forwarding class sets—best-effort-pg, guar-deliver-pg, hpc-pg • Congestion notification profile—gd-cnp

Figure 210 on page 5610 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example. You can perform the configuration steps in a different sequence if you want.

Figure 210: Hierarchical Port Scheduling Components Block Diagram

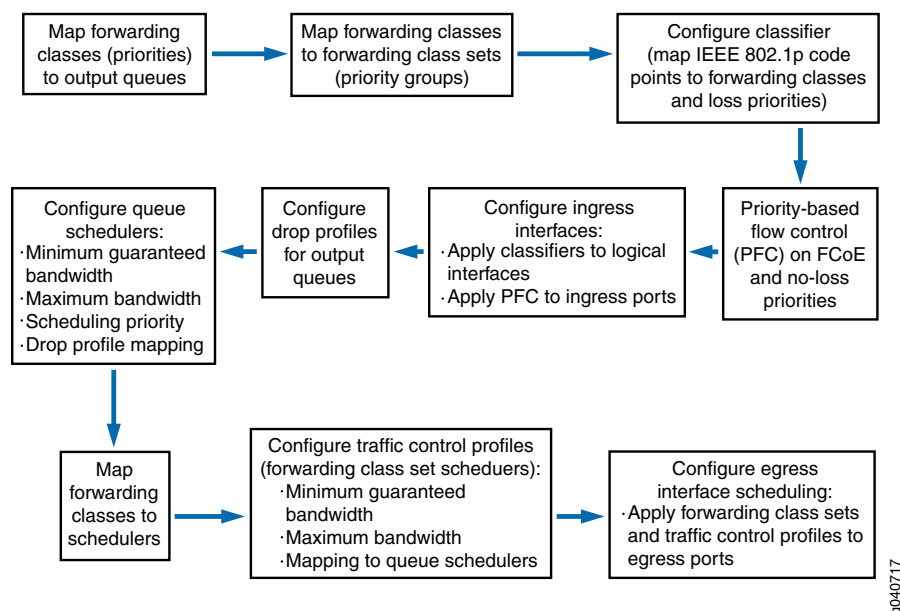
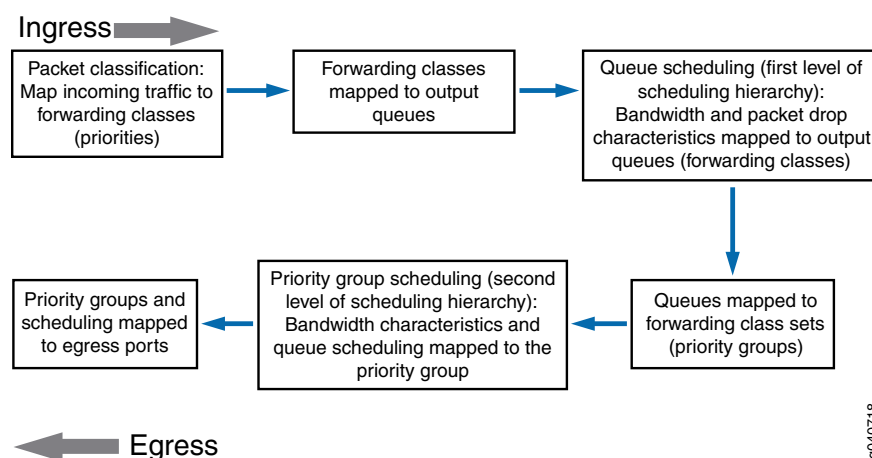


Figure 211 on page 5610 shows a block diagram of the hierarchical scheduling packet flow from ingress to egress.

Figure 211: Hierarchical Port Scheduling Packet Flow Block Diagram



Configuration

CLI Quick Configuration

To quickly configure hierarchical port scheduling, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network

configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets guar-delivery-pg class fcoe
set forwarding-class-sets guar-delivery-pg class no-loss
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points 000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority low code-points 011
set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority low code-points 100
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low code-points 110
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011 pfc
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100 pfc
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability 80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability 100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability 100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability 75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers fcoe-sched priority low transmit-rate 2500m
set schedulers fcoe-sched shaping-rate percent 100
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set schedulers nl-sched priority low transmit-rate 2g
set schedulers nl-sched shaping-rate percent 100
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate 4500m
set traffic-control-profiles gd-tcp shaping-rate percent 100
```

```

set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp

```

Step-by-Step Procedure

To perform a step-by-step configuration of the forwarding classes (priorities), forwarding class sets (priority groups), classifiers, queue schedulers, PFC, traffic control profiles, and interfaces to set up hierarchical port scheduling (ETS):

1. Configure the forwarding classes (priorities) and map them to unicast output queues (do not explicitly map the **fcoe** and **no-loss** forwarding classes to output queues; use the default configuration):

```

[edit class-of-service]
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class be2 queue-num 1
user@switch# set forwarding-classes class hpc queue-num 5
user@switch# set forwarding-classes class network-control queue-num 7

```

2. Configure forwarding class sets (priority groups) to group forwarding classes (priorities) that require similar CoS treatment:

```

[edit class-of-service]
user@switch# set forwarding-class-sets best-effort-pg class best-effort
user@switch# set forwarding-class-sets best-effort-pg class be2
user@switch# set forwarding-class-sets best-effort-pg class network-control
user@switch# set forwarding-class-sets guar-delivery-pg class fcoe
user@switch# set forwarding-class-sets guar-delivery-pg class no-loss
user@switch# set forwarding-class-sets hpc-pg class hpc

```

3. Configure a classifier to set the loss priority and IEEE 802.1 code points assigned to each forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort
loss-priority low code-points 000
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority
high code-points 001
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority
low code-points 011
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority
low code-points 100
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low
code-points 101
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control
loss-priority low code-points 110

```

4. Configure a congestion notification profile to enable PFC on the FCoE and no-loss queue IEEE 802.1 code points:

```

[edit class-of-service]
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011
pfc
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100
pfc

```


5. Assign the classifier to the interfaces:


```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
```
6. Apply the PFC configuration to the interfaces:


```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
```
7. Configure the drop profile for the best-effort low loss-priority queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50
drop-probability 0 drop-probability 80
```
8. Configure the drop profile for the best-effort high loss-priority queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40
drop-probability 0 drop-probability 100
```
9. Configure the drop profile for the network-control queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability
0 drop-probability 100
```
10. Configure the drop profile for the high-performance computing queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability
0 drop-probability 75
```
11. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profiles for the best-effort queue:


```
[edit class-of-service]
user@switch# set schedulers be-sched priority low transmit-rate 3g
user@switch# set schedulers be-sched shaping-rate percent 100
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile dp-be-low
user@switch# set schedulers be-sched drop-profile-map loss-priority high protocol any
drop-profile dp-be-high
```
12. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the FCoE queue:


```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 2500m
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
13. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the high-performance computing queue:


```
[edit class-of-service]
user@switch# set schedulers hpc-sched priority low transmit-rate 2g
user@switch# set schedulers hpc-sched shaping-rate percent 100
user@switch# set schedulers hpc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-hpc
```

14. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the network-control queue:

```
[edit class-of-service]
user@switch# set schedulers nc-sched priority low transmit-rate 500m
user@switch# set schedulers nc-sched shaping-rate percent 100
user@switch# set schedulers nc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-nc
```
15. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the no-loss queue:

```
[edit class-of-service]
user@switch# set schedulers nl-sched priority low transmit-rate 2g
user@switch# set schedulers nl-sched shaping-rate percent 100
```
16. Map the schedulers to the appropriate forwarding classes (queues):

```
[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class be2 scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class network-control scheduler
nc-sched
user@switch# set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
user@switch# set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
```
17. Define the traffic control profile for the best-effort priority group (queue scheduler to mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate
3500m
user@switch# set traffic-control-profiles be-tcp shaping-rate percent 100
```
18. Define the traffic control profile for the guaranteed delivery priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate
4500m
user@switch# set traffic-control-profiles gd-tcp shaping-rate percent 100
```
19. Define the traffic control profile for the high-performance computing priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):

```
[edit class-of-service]
user@switch# set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate
2g
user@switch# set traffic-control-profiles hpc-tcp shaping-rate percent 100
```
20. Apply the three priority groups (forwarding class sets) and the appropriate traffic control profiles to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```

```

user@switch# set interfaces xe-0/0/21 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp

```

Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** and **no-loss** lossless forwarding classes):

```

user@switch> show configuration class-of-service
classifiers {
  ieee-802.1 hsclassifier1 {
    forwarding-class best-effort {
      loss-priority low code-points 000;
    }
    forwarding-class be2 {
      loss-priority high code-points 001;
    }
    forwarding-class fcoe {
      loss-priority low code-points 011;
    }
    forwarding-class no-loss {
      loss-priority low code-points 100;
    }
    forwarding-class hpc {
      loss-priority low code-points 101;
    }
    forwarding-class network-control {
      loss-priority low code-points 110;
    }
  }
}
drop-profiles {
  dp-be-low {
    interpolate {
      fill-level [ 25 50 ];
      drop-probability [ 0 80 ];
    }
  }
  dp-be-high {
    interpolate {
      fill-level [ 10 40 ];
      drop-probability [ 0 100 ];
    }
  }
  dp-hpc {
    interpolate {
      fill-level [ 75 90 ];
      drop-probability [ 0 75 ];
    }
  }
  dp-nc {
    interpolate {

```

```
        fill-level [ 80 100 ];
        drop-probability [ 0 100 ];
    }
}
forwarding-classes {
    class best-effort queue-num 0;
    class be2 queue-num 1;
    class hpc queue-num 5;
    class network-control queue-num 7;
}
traffic-control-profiles {
    be-tcp {
        scheduler-map be-map;
        shaping-rate percent 100;
        guaranteed-rate 3500000000;
    }
    gd-tcp {
        scheduler-map gd-map;
        shaping-rate percent 100;
        guaranteed-rate 4500000000;
    }
    hpc-tcp {
        scheduler-map hpc-map;
        shaping-rate percent 100;
        guaranteed-rate 2g;
    }
}
forwarding-class-sets {
    guar-delivery-pg {
        class fcoe;
        class no-loss;
    }
    best-effort-pg {
        class best-effort;
        class be2;
        class network-control;
    }
    hpc-pg {
        class hpc;
    }
}
congestion-notification-profile {
    gd-cnp {
        input {
            ieee-802.1 {
                code-point 011 {
                    pfc;
                }
                code-point 100 {
                    pfc;
                }
            }
        }
    }
}
```

```

interfaces {
  xe-0/0/20 {
    forwarding-class-set {
      best-effort-pg {
        output-traffic-control-profile be-tcp;
      }
      guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
      }
      hpc-pg {
        output-traffic-control-profile hpc-tcp;
      }
    }
    congestion-notification-profile gd-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 hsclassifier1;
      }
    }
  }
  xe-0/0/21 {
    forwarding-class-set {
      best-effort-pg {
        output-traffic-control-profile be-tcp;
      }
      guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
      }
      hpc-pg {
        output-traffic-control-profile hpc-tcp;
      }
    }
    congestion-notification-profile gd-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 hsclassifier1;
      }
    }
  }
}
scheduler-maps {
  be-map {
    forwarding-class best-effort scheduler be-sched;
    forwarding-class network-control scheduler nc-sched;
    forwarding-class be2 scheduler be-sched;
  }
  gd-map {
    forwarding-class fcoe scheduler fcoe-sched;
    forwarding-class no-loss scheduler nl-sched;
  }
  hpc-map {
    forwarding-class hpc scheduler hpc-sched;
  }
}
schedulers {
  be-sched {

```

```
    transmit-rate 3g;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-be-low;
    drop-profile-map loss-priority high protocol any drop-profile dp-be-high;
  }
  fcoe-sched {
    transmit-rate 2500000000;
    shaping-rate percent 100;
    priority low;
  }
  hpc-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-hpc;
  }
  nc-sched {
    transmit-rate 500m;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-nc;
  }
  nl-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
  }
}
```



TIP: To quickly configure the interfaces, issue the `load merge terminal` command, and then copy the hierarchy and paste it into the switch terminal window.

Verification

To verify that the hierarchical port scheduling components have been created and are operating properly, perform these tasks:

- [Verifying That the Forwarding Classes \(Priorities\) Have Been Created on page 5619](#)
- [Verifying That the Forwarding Class Sets \(Priority Groups\) Have Been Created on page 5619](#)
- [Verifying That the Classifier Has Been Created on page 5620](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5620](#)
- [Verifying That the Output Queue Schedulers Have Been Created on page 5621](#)
- [Verifying That the Drop Profiles Have Been Created on page 5624](#)

- [Verifying That the Priority Group Output Schedulers \(Traffic Control Profiles\) Have Been Created on page 5625](#)
- [Verifying the Interface Configuration on page 5626](#)

Verifying That the Forwarding Classes (Priorities) Have Been Created

Purpose Verify that the forwarding classes have been created and mapped to the correct queues. (The system shows only the explicitly configured forwarding classes. It does not show default forwarding classes such as **fcoe** and **no-loss**.)

Action List the forwarding classes using the operational mode command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
be2	1	3	normal	Disabled
hpc	2	4	normal	Disabled
network-control	3	7	normal	Disabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command lists all of the configured forwarding classes, the internal identification number of each forwarding class, the queues that are mapped to the forwarding classes, the policing priority, and whether the forwarding class is lossless (no-loss packet drop attribute enabled) or lossy forwarding class (no-loss packet drop attribute disabled). The command output shows that:

- Forwarding class **best-effort** maps to queue **0** and is lossy
- Forwarding class **be2** maps to queue **1** and is lossy
- Forwarding class **hpc** maps to queue **5** and is lossy
- Forwarding class **network-control** maps to queue **7** and is lossy

In addition, the command lists the default multicast (multidestination) forwarding class and the default queue to which it is mapped.

Verifying That the Forwarding Class Sets (Priority Groups) Have Been Created

Purpose Verify that the priority groups have been created and that the correct priorities (forwarding classes) belong to the appropriate priority group.

Action List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set**:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: best-effort-pg, Type: normal-type, Forwarding class set
index: 19907
```

Forwarding class	Index
best-effort	0

be2	1
network-control	5

Forwarding class set: guar-delivery-pg, Type: normal-type, Forwarding class set index: 43700

Forwarding class	Index
fcoe	2
no-loss	3

Forwarding class set: hpc-pg, Type: normal-type, Forwarding class set index: 60758

Forwarding class	Index
hpc	4

Meaning The **show class-of-service forwarding-class-set** command lists all of the configured forwarding class sets (priority groups), the forwarding classes (priorities) that belong to each priority group, and the internal index number of each priority group. The command output shows that:

- The forwarding class set **best-effort-pg** includes the forwarding classes **best-effort**, **be2**, and **network-control**.
- The forwarding class set **guar-delivery-pg** includes the forwarding classes **fcoe** and **no-loss**.
- The forwarding class set **hpc-pg** includes the forwarding class **hpc**.

Verifying That the Classifier Has Been Created

Purpose Verify that the classifier maps forwarding classes to the correct IEEE 802.1p code points and packet loss priorities.

Action List the classifier configured for hierarchical port scheduling using the operational mode command **show class-of-service classifier name hsclassifier1**:

```
user@switch> show class-of-service classifier name hsclassifier1
Classifier: hsclassifier1, Code point type: ieee-802.1, Index: 43607
  Code point      Forwarding class      Loss priority
  ---
  000             best-effort             low
  001             be2                 high
  011             fcoe                 low
  100             no-loss             low
  101             hpc                 low
  110             network-control     low
```

Meaning The **show class-of-service classifier name hsclassifier1** command lists all of the IEEE 802.1p code points and the loss priorities mapped to all of the forwarding classes in the classifier. The command output shows that the forwarding classes **best-effort**, **be2**, **no-loss**, **fcoe**, **hpc**, and **network-control** have been created and mapped to IEEE 802.1p code points and loss priorities.

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the correct priorities for lossless transport.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: gd-cnp, Index: 51687
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2500
100	Enabled	2500
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	0
001	0
010	1
011	2
100	3
101	4
110	5
111	6
	7

Meaning The **show class-of-service congestion-notification** command lists all of the congestion notification profiles and the IEEE 802.1p code points with PFC enabled. The command output shows that PFC is enabled for code points **011** (**fcoe** priority and queue) and **100** (**no-loss** priority and queue) for the **gd-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying That the Output Queue Schedulers Have Been Created

Purpose Verify that the output queue schedulers have been created with the correct bandwidth parameters and priorities, mapped to the correct queues, and mapped to the correct drop profiles.

Action List the scheduler maps using the operational mode command **show class-of-service scheduler-map**:

```
user@switch> show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 64023
```

```
Scheduler: be-sched, Forwarding class: best-effort, Index: 13005
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
```

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	55387	dp-be-low
Medium high	any	1	<default-drop-profile>
High	any	4369	dp-be-high

Scheduler: be-sched, Forwarding class: be2, Index: 13005

Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	55387	dp-be-low
Medium high	any	1	<default-drop-profile>
High	any	4369	dp-be-high

Scheduler: nc-sched, Forwarding class: network-control, Index: 45740

Transmit rate: 5000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	dp-nc
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler map: gd-map, Index: 61447

Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289

Transmit rate: 25000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler: nl-sched, Forwarding class: no-loss, Index: 29359

Transmit rate: 20000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	44207	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

Scheduler map: hpc-map, Index: 56941

```

Scheduler: hpc-sched, Forwarding class: hpc, Index: 55900
Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       57716  dp-hpc
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>

```

Meaning The **show class-of-service scheduler-map** command lists all of the configured scheduler maps. For each scheduler map, the command output includes:

- The name of the scheduler map (**scheduler-map** field)
- The name of the scheduler (**scheduler** field)
- The forwarding classes mapped to the scheduler (**forwarding-class** field)
- The minimum guaranteed queue bandwidth (**transmit-rate** field)
- The scheduling priority (**priority** field)
- The maximum bandwidth in the priority group the queue can consume (**shaping-rate** field)
- The drop profile loss priority (**loss priority** field) for each drop profile name (**name** field)

The command output shows that:

- The scheduler map **be-map** has been created and has these properties:
 - There are two schedulers, **be-sched** and **nc-sched**.
 - The scheduler **be-sched** has two forwarding classes, **best-effort** and **be2**.
 - Scheduler **be-sched** forwarding classes **best-effort** and **be2** share a minimum guaranteed bandwidth of **3000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and use the drop profile **dp-be-low** for low loss-priority traffic, the default drop profile for medium-high loss-priority traffic, and the drop profile **dp-be-high** for high loss-priority traffic.
 - The scheduler **nc-sched** has one forwarding class, **network-control**.
 - The **network-control** forwarding class has a minimum guaranteed bandwidth of **5000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-nc** for low loss-priority traffic and the default drop profile for medium-high and high loss priority traffic.
- The scheduler map **gd-map** has been created and has these properties:
 - There are two schedulers, **fcoe-sched** and **nl-sched**.
 - The scheduler **fcoe-sched** has one forwarding class, **fcoe**.

- The **fcoe** forwarding class has a minimum guaranteed bandwidth of **2500000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler **nl-sched** has one forwarding class, **no-loss**.
- The **no-loss** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler map **hpc-map** has been created and has these properties:
 - There is one scheduler, **hpc-sched**.
 - The scheduler **hpc-sched** has one forwarding class, **hpc**.
 - The **hpc** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-hpc** for low loss-priority traffic and the default drop profile for medium-high and high loss-priority traffic.

Verifying That the Drop Profiles Have Been Created

Purpose Verify that the drop profiles **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc** have been created with the correct fill levels and drop probabilities.

Action List the drop profiles using the operational mode command **show configuration class-of-service drop-profiles**:

```
user@switch> show configuration class-of-service drop-profiles
dp-be-low {
    interpolate {
        fill-level [ 25 50 ];
        drop-probability [ 0 80 ];
    }
}
dp-be-high {
    interpolate {
        fill-level [ 10 40 ];
        drop-probability [ 0 100 ];
    }
}
dp-hpc {
    interpolate {
        fill-level [ 75 90 ];
        drop-probability [ 0 75 ];
    }
}
dp-nc {
    interpolate {
        fill-level [ 80 100 ];
        drop-probability [ 0 100 ];
    }
}
```

Meaning The **show configuration class-of-service drop-profiles** command lists the drop profiles and their properties. The command output shows that there are four drop profiles configured, **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc**. The output also shows that:

- For **dp-be-low**, the drop start point (the first fill level) is when the queue is 25 percent filled, the drop end point (the second fill level) occurs when the queue is 50 percent filled, and the drop probability at the drop end point is 80 percent.
- For **dp-be-high**, the drop start point (the first fill level) is when the queue is 10 percent filled, the drop end point (the second fill level) occurs when the queue is 40 percent filled, and the drop probability at the drop end point is 100 percent.
- For **dp-hpc**, the drop start point (the first fill level) is when the queue is 75 percent filled, the drop end point (the second fill level) occurs when the queue is 90 percent filled, and the drop probability at the drop end point is 75 percent.
- For **dp-nc**, the drop start point (the first fill level) is when the queue is 80 percent filled, the drop end point (the second fill level) occurs when the queue is 100 percent filled, and the drop probability at the drop end point is 100 percent.

Verifying That the Priority Group Output Schedulers (Traffic Control Profiles) Have Been Created

Purpose Verify that the traffic control profiles **be-tcp**, **gd-tcp**, and **hpc-tcp** have been created with the correct bandwidth parameters and scheduler mapping.

Action List the traffic control profiles using the operational mode command **show class-of-service traffic-control-profile**:

```
user@switch> show class-of-service traffic-control-profile
Traffic control profile: be-tcp, Index: 40535
  Shaping rate: 100 percent
  Scheduler map: be-map
  Guaranteed rate: 3500000000

Traffic control profile: gd-tcp, Index: 37959
  Shaping rate: 100 percent
  Scheduler map: gd-map
  Guaranteed rate: 4500000000

Traffic control profile: hpc-tcp, Index: 47661
  Shaping rate: 100 percent
  Scheduler map: hpc-map
  Guaranteed rate: 2000000000
```

Meaning The **show class-of-service traffic-control-profile** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**traffic-control-profile**)
- The maximum port bandwidth the priority group can consume (**shaping-rate**)
- The scheduler map associated with the traffic control profile (**scheduler-map**)
- The minimum guaranteed priority group port bandwidth (**guaranteed-rate**)

The command output shows that:

- The traffic control profile **be-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **be-map**, and has a minimum guaranteed bandwidth of **3500000000 bps**.
- The traffic control profile **gd-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **gd-map**, and has a minimum guaranteed bandwidth of **4500000000 bps**.
- The traffic control profile **hpc-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **hpc-map**, and has a minimum guaranteed bandwidth of **2000000000 bps**.

Verifying the Interface Configuration

Purpose Verify that the classifier, the congestion notification profile, and the forwarding class sets are configured on interfaces **xe-0/0/20** and **xe-0/0/21**.

Action List the interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hscclassifier1;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hscclassifier1;
    }
}
```

Meaning The `show configuration class-of-service interfaces interface-name` command shows that each interface includes the forwarding class sets **best-effort-pg**, **guar-delivery-pg**, and **hpc-pg**, congestion notification profile **gd-cnp**, and the IEEE 802.1p classifier **hsclassifier1**.

- Related Documentation**
- [Defining CoS Unicast BA Classifiers on page 5784](#)
 - [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
 - [Assigning CoS Components to Interfaces on page 5807](#)
 - [Example: Configuring Tail-Drop Profiles on page 5663](#)
 - [Example: Configuring Drop Profile Maps on page 5666](#)
 - [Example: Configuring Forwarding Classes on page 5668](#)
 - [Example: Configuring Forwarding Class Sets on page 5671](#)
 - [Example: Configuring Queue Schedulers on page 5674](#)
 - [Example: Configuring Queue Scheduling Priority on page 5679](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
 - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
 - [Example: Configuring Maximum Output Bandwidth on page 5689](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
 - [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
 - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
 - [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5502](#)
 - [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
 - [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)

Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS flag in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

To configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic.
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.
- Configure the bandwidth scheduling for the FCoE forwarding class output queue.
- Create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- Configure the bandwidth scheduling for the FCoE priority group.
- Apply the scheduling to the egress interfaces.



NOTE: If you are using Junos OS Release 12.2 or later, use the default forwarding classes for the lossless fcoe forwarding class. If you explicitly configure default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 5628](#)
- [Overview on page 5628](#)
- [Configuration on page 5630](#)
- [Verification on page 5633](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to:

- Assign FCoE traffic to the FCoE priority at the ingress.
- Create and apply CoS for the FCoE traffic using ETS (hierarchical port scheduling).

- Apply PFC to the FCoE traffic.
- Apply the configuration to ingress and egress interfaces.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Each interface in this example is configured as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and port scheduling are applied to all of the interfaces.

Topology

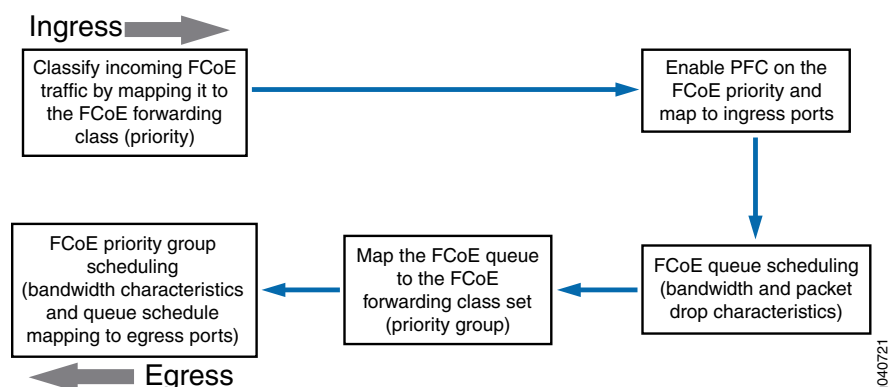
Table 405 on page 5114 shows the configuration components for this example.

Table 526: Components of the PFC for FCoE Traffic Configuration Topology

Component	Settings
Hardware	QFX3500 switch
Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point)	Code point 011 to forwarding class fcoe and loss priority low Ingress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34
PFC congestion notification profile	fcoe-cnp: Code point 011 Ingress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34
FCoE queue scheduler	fcoe-sched: Minimum bandwidth 3g Maximum bandwidth 100% Priority low
Forwarding class-to-scheduler mapping	Scheduler map fcoe-map: Forwarding class fcoe Scheduler fcoe-sched
Forwarding class set (FCoE priority group)	fcoe-pg: Forwarding class fcoe Egress interfaces: xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34
Traffic control profile	fcoe-tcp: Scheduler map fcoe-map Minimum bandwidth 3g Maximum bandwidth 100%

Figure 196 on page 5115 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 212: PFC for FCoE Traffic Configuration Components Block Diagram



Configuration

CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

Step-by-Step Procedure

To configure the FCoE forwarding class (priority), ingress classifier, output queue scheduling, forwarding class set (priority group) and its output port scheduling, PFC application, and interfaces to set up PFC for FCoE traffic:

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:


```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
```
2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:

- ```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
3. Apply the PFC configuration to the ingress interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
```
  4. Assign the classifier to the ingress interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
```
  5. Configure output scheduling for the FCoE queue:
 

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
  6. Map the FCoE forwarding class to the FCoE scheduler:
 

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
  7. Configure the forwarding class set for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
  8. Define the traffic control profile for the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
  9. Apply the FCoE forwarding class set and traffic control profile to the egress ports:
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

### Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class):

```
user@switch> show configuration class-of-service
classifiers {
 ieee-802.1 fcoe-classifier {
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 }
}
```

```
}
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/31 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
 xe-0/0/32 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
 xe-0/0/33 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 }
}
```

```

 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
xe-0/0/34 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}
}

```



**TIP:** To quickly configure the interfaces, issue the **load merge terminal** command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5633](#)
- [Verifying the Ingress Interface PFC Configuration on page 5634](#)

### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |
| 100      |                     |
|          | 4                   |
| 101      |                     |
|          | 5                   |
| 110      |                     |
|          | 6                   |
| 111      |                     |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point **011** for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

#### *Verifying the Ingress Interface PFC Configuration*

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
```

```
congestion-notification-profile fcoe-cnp;
```

```
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
```

```

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

```

**Meaning** The `show configuration class-of-service interfaces` commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX3500 switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree protocol (STP).

You can use an MC-LAG to provide a redundant aggregation layer for Fiber Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX3500 switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.



**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX3500 switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX3500 switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see [“Example: Configuring Multichassis Link Aggregation” on page 2462](#). However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.

QFX3500 Node devices support MC-LAGs. QFabric system Node devices do not support MC-LAGs.

This topic describes:

- [Requirements on page 5636](#)
- [Overview on page 5636](#)
- [Configuration on page 5641](#)
- [Verification on page 5649](#)

---

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 Switches that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX3500 Switches that provide FCoE server access in transit switch mode, and connect to the MC-LAG switches. These switches can be standalone QFX3500 switches or they can be Node devices in a QFabric system.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 12.2 or later for the QFX Series.

---

## Overview

FCoE traffic requires lossless transport. This example shows you how to:



- Configure CoS for FCoE traffic on the two QFX3500 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.


- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Disable IGMP snooping on the FCoE VLAN.
- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

This example does not show how to configure the MC-LAG itself. This example does include a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

### Topology

QFX3500 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 197 on page 5123](#):

**Figure 213: Supported Topology for an MC-LAG on an FCoE Transit Switch**



[Table 406 on page 5123](#) shows the configuration components for this example.

Table 527: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

| Component                                                                  | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                   | Four QFX3500 switches (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Forwarding class (all switches)                                            | Default <b>fcoe</b> forwarding class.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| LAGs and MC-LAG                                                            | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2.<br/>Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1.<br/>Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.<br/>Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.<br/>Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                                        | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Forwarding class-to-scheduler mapping (all switches)                       | Scheduler map <b>fcoe-map:</b><br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 527: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)**

| Component                                                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class set (FCoE priority group, all switches) | <p><b>fcoe-pg:</b><br/>Forwarding class <b>fcoe</b></p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                     |
| Traffic control profile (all switches)                   | <p><b>fcoe-tcp:</b><br/>Scheduler map <b>fcoe-map</b><br/>Minimum bandwidth <b>3g</b><br/>Maximum bandwidth <b>100%</b></p>                                                                                                                                                                                                                                                                                                                                                                                                    |
| PFC congestion notification profile (all switches)       | <p><b>fcoe-cnp:</b><br/>Code point <b>011</b></p> <p>Ingress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                          |
| FCoE VLAN name and tag ID                                | <p>Name—<b>fcoe_vlan</b><br/>ID—<b>100</b></p> <p>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.</p> <p>Disable IGMP snooping on the FCoE VLAN on all four switches.</p>                                                                                                                                                                                                                                                                                                                |
| FIP snooping                                             | <p>Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.</p> <p>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.</p> |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode and tagged access mode ports if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is `fcoe`, and the default output queue is queue `3`.



**NOTE:** In Junos OS Release 12.2, traffic mapped to explicitly configured forwarding classes, even lossless forwarding classes such as `fcoe`, is treated as lossy (best-effort) traffic and does *not* receive lossless treatment. To receive lossless treatment in release 12.2, traffic must use one of the default lossless forwarding classes (`fcoe` or `no-loss`).

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.

- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk and tagged-access port modes. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (011) to the FCoE forwarding class. You can configure the BA classifier if you wish, but you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point 011 in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers and how to disable IGMP

snooping on the FCoE VLAN. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. ([“Example: Configuring Multichassis Link Aggregation” on page 2462](#) describes how to configure MC-LAGs.)
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. ([“Configuring Link Aggregation” on page 2537](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

## Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [Configuring MC-LAG Switches S1 and S2 on page 5643](#)
- [Configuring FCoE Transit Switches TS1 and TS2 on page 5644](#)
- [Results on page 5646](#)

### CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the [edit] hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
```

```
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
```

```
set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the [edit] hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access vlan members fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access vlan members fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access vlan members fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access vlan members fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
set ethernet-switching-options secure-access-port vlan fcoe_vlan examine-fip examine-vn2v2 beacon-period 90000
```

*Configuring MC-LAG Switches S1 and S2*

**Step-by-Step Procedure** To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:
 

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):
 

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:
 

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```
6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:
 

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
7. Apply the PFC configuration to the LAG and MC-LAG interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):
 

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
9. Disable IGMP snooping on the FCoE VLAN:
 

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG between the two MC-LAG switches:

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```

11. Add the member interfaces to the MC-LAG:

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```

12. Configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and MC-LAG interfaces (2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes):

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```

14. Set the LAG and MC-LAG interfaces as FCoE trusted ports (ports that connect to other switches should be trusted and should not perform FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

### *Configuring FCoE Transit Switches TS1 and TS2*

#### **Step-by-Step Procedure**

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:



- ```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:


```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
 5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```
 6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:


```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
 7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile
fcoe-cnp
```
 8. Configure the VLAN for FCoE traffic (**fcoe_vlan**):


```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
 9. Disable IGMP snooping on the FCoE VLAN:


```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```
 10. Add the member interfaces to the LAG:


```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```
 11. On the LAG (**ae1**), configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe_vlan**):

```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

12. On the FCoE access interfaces (xe-0/0/30, xe-0/0/31, xe-0/0/32, xe-0/0/33), configure the port mode as **tagged-access** and membership in the FCoE VLAN (**fcoe_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and FCoE access interfaces (2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes):

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

14. Set the LAG interface as an FCoE trusted port (ports that connect to other switches should be trusted and should not perform FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

15. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN_Port FIP snooping; the example is equally valid if you use VN2VF_Port FIP snooping):

```
user@switch# set ethernet-switching-options secure-access-port vlan fcoe_vlan
examine-fip examine-vn2v2 beacon-period 90000
```

Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are similar):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
```

```

}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  ae0 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  ae1 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}

```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see [“Example: Configuring Multichassis Link Aggregation” on page 2462](#).

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are similar):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  xe-0/0/30 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/31 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/32 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/33 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
```

```

}
ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
}

```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 5650](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created on page 5650](#)
- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created on page 5651](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5651](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 5652](#)
- [Verifying That the Interfaces Are Correctly Configured on page 5654](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 5656](#)
- [Verifying That the FIP Snooping Mode is Correct on FCoE Transit Switches TS1 and TS2 on page 5657](#)
- [Verifying That IGMP Snooping Is Disabled on the FCoE VLAN on page 5657](#)

Verifying That the Output Queue Schedulers Have Been Created

Purpose Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

Action List the scheduler map using the operational mode command **show class-of-service scheduler-map fcoe-map**:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      <default-drop-profile>
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created

Purpose Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

Action List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
Shaping rate: 100 percent
```

Scheduler map: fcoe-map
Guaranteed rate: 3000000000

Meaning The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

Verifying That the Forwarding Class Set (Priority Group) Has Been Created

Purpose Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

Action List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class          Index
  fcoe                      1
```

Meaning The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

Action List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority  PFC          MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled    2500
  100      Disabled
  101      Disabled
```

110	Disabled
111	Disabled
Type: Output	
Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	
	7

Meaning The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying That the Interface Class of Service Configuration Has Been Created

Purpose Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

Action List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}

ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
```


List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
ae1 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
```

Meaning The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



NOTE: Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

Verifying That the Interfaces Are Correctly Configured

Purpose Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

Action List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/20 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
```

```

        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```

user@switch> show configuration interfaces
xe-0/0/25 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/26 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/30 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/31 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/32 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {

```

```
        port-mode tagged-access;
        vlan {
            members fcoe_vlan;
        }
    }
}

ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
```

Meaning The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The port mode (**trunk** mode for interfaces that connect two switches, **tagged-access** mode for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe_vlan**)

Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

Purpose Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Switch TS1 and TS2 FCoE access ports.

Action List the port security configuration on FCoE transit switches TS1 and TS2 using the operational mode command **show configuration ethernet-switching-options secure-access-port**:

```
user@switch> show configuration ethernet-switching-options secure-access-port
interface ae1.0 {
    fcoe-trusted;
}
vlan fcoe_vlan {
```

```
    examine-fip;
}
```

Meaning The **show configuration ethernet-switching-options secure-access-port** command lists port security information, including whether a port is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- FIP snooping is enabled (**examine-fip**) on the FCoE VLAN (**fcoe_vlan**). On Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

Verifying That the FIP Snooping Mode is Correct on FCoE Transit Switches TS1 and TS2

Purpose Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Switch TS1 and TS2 FCoE access ports.

Action List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,    Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
...
```



NOTE: The output has been truncated to show only the relevant information.

Meaning The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe_vlan**
- The FIP snooping mode is VN2VN_Port FIP snooping (**VN2VN Snooping**)

Verifying That IGMP Snooping Is Disabled on the FCoE VLAN

Purpose Verify that IGMP snooping is disabled on the FCoE VLAN on all four switches.

Action List the IGMP snooping protocol information on each of the four switches using the **show configuration protocols igmp-snooping** command:

```
user@switch> show configuration protocols igmp-snooping
vlan fcoe_vlan {
    disable;
}
```

Meaning The **show configuration protocols igmp-snooping** command lists the IGMP snooping configuration for the VLANs configured on the switch. The command output shows that IGMP snooping is disabled on the FCoE VLAN (**fcoe_vlan**).

- Related Documentation**
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
 - [Configuring Link Aggregation on page 2537](#)
 - [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
 - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Understanding Multichassis Link Aggregation on page 2435](#)
 - [Understanding MC-LAGs on an FCoE Transit Switch on page 5047](#)

Example: Configuring Unicast Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. You apply classifiers to ingress interfaces.

- [Requirements on page 5658](#)
- [Overview on page 5658](#)
- [Configuring Unicast Classifiers on page 5659](#)
- [Verification on page 5660](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP or DSCP IPv6) value or IEEE 802.1p value.

- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.



NOTE: You must assign unicast traffic and multidestination (multicast, broadcast, and destination lookup fail) traffic to different classifiers. One classifier cannot include both unicast and multidestination forwarding classes. A unicast classifier can include only forwarding classes for unicast traffic.

This example describes how to configure a BA classifier called **ba-ucast-classifier** as the default IEEE 802.1 map and apply it to ingress interface **xe-0/0/10**. The BA classifier assigns loss priorities, as shown in [Table 528 on page 5659](#), to incoming packets in the four forwarding classes.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

Table 528: ba-ucast-classifier Loss Priority Assignments

Unicast Forwarding Class	For CoS Traffic Type	ba-ucast-classifier Loss Priority to IEEE 802.1p Code Point Mapping	Packet Drop Attribute
be	Best-effort traffic	Low loss priority code point: 000	drop
fcoe	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic	Low loss priority code point: 011	no-loss
no-loss	Guaranteed delivery for TCP traffic	Low loss priority code point: 100	no-loss
nc	Network-control traffic	Low loss priority code point: 110	drop

Configuring Unicast Classifiers

To configure a unicast IEEE 802.1 BA classifier named **ba-ucast-classifier** as the default IEEE 802.1 map:

1. Associate code point 000 with forwarding class **be** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier import default forwarding-class be
loss-priority low code-points 000
```
2. Associate code point 011 with forwarding class **fcoe** and loss priority **low**:

```
[edit class-of-service classifiers]
```

```
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class fcoe loss-priority low
code-points 011
```

3. Associate code point **100** with forwarding class **no-loss** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class no-loss loss-priority low
code-points 100
```

4. Associate code point **110** with forwarding class **nc** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class nc loss-priority low
code-points 110
```

5. Apply the unicast classifier to ingress interface **xe-0/0/10**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/10 unit 0 classifiers ieee-802.1 ba-ucast-classifier
```

Verification

To verify the unicast classifier configuration, perform these tasks:

- [Verifying the Unicast Classifier Configuration on page 5660](#)
- [Verifying the Ingress Interface Configuration on page 5660](#)

Verifying the Unicast Classifier Configuration

Purpose Verify that you configured the unicast classifier with the correct forwarding classes, loss priorities, and code points.

Action List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier
  forwarding-class be {
    loss-priority low code-points 000;
  }
  forwarding-class fcoe {
    loss-priority low code-points 011;
  }
  forwarding-class no-loss {
    loss-priority low code-points 100;
  }
  forwarding-class nc
    loss-priority low code-points 110;
}
```

Verifying the Ingress Interface Configuration

Purpose Verify that the unicast classifier **ba-ucast-classifier** is attached to ingress interface **xe-0/0/10**.

Action List the ingress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/10**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/10
congestion-notification-profile fcoe-cnp;
unit 0 {
  classifiers {
```



```

        ieee-802.1 ba-ucast-classifier;
    }
}

```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)
- [Monitoring CoS Classifiers on page 5915](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class.

- [Requirements on page 5661](#)
- [Overview on page 5661](#)
- [Configuring Multidestination Classifiers on page 5662](#)
- [Verification on page 5662](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.



NOTE: DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.

- Multifield traffic classifiers—Examine multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.



NOTE: You must assign unicast traffic and multicast traffic to different classifiers. One classifier cannot include both unicast and multicast forwarding classes. A multidestination classifier can include only forwarding classes for multicast traffic.

The following example describes how to configure a BA classifier called **ba-mcast-classifier**, which is applied to all of the switch interfaces. The BA classifier assigns loss priorities, as shown in [Table 529 on page 5662](#), to incoming packets in the multidestination forwarding class.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

Table 529: BA-mcast-classifier Loss Priority Assignments

Multicast Forwarding Class	For CoS Traffic Type	ba-mcast-classifier Assignment
mcast	Best-effort multicast traffic	Low loss priority code point: 000

Configuring Multidestination Classifiers

To configure a multicast IEEE 802.1 BA classifier named **ba-mcast-classifier**:

1. Associate code point 000 with forwarding class **mcast** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-mcast-classifier forwarding-class mcast-be loss-priority
low code-points 000
```

2. Configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers ieee-802.1 ba-mcast-classifier
```

Verification

To verify the multidestination classifier configuration, perform these tasks:

- [Verifying the IEEE 802.1 Multidestination Classifier on page 5663](#)
- [Verifying the Multidestination Classifier Configuration on page 5663](#)

Verifying the IEEE 802.1 Multidestination Classifier

Purpose Verify that the classifier **ba-mcast-classifier** is configured as the IEEE 802.1 multidestination classifier:

Action Verify the results of the classifier configuration using the operational mode command **show configuration class-of-service multi-destination classifiers ieee-802.1**:

```
user@switch> show configuration class-of-service multi-destination classifiers ieee-802.1
ba-mcast-classifier;
```

Verifying the Multidestination Classifier Configuration

Purpose Verify that you configured the multidestination classifier with the correct forwarding classes, loss priorities, and code points.

Action List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier
  forwarding-class mcast {
    loss-priority low code-points 000;
  }
```

- Related Documentation**
- [Example: Configuring Unicast Classifiers on page 5658](#)
 - [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)
 - [Monitoring CoS Classifiers on page 5915](#)
 - [Understanding CoS Classifiers on page 5455](#)
 - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

Example: Configuring Tail-Drop Profiles

You can configure an interpolated weighted random early detection (WRED) tail-drop profile to control packet drop characteristics for different traffic loss priorities.



NOTE: You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

- [Requirements on page 5664](#)
- [Overview on page 5664](#)
- [Configuring a Drop Profile on page 5666](#)
- [Verification on page 5666](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

You associate a tail-drop profile with a loss priority in a scheduler. When you attach the scheduler to a forwarding class (queue), you apply the interpolated drop profile to traffic of the specified loss priority in that queue. *Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point:

- Drop start point—Percentage of average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop.
- Drop end point—Average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.)
- Maximum drop rate—Drop probability when the average queue fill level reaches the drop end point.


You set the drop start point and the drop end point by specifying two queue fill level percentage values. The first value is the drop start point and the second value is the drop end point.

You set the maximum drop rate by specifying two drop probability percentage values. The first value is always zero (0), which is the minimum drop rate, the probability of dropping a packet at the drop start point. The second value is the maximum drop rate at the drop end point.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

[Figure 214 on page 5665](#) shows the graph for a drop profile with a drop start point of 30 percent, a drop end point of 50 percent, and a maximum drop rate of 80 percent.

Figure 214: Tail-Drop Profile Packet Drop Example



The graph shows that when the queue fill level is less than 30 percent, the packet drop rate is zero. When the queue fill level reaches 30 percent, packets begin to drop. As the queue fills, a higher percentage of packets drop. When the queue fill level reaches 50 percent, the packet drop rate has climbed to 80 percent. When the queue fill level exceeds 50 percent, all packets drop.

This example describes how to configure the drop profile shown in [Figure 214 on page 5665](#). The drop profile will have:

- The name **be-dp1**
- 30 percent for the drop start point (first **fill-level** setting)
- 50 percent for the drop end point (second **fill-level** setting)
- 0 percent for the minimum drop rate (first **drop-probability** setting)
- 80 percent for the maximum drop rate (second **drop-probability** setting)

You apply a drop profile by configuring a drop profile map that maps the drop profile to a packet loss priority and associates the drop profile and packet loss priority with a scheduler. When you associate the scheduler with a forwarding class (queue), the switch applies the drop profile to the packets in the forwarding class that have a matching packet loss priority.

Configuring a Drop Profile

1. Set the drop start point at 30 percent, the drop end point at 50 percent, the minimum drop rate at 0 percent, and the maximum drop rate at 80 percent for the drop profile **be-dp1**:

```
[edit class-of-service]
user@switch# set drop-profile be-dp1 interpolate fill-level 30 fill-level 50 drop-probability
0 drop-probability 80
```

Verification

Verifying the Drop Profile Configuration

Purpose Verify that you configured the drop profile **be-dp1** with the correct drop start and end points and with the correct drop rates.

Action Verify the results of the drop profile configuration using the operational mode command **show configuration class-of-service drop-profiles be-dp1**:

```
user@switch> show configuration class-of-service drop-profiles be-dp1
interpolate {
    fill-level [ 30 50 ];
    drop-probability [ 0 80 ];
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

Example: Configuring Drop Profile Maps

A drop-profile map associates a tail-drop profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified tail-drop profile to traffic in the forwarding class that matches the specified loss priority.

- [Requirements on page 5666](#)
- [Overview on page 5667](#)
- [Configuring a Drop Profile Map on page 5667](#)
- [Verification on page 5667](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

The following example describes how to configure a drop profile map for a scheduler named **mylan** that includes:

- A drop profile called **lp-profile** for low-priority traffic
- A drop profile called **mh-profile** for medium-high priority traffic
- A drop profile called **h-profile** for high-priority traffic

You apply the drop profiles in the drop profile map to a forwarding class by associating the scheduler **mylan** with a forwarding class in a scheduler map.

Configuring a Drop Profile Map

1. Configure the drop profile for low-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority low protocol any
drop-profile lp-profile
```

2. Configure the drop profile for medium-high priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority medium-high protocol
any drop-profile mh-profile
```

3. Configure the drop profile for high-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority high protocol any
drop-profile h-profile
```

Verification

Verifying the Drop Profile Map Configuration

Purpose Verify that you configured the drop profile map for the scheduler **mylan** with the correct loss priorities and drop profiles.

Action Verify the results of the drop profile map configuration using the operational mode command **show configuration class-of-service schedulers mylan**:

```
user@switch> show configuration class-of-service schedulers mylan
transmit-rate 3g;
shaping-rate percent 100;
priority low;
drop-profile-map loss-priority low protocol any drop-profile lp-profile;
drop-profile-map loss-priority medium-high protocol any drop-profile mh-profile;
drop-profile-map loss-priority high protocol any drop-profile h-profile;
```



NOTE: This example does not include configuring scheduler bandwidth and priority. This information (transmit rate, shaping rate, and priority) is shown for completeness.

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

Example: Configuring Forwarding Classes

Forwarding classes allow you to group packets for transmission. You assign packets to unicast or multdestination (multicast, broadcast, and destination lookup fail) output queues based on forwarding classes.

- [Requirements on page 5668](#)
- [Overview on page 5668](#)
- [Configuring Forwarding Classes on page 5670](#)
- [Verification on page 5670](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The switch supports a total of 12 forwarding classes. In order to forward traffic, you must map (assign) the forwarding classes to unicast or multdestination output queues. The switch has 12 queues. Queues 0 through 7 are for unicast traffic, and queues 8 through 11 are for multdestination traffic. The switch supports up to two lossless forwarding classes.

By default, four categories of unicast forwarding classes and one multdestination forwarding class are defined. You can define the remaining seven forwarding classes and configure them as unicast or multdestination by mapping them to unicast or multdestination queues. The type of queue, unicast or multdestination, determines the type of forwarding class.

The four default unicast forwarding classes are:

- **be**—Best-effort traffic
- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **nc**—Network control traffic

The default multidestination forwarding class is:

- **mcast**—Multidestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue. The statement format is:

```
[edit class-of-service forwarding-classes]
user@switch# class class-name queue-num queue-number;
```



NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.



NOTE: Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

```
[edit class-of-service forwarding-classes]
user@switch# queue queue-number class-name
```

The **queue** statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the **queue** statement to map forwarding classes to queues, edit the configuration to replace the **queue** statement with the **class** statement.



NOTE: The QFX Series uses hierarchical scheduling to control output queue forwarding. When you define a forwarding class that will carry traffic on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

Table 530 on page 5670 shows the configuration forwarding-class-to-queue mapping for this example:

Table 530: Forwarding-Class-to-Queue Example Configuration

Forwarding Class	Queue
best-effort	0
nc	7
mcast	8

Configuring Forwarding Classes

To configure CoS forwarding classes, map the forwarding classes to queues:

1. Map the **best-effort** forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]
user@switch# set class best-effort queue-num 0
```

2. Map the **nc** forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]
user@switch# set class nc queue-num 7
```

3. Map the **mcast-be** forwarding class to queue 8:

```
[edit class-of-service forwarding-classes]
user@switch# set class mcast-be queue-num 8
```

Verification

Verifying the Forwarding-Class-to-Queue Mapping

Purpose Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as **fcoe** and **no-loss**.)

Action Verify the results of the forwarding class configuration using the operational mode command **show configuration class-of-service forwarding-classes**:

```
user@switch> show configuration class-of-service forwarding-classes
class best-effort queue-num 0;
class network-control queue-num 7;
class mcast queue-num 8;
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Defining CoS Forwarding Classes on page 5787](#)
- [Monitoring CoS Forwarding Classes on page 5916](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 5420](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS Forwarding Classes on page 5469](#)

Example: Configuring Forwarding Class Sets

A forwarding class set (fc-set) is a priority group for enhanced transmission selection (ETS) traffic control. Each fc-set consists of one or more forwarding classes (output queues).

ETS enables you to configure link resources (bandwidth and bandwidth sharing characteristics) for a priority group, and then allocate the priority group's resources among the forwarding classes that belong to the priority group. This is called two-tier, or hierarchical, scheduling. Traffic control profiles control the scheduling for the priority group, and schedulers control the scheduling for individual forwarding classes.

- [Requirements on page 5671](#)
- [Overview on page 5671](#)
- [Configuring Forwarding Class Sets on page 5672](#)
- [Verification on page 5673](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

You can configure up to three unicast fc-sets and one multicast fc-set. A common way to configure unicast priority groups is to configure separate fc-sets for local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and then assign the appropriate forwarding classes to each fc-set.



NOTE: If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

To apply ETS, you map one or more fc-sets to a physical egress port. You can map up to three forwarding class sets to each port. When you map an fc-set to a port, the port uses hierarchical scheduling to allocate port resources to the priority group (fc-set) and to allocate the priority group resources to the queues (forwarding classes) that belong to the priority group.

This example describes how to:

- Configure three fc-sets called **lan-pg**, **san-pg**, and **hpc-pg**.
- Assign forwarding classes to each of the fc-sets.
- Apply the fc-sets and their output traffic control profiles to an egress interface.

This example does not describe how to configure the forwarding classes assigned to the fc-sets or how to configure traffic control profiles. [Table 531 on page 5672](#) shows the configuration components for this example:

Table 531: Components of the Forwarding Class Sets Configuration Example

Component	Settings
Hardware	QFX3500 switch
LAN traffic priority group	Forwarding class set: lan-pg Forwarding classes: best-effort-1 , best-effort-2
SAN traffic priority group	Forwarding class set: san-pg Forwarding classes: fcoe , fcoe-2
HPC traffic priority group	Forwarding class set: hpc-pg Forwarding classes: nc , high-perf
Egress interface	xe-0/0/7

Configuring Forwarding Class Sets

1. Define the **lan-pg** priority group (fc-set) and assign to it the forwarding classes **best-effort-1** and **best-effort-2**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets lan-pg class best-effort-1
user@switch# set forwarding-class-sets lan-pg class best-effort-2
```
2. Define the **san-pg** priority group and assign to it the forwarding classes **fcoe** and **fcoe-2**:

```
[edit class-of-service]
```

```

user@switch# set forwarding-class-sets san-pg class fcoe
user@switch# set forwarding-class-sets san-pg class fcoe-2

```

3. Define the **hpc-pg** priority group and assign to it the forwarding classes **nc** and **high-perf**:

```

[edit class-of-service]
user@switch# set forwarding-class-sets hpc-pg class nc
user@switch# set forwarding-class-sets hpc-pg class high-perf

```

4. Map the three forwarding class sets to an interface (the output traffic control profiles associated with the forwarding class sets determine the class of service scheduling for the priority groups):

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile lan-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set san-pg
output-traffic-control-profile san-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp

```

Verification

To verify the priority group configuration, perform these tasks:

- [Verifying Forwarding Class Set Membership on page 5673](#)
- [Verifying the Egress Interface Configuration on page 5673](#)

Verifying Forwarding Class Set Membership

Purpose Verify that you configured the **lan-pg**, **san-pg**, and **hpc-pg** priority groups with the correct forwarding classes.

Action List the forwarding class set member configuration using the operational mode command **show configuration class-of-service forwarding-class-sets**:

```

user@switch> show configuration class-of-service forwarding-class-sets
lan-pg {
    class best-effort-1;
    class best-effort-2;
}
san-pg {
    class fcoe;
    class fcoe-2;
}
hpc-pg {
    class high-perf;
    class nc;
}

```

Verifying the Egress Interface Configuration

Purpose Verify that egress interface **xe-0/0/7** is associated with the **lan-pg**, **san-pg**, and **hpc-pg** priority groups and with the correct output traffic control profiles.

Action Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
  lan-pg {
    output-traffic-control-profile lan-tcp;
  }
  san-pg {
    output-traffic-control-profile san-tcp;
  }
  hpc-pg {
    output-traffic-control-profile hpc-tcp;
  }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Example: Configuring Queue Schedulers on page 5674](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
 - [Defining CoS Forwarding Class Sets on page 5789](#)
 - [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)

Example: Configuring Queue Schedulers

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, and the tail-drop profiles associated with the queue.

- [Requirements on page 5674](#)
- [Overview on page 5674](#)
- [Configuring a CoS Scheduler on page 5677](#)
- [Verification on page 5678](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Scheduler parameters define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each queue mapped to the scheduler receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue. You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.



NOTE: The **transmit-rate** setting works only if you also configure the **guaranteed-rate** in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the **guaranteed-rate**, the **transmit-rate** does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate. If you configure transmit rates whose sum exceeds the forwarding class set guaranteed rate, the commit check fails, and the system rejects the configuration.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each queue receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:
 - **low**—The scheduler has low priority.
 - **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.
- **drop-profile-map**—Mapping of a drop profile to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.



NOTE: Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



NOTE: Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Scheduler maps associate schedulers with forwarding classes (queues). After defining schedulers and mapping them to queues in a scheduler map, to configure hardware queue scheduling (port scheduling) you:

1. Associate a scheduler map with a traffic control profile (a traffic control profile schedules resources for a group of forwarding classes, called a *forwarding class set* or *priority group*).
2. Attach a forwarding class and a traffic control profile to an interface.

You can associate up to four user-defined scheduler maps with forwarding class sets.

This process configures the hardware queues, packet schedulers, and WRED characteristics that operate according to the scheduler mapping. The traffic control profile uses the scheduler CoS properties to determine the resources that should be allocated to the individual output queues from the total resources available to the priority group.

[Table 532 on page 5676](#) shows the configuration components for this example.

Table 532: Components of the Queue Scheduler Configuration Example

Component	Settings
Hardware	QFX3500 switch

Table 532: Components of the Queue Scheduler Configuration Example (*continued*)

Component	Settings
Scheduler	Name: be-sched Transmit rate: 20% Shaping rate: 40% Buffer size: 20% Priority: low Drop profile: be-dp
Scheduler map	Name: be-map Forwarding class to associate with the be-sched scheduler: best-effort
Traffic control profile	Name: be-tcp NOTE: This topic does not describe how to define a traffic control profile.
Forwarding class set	Name: lan-pg

Configuring a CoS Scheduler

To configure a CoS scheduler using the CLI:

- Create a scheduler (**be-sched**) with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, low priority, and map it to the drop profile **be-dp**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate percent 20
user@switch# set be-sched shaping-rate percent 40
user@switch# set be-sched buffer-size percent 20
user@switch# set be-sched priority low
user@switch# set be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
```
- Configure a scheduler map (**be-map**) that associates the scheduler (**be-sched**) with the forwarding class (**best-effort**):

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```
- Associate the scheduler map **be-map** with a traffic control profile (**be-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```
- Associate the traffic control profile **be-tcp** with a forwarding class set (**lan-pg**) and a 10-Gigabit Ethernet interface (**xe-0/0/7**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile be-tcp
```
- Alternatively, you can assign the scheduler map (**be-map**) to all the 10-Gigabit Ethernet interfaces using wildcards (**xe-***):

```
[edit class-of-service interfaces]
user@switch# set xe-* forwarding-class-set lan-pg output-traffic-control-profile be-tcp
```

Verification

To verify that the queue scheduler has been created and is mapped to the correct interfaces, perform these tasks:

- [Verifying the Scheduler Configuration on page 5678](#)
- [Verifying the Scheduler Map Configuration on page 5678](#)
- [Verifying That the Scheduler Is Associated with the Interface on page 5679](#)

Verifying the Scheduler Configuration

Purpose Verify that the queue scheduler **be-sched** has been created with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, the priority set to **low**, and the drop profile **be-dp**.

Action Display the scheduler using the operational mode command **show configuration class-of-service schedulers be-sched**:

```
user@switch> show configuration class-of-service schedulers be-sched
transmit-rate percent 20;
shaping-rate percent 40;
buffer-size percent 20
priority low;
drop-profile-map loss-priority low protocol any drop-profile be-dp;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** has been created and associates the forwarding class **best-effort** with the scheduler **be-sched**, and also that the scheduler map is attached to the traffic control profile **be-tcp**.

Action Display the scheduler map using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Display the traffic control profile to verify that the scheduler map **be-map** is attached using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```



NOTE: This topic does not describe how to configure a traffic control profile or its allocation of port bandwidth. Using a traffic control profile to configure the port resource allocation to the priority group is necessary to implement hierarchical scheduling.

Verifying That the Scheduler Is Associated with the Interface

Purpose Verify that the forwarding class set (**lan-pg**) and the traffic control profile (**be-tcp**) that are associated with the queue scheduler are attached to the interface **xe-0/0/7**.

Action List the interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    lan-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
 - [Example: Configuring Maximum Output Bandwidth on page 5689](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
 - [Example: Configuring Tail-Drop Profiles on page 5663](#)
 - [Defining CoS Queue Schedulers on page 5789](#)
 - [Monitoring CoS Scheduler Maps on page 5919](#)
 - [Understanding CoS Output Queue Schedulers on page 5486](#)
 - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)
 - [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 5591](#)
 - [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 5595](#)
 - [Understanding CoS Buffer Configuration on page 5527](#)

Example: Configuring Queue Scheduling Priority

You can configure the bandwidth scheduling priority of individual queues by specifying the priority in a scheduler, and then using a scheduler map to associate the scheduler with a queue.

- [Requirements on page 5680](#)
- [Overview on page 5680](#)
- [Configuring Queue Scheduling Priority on page 5681](#)
- [Verification on page 5681](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Queues can have one of two bandwidth priorities:

- **strict-high**—You can configure only one queue as a strict-high or high-priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.



NOTE: If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

- **low**—Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.

Table 533 on page 5680 shows the configuration components for this example.

This example describes how to set the queue priority for two forwarding classes (queues) named **fcoe** and **no-loss**. Both queues have a priority of **low**. The scheduler for the **fcoe** queue is named **fcoe-sched** and the scheduler for the **no-loss** queue is named **nl-sched**. One scheduler map, **schedmap1**, associates the schedulers to the queues.

Table 533: Components of the Queue Scheduler Priority Configuration Example

Component	Settings
Hardware	QFX3500 switch
Schedulers	fcoe-sched for FCoE traffic nl-sched for no-loss traffic
Priority	low for FCoE traffic low for no-loss traffic
Scheduler map	schedmap1 : FCoE mapping: scheduler fcoe-sched to forwarding class fcoe No-loss mapping: scheduler nl-sched to forwarding class no-loss

Configuring Queue Scheduling Priority

To configure queue priority using the CLI:

1. Create the FCoE scheduler with **low** priority:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low
```

2. Create the no-loss scheduler with **low** priority:

```
[edit class-of-service]
user@switch# set schedulers nl-sched priority low
```

3. Associate the schedulers with the desired queues in the scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

Verification

To verify that you configured the queue scheduling priority for bandwidth and mapped the schedulers to the correct forwarding classes, perform these tasks:

- [Verifying the Queue Scheduling Priority on page 5681](#)
- [Verifying the Scheduler-to-Forwarding-Class Mapping on page 5681](#)

Verifying the Queue Scheduling Priority

Purpose Verify that you configured the queue schedulers **fcoe-sched** and **nl-sched** with **low** queue scheduling priority.

Action Display the **fcoe-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers fcoe-sched priority**:

```
user@switch> show configuration class-of-service schedulers fcoe-sched priority
priority low;
```

Display the **nl-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers nl-sched priority**:

```
user@switch> show configuration class-of-service schedulers nl-sched priority
priority low;
```

Verifying the Scheduler-to-Forwarding-Class Mapping

Purpose Verify that you configured the scheduler map **schedmap1** to map scheduler **fcoe-sched** to forwarding class **fcoe** and schedule **nl-sched** to forwarding class **no-loss**.

Action Display the scheduler map **schedmap1** using the operational mode command **show configuration class-of-service scheduler-maps schedmap1**:

```
user@switch> show configuration class-of-service scheduler-maps schedmap1
forwarding-class fcoe scheduler fcoe-sched;
forwarding-class no-loss scheduler nl-sched;
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Defining CoS Queue Scheduling Priority on page 5792](#)
- [Monitoring CoS Scheduler Maps on page 5919](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

- [Requirements on page 5682](#)
- [Overview on page 5682](#)
- [Configuring a Traffic Control Profile on page 5684](#)
- [Verification on page 5684](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each priority group receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



NOTE: In order for the **transmit-rate** option (minimum bandwidth for a queue that you set using scheduler configuration) to work properly, you must configure the **guaranteed-rate** for the priority group. If a priority group does not have a guaranteed minimum bandwidth, the queues (forwarding classes) that belong to the priority group cannot have a guaranteed minimum bandwidth.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each priority group receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



NOTE: Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

For example, if you map three priority groups to a 10-Gigabit Ethernet port, the queues associated with all three of the priority groups share the 10-Gbps bandwidth as defined by the traffic control profiles. Therefore, the total combined guaranteed-rate value of the three priority groups should not exceed 10 Gbps. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue (forwarding class) transmit rates cannot exceed the total guaranteed-rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.

If you configure the guaranteed-rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Configuring a Traffic Control Profile

This example describes how to configure a traffic control profile named **san-tcp** with a scheduler map named **san-map1** and allocate to it a minimum bandwidth of 4 Gbps and a maximum bandwidth of 8 Gbps:

1. Create the traffic control profile and set the **guaranteed-rate** (minimum guaranteed bandwidth) to **4g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp guaranteed-rate 4g
```

2. Set the **shaping-rate** (maximum guaranteed bandwidth) to **8g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp shaping-rate 8g
```

3. Associate the scheduler map **san-map1** with the traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp scheduler-map san-map1
```

Verification

Verifying the Traffic Control Profile Configuration

Purpose Verify that the traffic control profile **san-tcp** has been created with a minimum guaranteed bandwidth of 4 Gbps, a maximum bandwidth of 8 Gbps, and the scheduler map **san-map1**.

Action List the traffic control profile using the operational mode command **show configuration class-of-service traffic-control-profiles san-tcp**:

```
user@switch> show configuration class-of-service traffic-control-profiles san-tcp
scheduler-map san-map1;
shaping-rate percent 8g;
guaranteed-rate 4g;
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)

Example: Configuring Minimum Guaranteed Output Bandwidth

Scheduling the minimum guaranteed output bandwidth for a queue (forwarding class) requires configuring both tiers of the two-tier hierarchical scheduler. One tier is scheduling

the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 5685](#)
- [Overview on page 5685](#)
- [Configuring Guaranteed Minimum Bandwidth on page 5687](#)
- [Verification on page 5687](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The priority group minimum guaranteed bandwidth defines the minimum total amount of bandwidth available for all of the queues in the priority group to meet their minimum bandwidth requirements.

The **transmit-rate** setting in the scheduler configuration determines the minimum guaranteed bandwidth for an individual queue. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.

The **guaranteed-rate** setting in the traffic control profile configuration determines the minimum guaranteed bandwidth for a priority group. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



NOTE: You must configure both the **transmit-rate** value for the queue and the **guaranteed-rate** value for the priority group in order to set a valid minimum bandwidth guarantee for a queue. (If the priority group does not have a guaranteed minimum bandwidth, there is no guaranteed bandwidth pool from which the queue can take its guaranteed minimum bandwidth.)

The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



NOTE: You cannot configure minimum guaranteed bandwidth on strict-high priority queues or on a priority group that contains strict-high priority queues.

This example describes how to:

- Configure a transmit rate (minimum guaranteed queue bandwidth) of 2 Gbps for queues in a scheduler named **be-sched**.
- Configure a guaranteed rate (minimum guaranteed priority group bandwidth) of 4 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the minimum guaranteed bandwidth scheduling to the egress interface **xe-0/0/7**.

Table 534 on page 5686 shows the configuration components for this example:

Table 534: Components of the Minimum Guaranteed Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Minimum guaranteed queue bandwidth	Transmit rate: 2g
Minimum guaranteed priority group bandwidth	Guaranteed rate: 4g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort

Table 534: Components of the Minimum Guaranteed Output Bandwidth Configuration Example (*continued*)

Component	Settings
Egress interface	xe-0/0/7

Configuring Guaranteed Minimum Bandwidth

To configure the minimum guaranteed bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the minimum guaranteed queue bandwidth of 2 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate 2g
```

2. Configure the minimum guaranteed priority group bandwidth of 4 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp guaranteed-rate 4g
```

3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Verification

To verify the minimum guaranteed output bandwidth configuration, perform these tasks:

- [Verifying the Minimum Guaranteed Queue Bandwidth on page 5688](#)
- [Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association on page 5688](#)
- [Verifying the Scheduler Map Configuration on page 5688](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 5688](#)
- [Verifying the Egress Interface Configuration on page 5689](#)

Verifying the Minimum Guaranteed Queue Bandwidth

Purpose Verify that you configured the minimum guaranteed queue bandwidth as **2g** in the scheduler **be-sched**.

Action Display the minimum guaranteed bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched transmit-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched transmit-rate
2g;
```

Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association

Purpose Verify that the minimum guaranteed priority group bandwidth is **4g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

Action Display the minimum guaranteed bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate
4g;
```

Display the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action Display the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action Display the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Example: Configuring Queue Schedulers on page 5674](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
 - [Example: Configuring Queue Scheduling Priority on page 5679](#)
 - [Example: Configuring Forwarding Class Sets on page 5671](#)
 - [Understanding CoS Traffic Control Profiles on page 5496](#)
 - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)

Example: Configuring Maximum Output Bandwidth

Scheduling the maximum output bandwidth for a queue (forwarding class) requires configuring both tiers of the hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 5689](#)
- [Overview on page 5689](#)
- [Configuring Maximum Bandwidth on page 5691](#)
- [Verification on page 5691](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The priority group maximum bandwidth defines the maximum total amount of bandwidth available for all of the queues in the priority group.

The **shaping-rate** setting in the scheduler configuration determines the maximum bandwidth for an individual queue.

The **shaping-rate** setting in the traffic control profile configuration determines the maximum bandwidth for a priority group.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



NOTE: When you set the maximum bandwidth (**shaping-rate**) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within ± 20 percent of the configured **shaping-rate** value.

This example describes how to:

- Configure a maximum rate of 4 Gbps for queues in a scheduler named **be-sched**.
- Configure a maximum rate of 6 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the bandwidth scheduling to the interface **xe-0/0/7**.

Table 535 on page 5690 shows the configuration components for this example:

Table 535: Components of the Maximum Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Maximum queue bandwidth	Shaping rate: 4g
Maximum priority group bandwidth	Shaping rate: 6g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp

Table 535: Components of the Maximum Output Bandwidth Configuration Example (*continued*)

Component	Settings
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Configuring Maximum Bandwidth

To configure the maximum bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the maximum queue bandwidth of 4 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched shaping-rate 4g
```

2. Configure the maximum priority group bandwidth of 6 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp shaping-rate 6g
```

3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Verification

To verify the maximum output bandwidth configuration, perform these tasks:

- [Verifying the Maximum Queue Bandwidth on page 5692](#)
- [Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association on page 5692](#)
- [Verifying the Scheduler Map Configuration on page 5692](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 5692](#)
- [Verifying the Egress Interface Configuration on page 5692](#)

Verifying the Maximum Queue Bandwidth

Purpose Verify that you configured the maximum queue bandwidth as **4g** in the scheduler **be-sched**.

Action List the maximum bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched shaping-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched shaping-rate
4g;
```

Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association

Purpose Verify that the maximum priority group bandwidth is **6g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

Action List the maximum bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp shaping-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp shaping-rate
4g;
```

List the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action List the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action List the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action List the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
 - [Example: Configuring Queue Schedulers on page 5674](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
 - [Example: Configuring Forwarding Class Sets on page 5671](#)
 - [Understanding CoS Traffic Control Profiles on page 5496](#)
 - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5481](#)

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your network uses a different priority than 3 for FCoE traffic, you need to configure lossless FCoE transport on that priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic instead of using priority 3.

- [Requirements on page 5693](#)
- [Overview on page 5693](#)
- [Configuration on page 5696](#)
- [Verification on page 5697](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

Overview

Although FCoE traffic typically uses IEEE 802.1p priority 3 on converged Ethernet networks, some networks use a different priority for FCoE traffic. Regardless of the priority used, FCoE traffic must receive lossless treatment. Supporting lossless behavior for FCoE traffic when your network does not use priority 3 requires configuring:

- A lossless forwarding class for FCoE traffic.
- A behavior aggregate (BA) classifier to map the FCoE forwarding class to the appropriate IEEE 802.1p priority.
- A congestion notification profile (CNP) to enable PFC on the FCoE code point at the interface ingress and to configure flow control on the interface egress. Flow control on the interface egress enables the interface to respond to PFC messages received from the connected peer and pause the correct IEEE 802.1p priority on the correct output queue.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- A DCBX application and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priority. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure one lossless FCoE traffic class, map it to a priority other than priority 3, and configure flow control to ensure lossless behavior on the interfaces. This example uses two Ethernet interfaces, xe-0/0/25 and xe-0/0/26. The interfaces connect to a converged Ethernet network that uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic.

The configuration on the two interfaces is the same. Both interfaces use the same explicitly configured lossless FCoE forwarding class and the same ingress classifier. Both interfaces enable PFC on priority 5 and enable flow control on the same output queue (which is mapped to the lossless FCoE forwarding class).

[Table 536 on page 5694](#) shows the configuration components for this example.

Table 536: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3

Component	Settings
Hardware	QFX3500 switch

Table 536: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3 (*continued*)

Component	Settings
Forwarding class	<p>Name—fcoe1</p> <p>Queue mapping—queue 5</p> <p>Packet drop attribute—no-loss</p> <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p>
BA classifier	<p>Name—fcoe_p5</p> <p>FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low.</p>
PFC configuration (CNPs)	<p>CNP name—fcoe_p5_cnp</p> <p>Input CNP code point—101</p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code point—101</p> <p>Output CNP flow control queue—5</p> <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>
DCBX application mapping	<p>Application name—fcoe_p5_app</p> <p>Application EtherType—0x8906</p> <p>Application map name—fcoe_p5_app_map</p> <p>Application map code points—101</p> <p>NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

Configuration

CLI Quick Configuration

To quickly configure a lossless FCoE forwarding class that uses a different priority than IEEE 802.1p priority 3 for FCoE traffic on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map
```

Configuring A Lossless FCoE Forwarding Class On IEEE 802.1p Priority 5

Step-by-Step Procedure

To configure a lossless forwarding class for FCoE traffic on IEEE 802.1p priority 5 (code point 101), classify FCoE traffic into the lossless forwarding class, configure a congestion notification profile to enable PFC on the FCoE priority and output queue, and configure DCBX application protocol TLV exchange for traffic on the FCoE priority:

1. Configure the lossless forwarding class (named **fcoe1** and mapped to output queue 5) for FCoE traffic on IEEE 802.1p priority 5:

```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```

2. Configure the ingress classifier (**fcoe_p5**). The classifier maps the FCoE priority (code point 101) to the lossless FCoE forwarding class **fcoe1**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

3. Apply the classifier to interfaces **xe-0/0/25** and **xe-0/0/26**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
```

4. Configure the CNP. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
```

```

user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5

```

5. Apply the CNP to the interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp

```

6. Configure the DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priority 5 instead of on the default priority 3:

```

[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906

```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priority:

```

[edit]
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101

```

8. Apply the application map to the Ethernet interfaces so that DCBX exchanges FCoE application TLVs on the correct code point:

```

[edit]
user@switch# set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map

```

Verification

To verify the configuration and proper operation of the lossless forwarding class and IEEE 802.1p priority, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 5697](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 5698](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 5698](#)
- [Verifying the Interface Configuration on page 5699](#)
- [Verifying the DCBX Application Configuration on page 5700](#)
- [Verifying the DCBX Application Map Configuration on page 5700](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 5700](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifier configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065
  Code point      Forwarding class      Loss priority
  101             fcoe1                      low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe_p5** maps code point 101 (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**, and all other priorities to the **best-effort** forwarding class with a packet loss priority of **high**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priority and that flow control is configured on the correct output queue in the CNP.

Action Display the congestion notification profile using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_p5_cnp, Index: 12137
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2240
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the configured CNPs.

The **fcoe_p5_cnp** CNP input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

Verifying the Interface Configuration

Purpose Verify that the correct classifier and congestion notification profile are configured on the interfaces.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/25** and **show configuration class-of-service interfaces xe-0/0/26**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/25
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_p5;
  }
}

user@switch> show configuration class-of-service interfaces xe-0/0/26
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_p5;
  }
}
```

Meaning Both the **show configuration class-of-service interfaces xe-0/0/25** command and the **show configuration class-of-service interfaces xe-0/0/26** command show that the

congestion notification profile **fcoe_p5_cnp** is configured on each interface, and that the IEEE 802.1p classifier associated with each interface is **fcoe_p5**.

Verifying the DCBX Application Configuration

Purpose Verify that the DCBX application for FCoE is configured.

Action List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application fcoe_p5_app {
    ether-type 0x8906;
```

Meaning The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe_p5_app** is configured with an EtherType of **0x8906**.

Verifying the DCBX Application Map Configuration

Purpose Verify that the application map is configured.

Action List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
fcoe_p5_app_map {
    application fcoe_p5_app code-points 101;
}
```

Meaning The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe_p5_app_map** consists of the application named **fcoe_p5_app**, which is mapped to IEEE 802.1p code point **101**.

Verifying the DCBX Application Protocol Exchange Interface Configuration

Purpose Verify that the application map is applied to the correct interfaces.

Action List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/25.0 {
    application-map fcoe_p5_app_map;
}
interface xe-0/0/26.0 {
    application-map fcoe_p5_app_map;
}
```

Meaning The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/25.0** and **xe-0/0/26.0** use application map **fcoe_p5_app_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
 - [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
 - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
 - [Example: Configuring Unicast Classifiers on page 5658](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
 - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
 - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your converged Ethernet network uses more than one priority for FCoE traffic, you need to configure lossless transport for each FCoE priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses both priority 3 (IEEE 802.1p code point 011) and priority 5 (IEEE 802.1p code point 101) for FCoE traffic.

- [Requirements on page 5701](#)
- [Overview on page 5701](#)
- [Configuration on page 5704](#)
- [Verification on page 5706](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, a converged Ethernet network might include two separate FCoE networks that use different priorities to identify traffic. Interfaces that carry traffic for both FCoE networks need to support lossless FCoE transport on both priorities.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class).
- A behavior aggregate (BA) classifier to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities).
- A congestion notification profile (CNP) to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifier, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes on an interface, map them to two different priorities, and configure flow control to ensure lossless behavior. This example uses two Ethernet interfaces, xe-0/0/20 and xe-0/0/21, that are connected to the converged Ethernet network. Both interfaces transport FCoE traffic on priorities 3 (011) and 5 (101), and must support lossless transport of that traffic.

[Table 537 on page 5702](#) shows the configuration components for this example.

Table 537: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology

Component	Settings
Hardware	QFX3500 switch

Table 537: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (*continued*)

Component	Settings
Forwarding classes	<p>Name—fcoe1 Queue mapping—queue 5 Packet drop attribute—no-loss</p> <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—fcoe This is the default lossless FCoE forwarding class, so no configuration required. The fcoe forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of no-loss.</p>
BA classifier	<p>Name—fcoe_classifier</p> <p>FCoE priority mapping for forwarding class fcoe—mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low.</p> <p>FCoE priority mapping for forwarding class fcoe1—mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low.</p>
PFC configuration (CNP)	<p>CNP name—fcoe_cnp</p> <p>Input CNP code points—011 and 101</p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code points—011 and 101</p> <p>Output CNP flow control queues—3 and 5</p> <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for PFC pause unless they are included in the explicitly configured output CNP. In this example, because the explicit output CNP overwrites the default output CNP, we must explicitly configure flow control on queue 3.</p>

Table 537: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (continued)

Component	Settings
DCBX application mapping	Application name— fcoe_app Application EtherType— 0x8906 Application map name— fcoe_app_map Application map code points— 011 and 101 NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.
Interfaces	Interfaces xe-0/0/20 and xe-0/0/21 use the same configuration: <ul style="list-style-type: none"> • Classifier—fcoe_classifier • CNP—fcoe_cnp • DCBX application map—fcoe_app_map



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

Configuration

CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 011 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 101 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 011
pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 101
pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps fcoe_app_map application fcoe_app code-points [011 101]
set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

Step-by-Step Procedure To configure two lossless forwarding classes for FCoE traffic on the same interface, classify FCoE traffic into the forwarding classes, configure CNPs to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```



NOTE: This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier. The classifier maps the FCoE priorities (IEEE 802.1p code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101
```

3. Apply the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
```

4. Configure the CNP. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues **3** and **5** on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

5. Apply the CNP to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
```

6. Configure a DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on both of the IEEE 802.1p priorities used for FCoE transport:

```
[edit]
```

```
user@switch# set applications application fcoe_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priorities:

```
[edit]
user@switch# set policy-options application-maps fcoe_app_map application fcoe_app
code-points [011 101]
```

8. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 5706](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 5707](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 5707](#)
- [Verifying the Interface Configuration on page 5708](#)
- [Verifying the DCBX Application Configuration on page 5708](#)
- [Verifying the DCBX Application Map Configuration on page 5709](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 5709](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_classifier, Code point type: ieee-802.1, Index: 10964
  Code point      Forwarding class      Loss priority
  011             fcoe                      low
  101             fcoe1                     low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe_classifier** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities.

Action List the CNPs using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_cnp, Index: 46504
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled   2240
  100      Disabled
  101      Enabled   2240
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  011      3
  101      5
```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the CNP.

The CNP **fcoe_cnp** input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the interface cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

Verifying the Interface Configuration

Purpose Verify that the classifier and congestion notification profile are configured on the interfaces. Both interfaces should show the same configuration.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_classifier;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_classifier;
    }
}
```

Meaning The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_classifier**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_classifier**.

Verifying the DCBX Application Configuration

Purpose Verify that the DCBX application for FCoE is configured.

Action List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application fcoe_app {
    ether-type 0x8906;
```

Meaning The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe_app** is configured with an EtherType of **0x8906**.

Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map is configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
fcoe_app_map {
 application fcoe_app code-points [011 101];
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe\_app\_map** consists of the application named **fcoe\_app**, which is mapped to IEEE 802.1p code points **011** and **101** (priorities 3 and 5, respectively).

**Verifying the DCBX Application Protocol Exchange Interface Configuration**

- Purpose** Verify that the application map is applied to the interfaces.
- Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/20.0 {
    application-map fcoe_app_map;
}
interface xe-0/0/21.0 {
    application-map fcoe_app_map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/20.0** and **xe-0/0/21.0** use application map **fcoe_app_map**.
- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
 - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
 - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
 - [Example: Configuring Unicast Classifiers on page 5658](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
 - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
 - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic if it uses different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and map them to two different priorities on an FCoE transit switch.

- [Requirements on page 5710](#)
- [Overview on page 5710](#)
- [Configuration on page 5714](#)
- [Verification on page 5717](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, when the QFX3500 switch acts as a transit switch, it could be connected to two QFX3500 switches in FCoE-FC gateway mode. Each of the gateway switches could connect a set of FCoE clients to a different SAN, and each set of FCoE clients could use a different priority for FCoE traffic to avoid fate sharing and maintain separation of the two FCoE networks. In this case, you need to configure two forwarding classes for FCoE traffic, each mapped to a different output queue and a different priority.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- Behavior aggregate (BA) classifiers to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes, map them to two different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses three Ethernet interfaces, xe-0/0/20, xe-0/0/21, and xe-0/0/22:

- Interface xe-0/0/20 connects to an FCoE-FC gateway that connects to Fibre Channel (FC) SAN 1. FCoE traffic to and from FC SAN 1 uses the default **fcoe** forwarding class and the default mapping to priority 3 (IEEE 802.1p code point 011) and output queue 3.
- Interface xe-0/0/21 connects to another FCoE-FC gateway that connects to Fibre Channel (FC) SAN 2. FCoE traffic to and from FC SAN-2 uses an explicitly configured FCoE forwarding class that is mapped to priority 5 (code point 101) and output queue 5.
- Interface xe-0/0/22 connects to FCoE devices on the converged Ethernet network and handles traffic destined for FC SAN 1 and FC SAN 2. Interface xe-0/0/22 must properly handle lossless FCoE traffic of both priorities (both FCoE forwarding classes), including pausing the traffic on ingress or egress as required.

Figure 215 on page 5712 shows the topology for this example, and Table 538 on page 5712 shows the configuration components for this example.

Figure 215: Topology of the Two Lossless FCoE Priorities Example




Table 538: Components of the Two Lossless FCoE Priorities Configuration Topology

Component	Settings
Hardware	QFX3500 switch
Forwarding classes	<p>Name—fcoe1 Queue mapping—queue 5 Packet drop attribute—no-loss</p> <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—fcoe This is the default lossless FCoE forwarding class, so no configuration required. The fcoe forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of no-loss</p>
BA classifiers	<p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> Interface xe-0/0/20 classifier: Name—fcoe_p3 FCoE priority mapping—Forwarding class fcoe mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low. Interface xe-0/0/21 classifier: Name—fcoe_p5 FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low. Interface xe-0/0/22 classifier: Name—fcoe_p3_p5 FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 and a packet loss priority of low, and forwarding class fcoe mapped to code point 011 and a packet loss priority of low.

Table 538: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)

Component	Settings
PFC configuration (CNPs)	<p>Each interface requires a different CNP because each interface handles a different subset of FCoE traffic and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> Interface xe-0/0/20 CNP: CNP name—fcoe_p3_cnp Input CNP code point—011 MRU—2240 bytes Cable length—100 meters <p>NOTE: Because interface xe-0/0/20 uses the default FCoE configuration, output queue 3 is paused by default and you do not need to configure the output stanza of the CNP.</p> <ul style="list-style-type: none"> Interface xe-0/0/21 CNP: CNP name—fcoe_p5_cnp Input CNP code point—101 MRU—2240 bytes Cable length—150 meters Output CNP code point—101 Output CNP flow control queue—5 Interface xe-0/0/22 CNP: CNP name—fcoe_p3_p5_cnp Input CNP code points—011 and 101 MRU—2240 bytes (both priorities) Cable length—100 meters Output CNP code points—011 (for queue 3) and 101 (for queue 5) Output CNP flow control queues—3 for priority 3 (code point 011) and 5 for priority 5 (code point 101) <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>

Table 538: Components of the Two Lossless FCoE Priorities Configuration Topology (continued)

Component	Settings
DCBX application mapping	<p>Interface xe-0/0/20 does not need an application map because DCBX exchanges application protocol TLVs only on the default FCoE priority (priority 3).</p> <p>Interface xe-0/0/21 requires an application map that enables DCBX application protocol TLV exchange on priority 5 (code point 101) for FCoE traffic. Interface xe-0/0/22 requires an application map that enables DCBX application protocol TLV exchange both on priority 3 (code point 011) and on priority 5 (code point 101) for FCoE traffic.</p> <ul style="list-style-type: none"> Interface xe-0/0/21 DCBX application mapping: Application name—fcoe_p5_app Application ether-type—0x8906 Application map name—fcoe_p5_app_map Application map code points—101 Interface xe-0/0/22 DCBX application mapping: Application name—fcoe_all_app Application ether-type—0x8906 Application map name—fcoe_all_app_map Application map code points—011 and 101 <p>NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

Configuration

CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set applications application fcoe_all_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set policy-options application-maps fcoe_all_app_map application fcoe_all_app code-points [011
101]
set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on different interfaces, classify FCoE traffic into the forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```



NOTE: This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier (**fcoe_p3**) for interface **xe-0/0/20**. The classifier maps the FCoE priority (IEEE 802.1p code point **011**) to lossless FCoE forwarding class **fcoe**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low code-points
011

```

3. Configure the ingress classifier (**fcoe_p5**) for interface **xe-0/0/21**. The classifier maps the FCoE priority (IEEE 802.1p code point **101**) to lossless FCoE forwarding class **fcoe1**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101

```

4. Configure the ingress classifier (**fcoe_p3_p5**) for interface **xe-0/0/22**. The classifier maps the two FCoE priorities (IEEE 802.1p code points **011** and **101**) to the two lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points 011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points 101
```

5. Apply each classifier to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
```

6. Configure the CNP input stanza for interface **xe-0/0/20** to enable PFC on the FCoE priority (IEEE 802.1p code point **011**), set the MRU value (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queue 3 is paused by default on priority 3, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100
```

7. Configure the CNP for interface **xe-0/0/21**. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point **101**), sets the MRU value (2240 bytes), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point 101 pfc flow-control-queue 5
```

8. Configure the CNP for interface **xe-0/0/22**. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point 101 pfc flow-control-queue 5
```

9. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
```



```
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp
```

10. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/21. Interface xe-0/0/21 uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on priority 5 on interface xe-0/0/21. Configure an application named **fcoe_p5_app** for FCoE traffic (EtherType 0x8906) and configure an application map named **fcoe_p5_app_map** to map the application to code point 101:

```
[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101
```



NOTE: Interface xe-0/0/20 uses the default FCoE configuration (priority 3). DCBX exchanges protocol TLVs for the FCoE application by default, so you do not need to configure DCBX explicitly on interface xe-0/0/20.

11. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/22. Interface xe-0/0/22 uses both priority 3 (IEEE 802.1p code point 011) and priority 5 for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on both priority 3 and priority 5. Configure an application named **fcoe_all_app** for FCoE traffic (EtherType 0x8906) and configure an application map named **fcoe_all_app_map** to map the application to code points 011 and 101:

```
[edit]
user@switch# set applications application fcoe_all_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_all_app_map application
fcoe_all_app code-points [011 101]
```

12. Apply the application maps to the interfaces xe-0/0/21 and xe-0/0/22 so that DCBX exchanges FCoE application TLVs on the correct code points on each interface:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 5718](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 5718](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 5719](#)
- [Verifying the Interface Configuration on page 5721](#)
- [Verifying the DCBX Application Configuration on page 5721](#)
- [Verifying the DCBX Application Map Configuration on page 5722](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 5722](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe_p3, Code point type: ieee-802.1, Index: 13913

Code point	Forwarding class	Loss priority
011	fcoe	low

Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065

Code point	Forwarding class	Loss priority
101	fcoe1	low

Classifier: fcoe_p3_p5, Code point type: ieee-802.1, Index: 10964

Code point	Forwarding class	Loss priority
011	fcoe	low
101	fcoe1	low

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are three classifiers, **fcoe_p3**, **fcoe_p5**, and **fcoe_p3_p5**.

Classifier **fcoe_p3** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**.

Classifier **fcoe_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe_p3_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240

```

100      Disabled
101      Enabled      2240
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
011
3
101
5

Name: fcoe_p5_cnp, Index: 12133
Type: Input
Cable Length: 150 m
Priority  PFC      MRU
000      Disabled
001      Disabled
010      Disabled
011      Disabled
100      Disabled
101      Enabled   2240
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
101
5

```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the three CNPs. For CNP **fcoe_p3_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows the default mapping of priorities to output queues.



NOTE: By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queue 3 responds to pause messages from the connected peer on interfaces that use CNP **fcoe_p3_cnp**, because the input stanza enables PFC priority 3 only.

For CNP **fcoe_p3_p5_cnp**, the input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

For CNP **fcoe_p5_cnp**, the input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **150** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

Verifying the Interface Configuration

Purpose Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20**, **show configuration class-of-service interfaces xe-0/0/21**, and **show configuration class-of-service interfaces xe-0/0/22**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/22
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_p5;
    }
}
```

Meaning The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe_p3_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p5**.

The **show configuration class-of-service interfaces xe-0/0/22** command shows that the congestion notification profile **fcoe_p3_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3_p5**.

Verifying the DCBX Application Configuration

Purpose Verify that the two DCBX applications for FCoE are configured.

Action List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application fcoe_all_app {
    ether-type 0x8906;
```

```
application fcoe_p5_app {  
    ether-type 0x8906;  
}
```

Meaning The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe_all_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic) and that the application **fcoe_p5_app** is also configured with an EtherType of **0x8906**.

Verifying the DCBX Application Map Configuration

Purpose Verify that the application maps are configured.

Action List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps  
fcoe_all_app_map {  
    application fcoe_all_app code-points [011 101];  
}  
fcoe_p5_app_map {  
    application fcoe_p5_app code-points 101;  
}
```

Meaning The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there are two application maps.

Application map **fcoe_all_app_map** consists of the application named **fcoe_all_app** mapped to IEEE 802.1p code points **011** (priority 3) and **101** (priority 5).

Application map **fcoe_p5_app_map** consists of the application named **fcoe_p5_app** mapped to IEEE 802.1p code point **101** (priority 5).

Verifying the DCBX Application Protocol Exchange Interface Configuration

Purpose Verify that the application maps are applied to the correct interfaces.

Action List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx  
interface xe-0/0/21.0 {  
    application-map fcoe_p5_app_map;  
}  
interface xe-0/0/22.0 {  
    application-map fcoe_all_app_map;  
}
```

Meaning The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interface **xe-0/0/21.0** uses application map **fcoe_p5_app_map** and interface **xe-0/0/22.0** uses application map **fcoe_all_app_map**.



NOTE: Because interface xe-0/0/20 uses the default lossless FCoE configuration, you do not configure application mapping to interface xe-0/0/20. The default configuration automatically exchanges application protocol TLVs for the default FCoE configuration on priority 3 (IEEE 802.1p code point 011).

Related Documentation

- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic on different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, and map the forwarding classes to three different priorities. (The converged Ethernet network includes two FCoE networks, each of which uses a different priority to identify FCoE traffic, and an iSCSI network.)

- [Requirements on page 5723](#)
- [Overview on page 5724](#)
- [Configuration on page 5728](#)
- [Verification on page 5732](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode

- Junos OS Release 12.3 or later for the QFX Series

Overview

Some converged Ethernet networks support FCoE on more than one IEEE 802.1p priority and also require supporting other lossless traffic classes. Interfaces that carry multiple lossless forwarding classes need to support lossless behavior for the priorities mapped to those forwarding classes. To support the two FCoE forwarding classes and the iSCSI forwarding class used in this example, you need to configure:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- A lossless forwarding class for iSCSI traffic
- Behavior aggregate (BA) classifiers to map the lossless forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE and iSCSI code points at the interface ingress, and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the FCoE and iSCSI traffic on the configured lossless priorities. By default, DCBX is enabled on all Ethernet interfaces for FCoE, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure two lossless FCoE traffic classes and one lossless iSCSI traffic class, map them to three different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses four Ethernet interfaces, xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34:

- Interface xe-0/0/31 handles FCoE traffic on priority 3 (IEEE 802.1p code point 011) and iSCSI traffic on priority 4 (code point 100).
- Interface xe-0/0/32 handles FCoE traffic on priority 5 (code point 101) and iSCSI traffic on priority 4.
- Interface xe-0/0/33 handles FCoE traffic on priority 3 and priority 5.
- Interface xe-0/0/34 handles iSCSI traffic on priority 4.

Figure 216 on page 5725 shows the topology for this example, and Table 539 on page 5725 shows the configuration components for this example.

Figure 216: Topology of the Lossless FCoE and iSCSI Priorities Example




Table 539: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology

Component	Settings
Hardware	QFX3500 switch

Table 539: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)

Component	Settings
Forwarding classes	<p>This example uses one explicitly configured lossless FCoE forwarding class, the default lossless FCoE forwarding class, and one explicitly configured iSCSI forwarding class.</p> <ul style="list-style-type: none"> iSCSI forwarding class: Name—iscsi Queue mapping—queue 4 Packet drop attribute—no-loss FCoE forwarding class (explicitly configured): Name—fcoe1 Queue mapping—queue 5 Packet drop attribute—no-loss <p>NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <ul style="list-style-type: none"> FCoE forwarding class (default) Name—fcoe The default fcoe forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of no-loss.
BA classifiers	<p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> Interface xe-0/0/31 classifier: Name—fcoe_p3_iscsi FCoE priority mapping—Forwarding class fcoe mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low. iSCSI priority mapping—Forwarding class iscsi mapped to code point 100 (priority 4) and a packet loss priority of low. Interface xe-0/0/32 classifier: Name—fcoe_p5_iscsi FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low. iSCSI priority mapping—Forwarding class iscsi mapped to code point 100 (priority 4) and a packet loss priority of low. Interface xe-0/0/33 classifier: Name—fcoe_p3_p5 FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (priority 5) and a packet loss priority of low, and forwarding class fcoe mapped to code point 011 and a packet loss priority of low. Interface xe-0/0/34 classifier: Name—iscsi_classifier iSCSI priority mapping—Forwarding class iscsi mapped to code point 100 (priority 4) and a packet loss priority of low.

Table 539: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)

Component	Settings
PFC configuration (CNPs)	<p>Each interface requires a different CNP because each interface handles a different subset of FCoE and iSCSI traffic, and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> Interface xe-0/0/31 CNP: CNP name—fcoe_p3_cnp Input CNP code points—011 and 100 MRU—2240 bytes for code point 011, default value (2500 bytes) for code point 100 Cable length—100 meters <p>NOTE: On interface xe-0/0/31, the FCoE forwarding class is mapped to queue 3 and priority 3 (code point 011), and the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Therefore, interface xe-0/0/31 does not require an output CNP configuration because queue 3 and queue 4 are enabled for PFC flow control by default on code points 011 and 100, respectively.</p> <ul style="list-style-type: none"> Interface xe-0/0/32 CNP: CNP name—fcoe_p5_cnp Input CNP code points—100 and 101 MRU—Default value (2500 bytes) for code point 100, 2240 bytes for code point 101 Cable length—150 meters Output CNP code points—100 and 101 Output CNP flow control queues—4 and 5 Interface xe-0/0/33 CNP: CNP name—fcoe_p3_p5_cnp Input CNP code points—011 and 101 MRU—2240 bytes (both priorities) Cable length—100 meters Output CNP code points—011 and 101 Output CNP flow control queues—3 and 5 Interface xe-0/0/34 CNP: CNP name—iscsi_cnp Input CNP code point—100 MRU—2500 bytes (default value) Cable length—100 meters <p>NOTE: On interface xe-0/0/34, the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Interface xe-0/0/34 does not require an output CNP configuration because queue 4 is enabled for PFC flow control by default on code point 100.</p> <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>

Table 539: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)

Component	Settings
DCBX application mapping	<p>This example requires configuring applications for FCoE and iSCSI, including them in the same application map, and applying the application map to all four interfaces.</p> <p>Application map name—dcbx_iscsi_fcoe_app_map</p> <ul style="list-style-type: none"> FCoE application name—fcoe_app Application ether-type—0x8906 Application map code points—011 and 101 iSCSI application name—iscsi_app Application protocol type—tcp Application destination port—3260 Application map code point—100 <p>NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

Configuration

CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes and one lossless iSCSI forwarding class and map them to different priorities, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class iscsi queue-num 4 no-loss
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100
set class-of-service interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
set class-of-service interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
set class-of-service interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5set
class-of-service interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile iscsi_cnp input ieee-802.1 code-point 100 pfc
set class-of-service congestion-notification-profile iscsi_cnp input cable-length 100
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
set class-of-service interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
set applications application iscsi_app protocol tcp destination-port 3260
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps dcbx_iscsi_fcoe_app_map application iscsi_app code-points
100
set policy-options application-maps dcbx_iscsi_fcoe_app_map application fcoe_app code-points
[011 101]
set protocols dcbx interface xe-0/0/31 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/32 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/33 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/34 application-map dcbx_iscsi_fcoe_app_map

```

Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, classify the traffic into the three forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding classes **iscsi** for iSCSI traffic and **fcoe1** for FCoE traffic (this example uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class) and map them to output queues:

```

[edit class-of-service]
user@switch# set forwarding-classes class iscsi queue-num 4 no-loss
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```

2. Configure the ingress classifier (**fcoe_p3_iscsi**) for interface **xe-0/0/31**. The classifier maps the FCoE priority (code point **011**) to lossless FCoE forwarding class **fcoe** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]

```

```

user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100

```

3. Configure the ingress classifier (**fcoe_p5_iscsi**) for interface **xe-0/0/32**. The classifier maps the FCoE priority (code point **101**) to lossless FCoE forwarding class **fcoe1** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101

```

4. Configure the ingress classifier (**fcoe_p3_p5**) for interface **xe-0/0/33**. The classifier maps the two FCoE priorities (code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101

```

5. Configure the ingress classifier (**iscsi_classifier**) for interface **xe-0/0/34**. The classifier maps the iSCSI priority (code point **101**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100

```

6. Apply each classifier to the appropriate interface:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier

```

7. Configure the CNP input stanza for interface **xe-0/0/31** to enable PFC on the FCoE and iSCSI priorities that the interface handles (code points **011** and **100**), set the MRU value for the FCoE traffic (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queues 3 and 4 are paused by default on priorities 3 and 4, respectively, and we are not explicitly configuring output queue flow control for any other queues.

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100

```

8. Configure the CNP for interface **xe-0/0/32**. The input stanza enables PFC on the FCoE priority (code point **101**), sets the MRU value for FCoE traffic (2240 bytes), enables PFC on the iSCSI priority (code point **100**), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority and on output queue 4 on the iSCSI priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

9. Configure the CNP for interface xe-0/0/33. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

10. Configure the CNP input stanza for interface xe-0/0/34 to enable PFC on the iSCSI priority (code point 100) and set the cable length value (100 meters). No output stanza is needed because queue 4 is paused by default on priority 4, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile iscsi_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile iscsi_cnp input cable-length 100
```

11. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
```

12. Configure the DCBX applications for FCoE and iSCSI to map to the interfaces so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priorities used for FCoE and iSCSI traffic:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
user@switch# set applications application iscsi_app protocol tcp destination-port 3260
```

13. Configure a DCBX application map to map the FCoE and iSCSI applications to the correct priorities:

```
[edit]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
fcoe_app code-points [011 101]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
iscsi_app code-points 100
```

14. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/31 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/32 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/33 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/34 application-map
dcbx_iscsi_fcoe_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 5732](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 5733](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 5733](#)
- [Verifying the Interface Configuration on page 5736](#)
- [Verifying the DCBX Application Configuration on page 5737](#)
- [Verifying the DCBX Application Map Configuration on page 5737](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 5738](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding classes **iscsi** and **fcoe1** have been created and that the default lossless forwarding class **fcoe** is still enabled for lossless transport.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
iscsi	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **iscsi** and **fcoe1** forwarding classes are configured on output queues 4 and 5, respectively, with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default **fcoe** forwarding class, it remains in its default state (lossless configuration).

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the four classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p3_iscsi, Code point type: ieee-802.1, Index: 13915
  Code point  Forwarding class  Loss priority
  011         fcoe             low
  100         iscsi            low
```

```
Classifier: fcoe_p5_iscsi, Code point type: ieee-802.1, Index: 62035
  Code point  Forwarding class  Loss priority
  100         iscsi            low
  101         fcoe1           low
```

```
Classifier: fcoe_p3_p5, Code point type: ieee-802.1, Index: 17774
  Code point  Forwarding class  Loss priority
  011         fcoe             low
  101         fcoe1           low
```

```
Classifier: iscsi_classifier, Code point type: ieee-802.1, Index: 31635
  Code point  Forwarding class  Loss priority
  100         iscsi            low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are four classifiers, **fcoe_p3_iscsi**, **fcoe_p5_iscsi**, **fcoe_p3_p5**, and **iscsi_classifier**.

Classifier **fcoe_p3_iscsi** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and code point **100** (priority 4) to explicitly configured lossless forwarding class **iscsi**.

Classifier **fcoe_p5_iscsi** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**, and code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe_p3_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **iscsi_classifier** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Enabled	9216
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	
	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
011	
	3
101	
	5

```
Name: fcoe_p5_cnp, Index: 12133
```

```
Type: Input
```

```
Cable Length: 150 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	

```

011      Disabled
100      Enabled      9216
101      Enabled      2240
110      Disabled
111      Disabled
Type: Output
100
      4
101
      5

Name: iscsi_cnp, Index: 19342
Type: Input
Cable Length: 100 m
Priority  PFC      MRU
000      Disabled
001      Disabled
010      Disabled
011      Disabled
100      Enabled      9216
101      Disabled
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
000
      0
001
      1
010
      2
011
      3
100
      4
101
      5
110
      6
111
      7

```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the four CNPs.

For CNP **fcoe_p3_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3) with an MRU of **2240** bytes, and cable length of **100** meters. The input stanza also shows that PFC is enabled on code point **100** (priority 4) with the default MRU value of **9216** bytes. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.



NOTE: By default, only queues 3 and 4 are enabled respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queues 3 and 4 respond to pause messages from the connected peer on interfaces that use CNP `fcoe_p3_cnp` because the input stanza enables PFC only on priorities 3 and 4.

For CNP `fcoe_p3_p5_cnp`, the input stanza shows that PFC is enabled on code points 011 and 101 (priority 5), the MRU is 2240 bytes on both priorities, and the cable length is 100 meters. The CNP output stanza shows that output flow control is configured on queues 3 and 5 for code points 011 and 101, respectively.

For CNP `fcoe_p5_cnp`, the input stanza shows that PFC is enabled on code points 100 and 101. The MRU for code point 101 (FCoE traffic) is 2240 bytes and the MRU for code point 100 is 9216. The interface cable length is 150 meters. The CNP output stanza shows that output flow control is configured on queue 4 for code point 100 and on queue 5 for code point 101.

For CNP `iscsi_cnp`, the input stanza shows that PFC is enabled on code point 100, the MRU value is 9216 bytes, and the interface cable length is 100 meters. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.

Verifying the Interface Configuration

Purpose Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

Action List the ingress interfaces using the operational mode commands `show configuration class-of-service interfaces xe-0/0/31`, `show configuration class-of-service interfaces xe-0/0/32`, `show configuration class-of-service interfaces xe-0/0/33`, and `show configuration class-of-service interfaces xe-0/0/34`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_iscsi;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5_iscsi;
    }
}
```

```

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_p5;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile iscsi_cnp;
unit 0 {
    classifiers {
        ieee-802.1 iscsi_classifier;
    }
}

```

Meaning The `show configuration class-of-service interfaces xe-0/0/31` command shows that the congestion notification profile `fcoe_p3_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_p3_iscsi`.

The `show configuration class-of-service interfaces xe-0/0/32` command shows that the congestion notification profile `fcoe_p5_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_p5_iscsi`.

The `show configuration class-of-service interfaces xe-0/0/33` command shows that the congestion notification profile `fcoe_p3_p5_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_p3_p5`.

The `show configuration class-of-service interfaces xe-0/0/34` command shows that the congestion notification profile `iscsi_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `iscsi_classifier`.

Verifying the DCBX Application Configuration

Purpose Verify that the DCBX applications for FCoE and iSCSI are configured.

Action List the DCBX applications by using the configuration mode command `show applications`:

```

user@switch# show applications
application iscsi_app {
    protocol tcp;
    destination-port 3260;
}
application fcoe_app {
    ether-type 0x8906;
}

```

Meaning The `show applications` configuration mode command shows all of the configured applications. The output shows that the application `iscsi_app` is configured with a protocol value of `tcp` and a destination port value of `3260`, and that the application `fcoe_app` is configured with an EtherType of `0x8906` (the correct EtherType for FCoE traffic).

Verifying the DCBX Application Map Configuration

Purpose Verify that the application map is configured.

Action List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
    application iscsi_app code-points 100;
    application fcoe_app code-points [011 101];
}
```

Meaning The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there is one application map named **dcbx-iscsi-fcoe_app_map**. It consists of the application **iscsi_app** mapped to code point **100** and the application **fcoe_app** mapped to code points **011** and **101**.

Verifying the DCBX Application Protocol Exchange Interface Configuration

Purpose Verify that the application maps are applied to the correct interfaces.

Action List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/31.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/32.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/33.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/34.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
```

Meaning The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that all four interfaces use the application map **dcbx-iscsi-fcoe-app-map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
 - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
 - [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
 - [Example: Configuring Unicast Classifiers on page 5658](#)
 - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
 - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

FCoE traffic typically uses IEEE 802.1p priority 3 (code point 011). However, if your FCoE network uses a different IEEE 802.1p priority than priority 3 for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that priority. You specify the lossless forwarding class used for the FCoE traffic by configuring a fixed classifier and applying it to the native FC (NP_Port) interface. All traffic received from the FC SAN on that NP_Port interface is classified into the forwarding class specified in the fixed classifier.

When native FC interfaces on the FCoE-FC gateway encapsulate incoming FC traffic in Ethernet to create FCoE frames, by default they assign IEEE 802.1p code point 011 to the FCoE traffic, forward the traffic internally to the gateway Ethernet interfaces, and then forward the traffic to the FCoE network. Setting a rewrite value for the IEEE 802.1p code point configures the gateway native FC interface to assign the rewrite value priority to the FCoE frames when the native FC interface forwards the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

You can configure one rewrite value per native FC interface. Each native FC interface can use a different rewrite value.

This example shows how to configure FCoE priority remapping for a converged Ethernet network that uses priority 5 (IEEE code point 101) for FCoE traffic. If your network uses priority 3 for FCoE traffic, then you do not need to remap the FCoE priority, because the default configuration supports lossless FCoE transport on priority 3.

- [Requirements on page 5739](#)
- [Overview on page 5739](#)
- [Configuration on page 5743](#)
- [Verification on page 5744](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

Overview

Native FC interfaces on an FCoE-FC gateway receive native FC traffic from the FC SAN and encapsulate it in Ethernet to create FCoE frames. Priority remapping enables you to map the encapsulated FC traffic (the FCoE traffic) to any IEEE 802.1p priority.

To support lossless FCoE traffic flows, you must configure the remapped priority correctly on the native FC interfaces and also on the Ethernet interfaces that connect to the FCoE

network. Achieving lossless behavior for FCoE traffic when you remap the FCoE priority requires configuring:

- A lossless forwarding class for FCoE traffic (or using the default **fcoe** forwarding class)
- A behavior aggregate (BA) classifier on the FCoE Ethernet interfaces to map the FCoE forwarding class to the IEEE 802.1p code points (priority) used for FCoE traffic on the FCoE network (the ingress classifier priority for the forwarding class must be the same as the rewrite value priority)
- A fixed classifier on the FCoE-FC gateway FC interface that maps all traffic from the FC network into the lossless FCoE forwarding class (the forwarding class must be lossless)
- A priority rewrite value that remaps the IEEE 802.1p code point on the FCoE-FC gateway FC interface to the priority used for FCoE traffic on the FCoE network
- An input congestion notification profile (CNP) to enable priority-based flow control (PFC) on the FCoE code point (the code point used as the rewrite value) at the Ethernet interface ingress and an output CNP to configure flow control to pause the correct output queue at the Ethernet interface egress



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- A DCBX application and application map on the Ethernet interface to support DCBX application TLV exchange for the lossless FCoE traffic on the FCoE priority

The priority specified in the BA classifier, CNP, and DCBX application map on the Ethernet ingress interfaces must match the priority specified in the fixed classifier and rewrite value configurations on the FC interfaces. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure priority remapping of FCoE traffic on one native FC interface (fc-0/0/2) connected to the FC SAN and on one Ethernet interface (xe-0/0/27) connected to the converged Ethernet (FCoE) network. Both the native FC interface and the Ethernet interface belong to the same local FC fabric on the FCoE-FC gateway.

The converged Ethernet network uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic. The native FC interface on the FCoE-FC gateway receives FC traffic from the FC SAN. The native FC interface encapsulates the FC traffic in Ethernet to create FCoE frames, tags the frames with the IEEE 802.1p priority value 101, and then forwards the FCoE frames to the FCoE-FC gateway Ethernet interface. Because traffic marked with IEEE 802.1p priority 5 is mapped to a lossless FCoE forwarding class, the traffic receives

lossless treatment. The Ethernet interface forwards the FCoE traffic on to the Ethernet network.

FCoE traffic (tagged with priority 5) arriving at the FCoE-FC gateway from the Ethernet network receives lossless treatment and is forwarded to the native FC interface. The native FC interface removes the Ethernet encapsulation from the FCoE frames and forwards the resulting native FC traffic to the FC SAN.

Figure 217 on page 5741 shows the topology for this example, and Table 540 on page 5741 shows the configuration components for this example.

Figure 217: Topology of the IEEE 802.1p Priority Remapping Example




Table 540: Components of the IEEE 802.1p Priority Remapping Configuration Topology

Component	Settings
Hardware	QFX3500 switch
Forwarding class configuration	Name— fcoe1 Queue mapping—queue 5 Packet drop attribute— no-loss NOTE: The lossless forwarding class can be mapped to any output queue. However, because FCoE uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.
BA classifier (Ethernet interface)	Name— fcoe_gw_classifier Maps code point 101 (IEEE 802.1p priority 5) to the fcoe1 forwarding class and assigns traffic a packet loss priority of low . The classifier is applied to Ethernet interface xe-0/0/27 .
Fixed classifier (native FC interface)	Forwarding class— fcoe1 The classifier is applied to native FC interface fc-0/0/2
Rewrite value	IEEE 802.1p code point— 101 The rewrite value is applied to native FC interface fc-0/0/2

Table 540: Components of the IEEE 802.1p Priority Remapping Configuration Topology (*continued*)

Component	Settings
PFC configuration (CNP on Ethernet interface)	Name— fcoe1_p5_rewrite_cnp Input CNP code point— 101 Output CNP code point— 101 Output CNP flow control queue— 5 Interface— xe-0/0/27
DCBX application mapping	Application name— myfcoe5 Application ether-type— 0x8906 Application map name— myfcoe5_map Application map code points— 101 Interface— xe-0/0/27 NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.

The priority used to identify FCoE traffic (5, IEEE 802.1p code point 101) is configured for lossless transport across the QFX device on interfaces xe-0/0/27 and fc-0/0/2, which belong to the same local FC fabric on the FCoE-FC gateway.

On the Ethernet interface, the classifier maps priority 5 to a lossless forwarding class (fcoe1), the input CNP enables PFC on incoming priority 5 traffic, and the output CNP enables output queue 5 to respond to pause messages received from the peer on traffic tagged with priority 5. On the native FC interface, FC traffic is remapped from priority 3 (the default mapping) to priority 5 and assigned to the same lossless forwarding class, fcoe1, because of the fixed classifier configuration. In this way, traffic tagged with priority 5 on interfaces xe-0/0/27 and fc-0/0/2 receives lossless treatment.



NOTE: To avoid fate sharing, ensure that the remapped priority is classified only to the forwarding class used in the fixed classifier on all other interfaces. For example, if you configure a fixed classifier on an FC interface that classifies all of the traffic into lossless forwarding class fcoe1 and remaps the priority to priority 5 (IEEE 802.1p code point 101), then in all other classifier configurations on all other interfaces, priority 5 should always be classified to forwarding class fcoe1. If you classify priority 6 on another interface to forwarding class fcoe1, then congestion on priority 6 traffic affects priority 5 traffic unfairly.



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the local FC fabric configuration. This examples focuses only on priority remapping.

Configuration

CLI Quick Configuration

To quickly configure IEEE 802.1p priority remapping on an FCoE-FC gateway, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_gw_classifier forwarding-class fcoe1 loss-priority
low code-points 101
set class-of-service interfaces xe-0/0/27 unit 0 classifiers ieee-802.1 fcoe_gw_classifier
set class-of-service interfaces fc-0/0/2 forwarding-class fcoe1
set class-of-service interfaces fc-0/0/2 rewrite-value input ieee-802.1p code-point 101
set class-of-service congestion-notification-profile fcoe1_p5_rewrite_cnp input ieee-802.1
code-point 101 pfc
set class-of-service congestion-notification-profile fcoe1_p5_rewrite_cnp output ieee-802.1
code-point 101 flow-control-queue 5
set class-of-service interfaces xe-0/0/27 congestion-notification-profile fcoe1_p5_rewrite_cnp
set applications application myfcoe5 ether-type 0x8906
set policy-options application-maps myfcoe5_app_map application myfcoe5 code-points 101
set protocols dcbx interface xe-0/0/27 application-map myfcoe5_app_map
```

Step-by-Step Procedure

To configure a lossless forwarding class for FCoE traffic, classify FCoE traffic into that forwarding class, configure a rewrite value on the native FC interface for the FCoE traffic, and enable PFC on the Ethernet interface, and configure DCBX application protocol TLV exchange for FCoE traffic:

1. Configure the lossless forwarding class (named **fcoe1** and mapped to output queue 5) for FCoE traffic that uses IEEE 802.1p priority 5:


```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```
2. Configure an ingress classifier named **fcoe_gw_classifier** to map the FCoE priority (IEEE 802.1p code point 101) to the lossless FCoE forwarding class (**fcoe1**):


```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_gw_classifier forwarding-class fcoe1 loss-priority low
code-points 101
```
3. Apply the classifier named **fcoe_gw_classifier** to Ethernet interface **xe-0/0/27**:


```
[edit class-of-service]
user@switch# set interfaces xe-0/0/27 unit 0 classifiers ieee-802.1 fcoe_gw_classifier
```
4. Configure the fixed classifier on the native FC interface, using the lossless FCoE forwarding class **fcoe1** (all traffic from the FC SAN is classified into the specified forwarding class). The traffic classified into this forwarding class is tagged with the priority value configured in the next step.


```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 forwarding-class fcoe1
```

5. Configure the rewrite value (IEEE 802.1p code point 101) applied to all incoming traffic from the FC SAN on the native FC interface. The rewrite value is the IEEE 802.1p priority that the encapsulated FCoE traffic classified into the **fcoe1** forwarding class uses on the converged Ethernet network.

```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 rewrite-value input ieee-802.1p code-point 101
```

6. Configure the input stanza of the CNP (named **fcoe1_p5_rewrite_cnp**) to enable PFC on the FCoE priority on the Ethernet interface:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe1_p5_rewrite_cnp input ieee-802.1p code-point 101 pfc
```

7. Configure the output stanza of the CNP to enable output queue 5 to respond to pause messages received from the peer on traffic tagged with priority 5:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe1_p5_rewrite_cnp output ieee-802.1p code-point 101 flow-control-queue 5
```

8. Apply the CNP named **fcoe1_p5_rewrite_cnp** to Ethernet interface **xe-0/0/27**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/27 congestion-notification-profile fcoe1_p5_rewrite_cnp
```

9. Configure a DCBX application for FCoE to map to the Ethernet interface, so that DCBX can exchange application protocol TLVs on the correct (remapped) IEEE 802.1p FCoE priority:

```
[edit]
user@switch# set applications application myfcoe5 ether-type 0x8906
```

10. Configure a DCBX application map to map the FCoE application to the correct (remapped) IEEE 802.1p FCoE priority:

```
[edit]
user@switch# set policy-options application-maps myfcoe5_app_map application myfcoe5 code-points 101
```

11. Apply the application map to the Ethernet interface so that DCBX exchanges FCoE application TLVs on the correct code point:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/27 application-map myfcoe5_app_map
```

Verification

To verify the configuration and proper operation of IEEE 802.1p priority remapping on an FCoE-FC gateway, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 5745](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 5745](#)
- [Verifying the FC Interface Configuration \(Fixed Classifier, Rewrite Value\) on page 5746](#)
- [Verifying the Ethernet Interface PFC Configuration \(CNP\) on page 5746](#)
- [Verifying the Ethernet Interface Configuration on page 5747](#)

- [Verifying the DCBX Application Configuration on page 5747](#)
- [Verifying the DCBX Application Map Configuration on page 5747](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 5748](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifier configured for priority remapping using the operational mode command **show class-of-service classifier name fcoe_gw_classifier**:

```
user@switch> show class-of-service classifier name fcoe_gw_classifier
Classifier: fcoe_gw_classifier, Code point type: ieee-802.1, Index: 13100
```

Code point	Forwarding class	Loss priority
101	fcoe1	low

Meaning The **show class-of-service classifier name fcoe_gw_classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in the classifier. The command output shows that the classifier maps forwarding class **fcoe1** to IEEE 802.1p code point 101 (priority 5) with a packet loss priority of **low**.

Verifying the FC Interface Configuration (Fixed Classifier, Rewrite Value)

Purpose Verify that the native FC interface (NP_Port) classifies incoming traffic into forwarding class **fcoe1** and that the interface rewrite value is priority 5 (IEEE code point 101).

Action Display the FC interface configuration using the operational mode command **show configuration class-of-service interfaces fc-0/0/2**:

```
user@switch> show configuration class-of-service interfaces fc-0/0/2
rewrite-value {
    input {
        ieee-802.1 {
            code-point {
                101;
            }
        }
    }
}
forwarding-class fcoe1;
```

Meaning The **show configuration class-of-service interfaces fc-0/0/2** command shows that the rewrite value for incoming (input) traffic is IEEE 802.1p code point **101** (priority 5), and that the interface uses forwarding class **fcoe1** as the fixed classifier for all incoming traffic.

Verifying the Ethernet Interface PFC Configuration (CNP)

Purpose Verify that PFC is enabled on the correct priority (IEEE 802.1p code point **101**) for lossless transport and that flow control is enabled on the correct output queue (queue **5**) on the Ethernet interface.

Action List the congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe1_p5_rewrite_cnp**:

```
user@switch> show class-of-service congestion-notification fcoe1_p5_rewrite_cnp
Name: fcoe1_p5_rewrite_cnp, Index: 7061
Type: Input
Cable Length: 100 m
  Priority  PFC          MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled    2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

Meaning The **show class-of-service congestion-notification fcoe1_p5_rewrite_cnp** command shows the input and output stanzas of the CNP. The input stanza shows that PFC is enabled on IEEE 802.1p code point 101 (priority 5). The input stanza also shows that the CNP uses

the default values of 100 meters for the cable length value and 2500 bytes for the maximum receive unit (MRU) value.

The output stanza shows that flow control is enabled on output queue 5 for IEEE 802.1p priority code point 101 (priority 5).

Verifying the Ethernet Interface Configuration

- Purpose** Verify that the classifier **fcoe_gw_classifier** and the congestion notification profile **fcoe1_p5_rewrite_cnp** are configured on Ethernet interface **xe-0/0/27**.
- Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces xe-0/0/27**:
- ```
user@switch> show configuration class-of-service interfaces xe-0/0/27
congestion-notification-profile fcoe1_p5_rewrite_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_gw_classifier;
 }
}
```
- Meaning** The **show configuration class-of-service interfaces xe-0/0/27** command shows that the congestion notification profile **fcoe1\_p5\_rewrite\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_gw\_classifier**.

### *Verifying the DCBX Application Configuration*

- Purpose** Verify that the DCBX application named **myfcoe5** for FCoE is configured.
- Action** List the DCBX applications by using the configuration mode command **show applications**:
- ```
user@switch# show applications
application myfcoe5 {
  ether-type 0x8906;
```
- Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **myfcoe5** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic).

Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map **myfcoe5_app_map** is configured.
- Action** List the application map by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
myfcoe5_app_map {
 application myfcoe5 code-points 101;
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map.

The output shows that there is one application map, **myfcoe5\_app\_map**, which consists of the application named **myfcoe5** mapped to IEEE 802.1p code point **101** (priority 5).

### *Verifying the DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application map is applied to the correct interface (xe-0/0/27).

**Action** List the application maps using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/27.0 {
 application-map myfcoe5_app_map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interface **xe-0/0/27** uses application map **myfcoe5\_app\_map**.

- Related Documentation**
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
  - [Example: Configuring Unicast Classifiers on page 5658](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
  - [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 5801](#)
  - [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)

## **Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic**

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) unicast traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

---

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.



After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 5749](#)
- [Overview on page 5749](#)
- [Configuration on page 5750](#)
- [Verification on page 5752](#)

---

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

---

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort unicast traffic, more buffer space needs to be allocated to lossy buffers, and less buffer space should be allocated to lossless buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly unicast traffic.

### Topology

[Table 541 on page 5750](#) shows the configuration components for this example.

**Table 541: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                           |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%                       |

### Configuration

**CLI Quick Configuration** To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text

file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

### *Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic*

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly best-effort unicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
user@switch# set egress buffer-partition multicast percent 20
```

### **Results**

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 95;
 }
 buffer-partition lossless-headroom {
 percent 0;
 }
}
```

```
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 75;
 }
 buffer-partition multicast {
 percent 20;
 }
}
```

---

### Verification

Verify that the shared buffer configuration has been created properly.

#### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
 Shared Buffer : 7202.00 KB
 Lossless : 360.10 KB
 Lossless Headroom : 0.00 KB
 Lossy : 6841.90 KB

 Lossless Headroom Utilization:
 Node Device Total Used Free
 0 0.00 KB 0.00 KB 0.00 KB

Egress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2704.00 KB
 Shared Buffer : 6656.00 KB
 Lossless : 332.80 KB
 Multicast : 1331.20 KB
 Lossy : 4992.00 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved,

ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.

- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic
  - No space to lossless headroom traffic
  - 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 1331.20 KB to multicast traffic
  - 4992 KB to lossy unicast traffic

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) traffic on links with Ethernet PAUSE (IEEE 802.3X) enabled. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

---

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 5754](#)
- [Overview on page 5754](#)
- [Configuration on page 5756](#)
- [Verification on page 5757](#)

### Requirements

---

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

### Overview

---

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic.

From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE, more buffer space needs to be allocated to ingress dedicated port buffers, and less buffer space should be allocated to ingress shared buffers. Also, more buffer space needs to be allocated to lossless-headroom buffers, and less space to ingress lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE.

### Topology

Table 542 on page 5756 shows the configuration components for this example.

**Table 542: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                           |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 70%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 80%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 15% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%                       |

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 70
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 15
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

#### Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links Enabled for Ethernet PAUSE

#### Step-by-Step Procedure

To configure the global ingress and egress shared buffer allocations and partitions:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 70
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:



```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 15
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```

4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
user@switch# set egress buffer-partition multicast percent 20
```

### Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 70;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 15;
 }
 buffer-partition lossless-headroom {
 percent 80;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 75;
 }
 buffer-partition multicast {
 percent 20;
 }
}
```

### Verification

Verify that the shared buffer configuration has been created properly.

#### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
```

```
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 4318.60 KB
 Shared Buffer : 5041.40 KB
 Lossless : 252.07 KB
 Lossless Headroom : 4033.12 KB
 Lossy : 756.21 KB

Egress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2704.00 KB
 Shared Buffer : 6656.00 KB
 Lossless : 332.80 KB
 Multicast : 1331.20 KB
 Lossy : 4992.00 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 4318.6 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 70 percent of the available (user-configurable) buffer space.
- With the ingress shared buffer pool configured as 70 percent of the available buffers, the total size of the ingress shared buffer pool is 5041.4 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 252.07 KB to lossless traffic
  - 4033.12 KB to lossless headroom traffic
  - 756.21 KB to lossy unicast traffic

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:

- 332.80 KB to lossless traffic
- 1331.20 KB to multicast traffic
- 4992 KB to lossy unicast traffic

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

### Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly multicast traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 5760](#)
- [Overview on page 5760](#)
- [Configuration on page 5761](#)
- [Verification on page 5763](#)

## Requirements

---

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

---

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.

- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly multicast traffic, more buffer space needs to be allocated to lossy buffers, less buffer space should be allocated to lossless buffers, and more space needs to be allocated to egress multicast buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly multicast traffic.

### Topology

[Table 543 on page 5761](#) shows the configuration components for this example.

**Table 543: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                           |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 20%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 75%                       |

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly multicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
```

```
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 20
set egress buffer-partition multicast percent 75
```

### *Configuring the Global Shared Buffer Pool for Networks with Mostly Multicast Traffic*

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly multicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 20
user@switch# set egress buffer-partition multicast percent 75
```

### *Results*

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 95;
 }
 buffer-partition lossless-headroom {
 percent 0;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 20;
 }
 buffer-partition multicast {
```

```

 percent 75;
 }
}

```

## Verification

Verify that the shared buffer configuration has been created properly.

### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```

user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
 Shared Buffer : 7202.00 KB
 Lossless : 360.10 KB
 Lossless Headroom : 0.00 KB
 Lossy : 6841.90 KB

 Lossless Headroom Utilization:
 Node Device Total Used Free
 0 0.00 KB 0.00 KB 0.00 KB

Egress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2704.00 KB
 Shared Buffer : 6656.00 KB
 Lossless : 332.80 KB
 Multicast : 4992.00 KB
 Lossy : 1331.20 KB

```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic

- No space to lossless headroom traffic
- 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 4992 KB to multicast traffic
  - 1331.20 KB to lossy unicast traffic

**Related  
Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

### Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.



This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly lossless traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.



**NOTE:** When we discuss lossless buffers, we mean buffers that handle traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x).

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 5765](#)
- [Overview on page 5765](#)
- [Configuration on page 5767](#)
- [Verification on page 5768](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly lossless traffic, more buffer space needs to be allocated to lossless buffers, and less buffer space should be allocated to lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly lossless traffic.

### ***Topology***

[Table 544 on page 5767](#) shows the configuration components for this example.

Table 544: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                            |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 15%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless headroom buffer partition): 80%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 5% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 90%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 5%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 5%                         |

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly lossless traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 15
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 5
set egress percent 100
set egress buffer-partition lossless percent 90
set egress buffer-partition lossy percent 5
set egress buffer-partition multicast percent 5
```

#### Configuring the Global Shared Buffer Pool for Networks with Mostly Lossless Traffic

#### Step-by-Step Procedure

To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly lossless traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```

2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 15
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 5
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 90
user@switch# set egress buffer-partition lossy percent 5
user@switch# set egress buffer-partition multicast percent 5
```

### Results

Display the results of the configuration:

```
rroot@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 15;
 }
 buffer-partition lossy {
 percent 5;
 }
 buffer-partition lossless-headroom {
 percent 80;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 90;
 }
 buffer-partition lossy {
 percent 5;
 }
 buffer-partition multicast {
 percent 5;
 }
}
```

---

### Verification

Verify that the shared buffer configuration has been created properly.

#### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
```

```

Shared Buffer : 7202.00 KB
Lossless : 1080.30 KB
Lossless Headroom : 5761.60 KB
Lossy : 360.10 KB

```

Lossless Headroom Utilization:

| Node Device | Total      | Used    | Free       |
|-------------|------------|---------|------------|
| 0           | 5761.60 KB | 0.00 KB | 5761.60 KB |

Egress:

```

Total Buffer : 9360.00 KB
Dedicated Buffer : 2704.00 KB
Shared Buffer : 6656.00 KB
 Lossless : 5990.40 KB
 Multicast : 332.80 KB
 Lossy : 332.80 KB

```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 1080 KB to lossless traffic
  - 5761.60 KB to lossless headroom traffic
  - 360.10 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Of the total available lossless headroom buffer space of 5761.60 KB, currently no buffer space is being used, so all 5761.60 KB of buffer space is free.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.

- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 5990.40 KB to lossless traffic
  - 332.80 KB to multicast traffic
  - 332.80 KB to lossy unicast traffic

**Related Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The QFX Series handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the QFX Series exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the

FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).

- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 5771](#)
- [Overview on page 5771](#)
- [Configuration on page 5775](#)
- [Verification on page 5776](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device
- Junos OS Release 12.1 or later for the QFX Series

## Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 5204](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.

- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 407 on page 5146](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 408 on page 5146](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 545: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

**Table 546: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |



Table 546: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier) (*continued*)

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

**Topology**

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 409 on page 5147 shows the configuration components for this example.

Table 547: Components of DCBX Application Protocol Exchange Configuration Topology

| Component                   | Settings                                                                                                           |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------|
| Hardware                    | QFX Series device                                                                                                  |
| LLDP                        | Enabled by default on Ethernet interfaces                                                                          |
| DCBX                        | Enabled by default on Ethernet interfaces                                                                          |
| iSCSI application (Layer 4) | Application name— <b>iscsi</b><br>protocol— <b>TCP</b><br>destination-port— <b>3260</b><br>code-points— <b>111</b> |
| PTP application (Layer 2)   | Application name— <b>ptp</b><br>ether-type— <b>0x88F7</b><br>code-points— <b>001, 101</b>                          |

**Table 547: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)**

| Component                                                                                                         | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCoE application (Layer 2)                                                                                        | <p>Application name—<b>fcoe</b></p> <p>ether-type—<b>0x8906</b></p> <p>code-points—<b>011</b></p> <p><b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Application maps                                                                                                  | <p><b>dcbx-iscsi-fcoe-app-map</b>—Maps the iSCSI and FCoE applications to IEEE 802.1p code points</p> <p><b>dcbx-iscsi-ptp-app-map</b>—Maps iSCSI and PTP applications to IEEE 802.1p code points</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Interfaces                                                                                                        | <p><b>xe-0/0/10</b>—Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b>, carries FCoE traffic, and has PFC enabled on the FCoE priority)</p> <p><b>xe-0/0/11</b>—Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b>)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| PFC congestion notification profile for FCoE application exchange                                                 | <p><b>fcoe-cnp:</b></p> <ul style="list-style-type: none"> <li>Code point—<b>011</b></li> <li>Interface—<b>xe-0/0/10</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point) | <p><b>fcoe-iscsi-cl1:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (<b>011</b>) and a loss priority of <b>high</b></li> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>high</b></li> <li>Applied to interface <b>xe-0/0/10</b></li> </ul> <p><b>iscsi-ptp-cl2:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>low</b></li> <li>Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (<b>001</b> and <b>101</b>) and a loss priority of <b>low</b></li> <li>Applied to interface <b>xe-0/0/11</b></li> </ul> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

## Configuration

**CLI Quick Configuration** To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

### Configuring DCBX Application Protocol TLV Exchange

**Step-by-Step Procedure** To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  
  

```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  
  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  
  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```

4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
```

5. Create the congestion notification profile to enable PFC on the FCoE code point (**011**), and apply the congestion notification profile to interface **xe-0/0/10**.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```

6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority
high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control
loss-priority high code-points 111
```

7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
```

8. Apply the classifiers to the appropriate interfaces.

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ntp-cl2
```

---

## Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 5776](#)
- [Verifying the Application Map Configuration on page 5777](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 5777](#)
- [Verifying the PFC Configuration on page 5778](#)
- [Verifying the Classifier Configuration on page 5779](#)

### *Verifying the Application Configuration*

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```
user@switch# show applications
```

```

application iSCSI {
 protocol tcp;
 destination-port 3260;
}

application fcoe {
 ether-type 0x8906;
}

application ptp {
 ether-type 0x88F7;
}

```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

#### *Verifying the Application Map Configuration*

**Purpose** Verify that the application maps have been configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```

user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
 application iSCSI code-points 111;
 application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
 application iSCSI code-points 111;
 application PTP code-points [001 101];
}

```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the FCoE application, which is mapped to IEEE 802.1p code point **011**.

The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the PTP application, which is mapped to IEEE 802.1p code points **001** and **101**.

#### *Verifying DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application maps have been applied to the correct interfaces.

**Action** List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
 application-map dcbx-iscsi-ptp-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.

#### *Verifying the PFC Configuration*

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

---

**Meaning** The `show class-of-service congestion-notification-profile` configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile `fcoe-cnp` has been configured and has enabled PFC on the IEEE 802.1p code point `011` (the default FCoE code point).

The `show class-of-service interfaces` configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile `fcoe-cnp`, which enables PFC on the FCoE code point, is applied to interface `xe-0/0/10`.

### *Verifying the Classifier Configuration*

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command `show class-of-service`:

```
user@switch# show class-of-service
classifiers {
 ieee-802.1 fcoe-iscsi-cl1 {
 import default;
 forwarding-class network-control {
 loss-priority high code-points 111;
 }
 forwarding-class fcoe {
 loss-priority high code-points 011;
 }
 }
 ieee-802.1 iscsi-ptp-cl2 {
 import default;
 forwarding-class network-control {
 loss-priority low code-points 111;
 }
 forwarding-class best-effort {
 loss-priority low code-points [001 101];
 }
 }
}
interfaces {
 xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-iscsi-cl1;
 }
 }
 }
 xe-0/0/11 {
 unit 0 {
 classifiers {
 ieee-802.1 iscsi-ptp-cl2;
 }
 }
 }
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ntp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ntp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ntp-cl2** is mapped to interface **xe-0/0/11.0**.

**Related Documentation**

- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [show dcbx on page 5299](#)
- [show dcbx neighbors on page 5300](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Using DCBX Protocol to Lower Costs](#)



---

## Configuration Tasks

---

- [Configuring CoS on page 5781](#)
- [Defining CoS Code-Point Aliases on page 5783](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Defining CoS Forwarding Classes on page 5787](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Defining CoS Queue Scheduling Priority on page 5792](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 5793](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 5801](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Defining CoS Rewrite Rules on page 5805](#)
- [Assigning CoS Components to Interfaces on page 5807](#)
- [Configuring the DCBX Mode on page 5808](#)
- [Configuring DCBX Autonegotiation on page 5809](#)
- [Disabling the ETS Recommendation TLV on page 5812](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5812](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5814](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5815](#)

## Configuring CoS

The class-of-service topics describe how to configure the Junos CoS components for the QFX Series. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components, you assign classifiers to the required physical and logical interfaces.

You can configure various CoS components individually or in combination to define CoS services.



**NOTE:** When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

The following topics describe how to configure CoS components :

- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring Drop Profile Maps on page 5666](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Queue Scheduling Priority on page 5679](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5121](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
- [Defining CoS Code-Point Aliases on page 5783](#)
- [Defining CoS Rewrite Rules on page 5805](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)

- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 5801](#)
- [Assigning CoS Components to Interfaces on page 5807](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 5793](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)

## Defining CoS Code-Point Aliases

You can use code-point aliases to streamline the process of configuring CoS features on your switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP or DSCP IPv6—Handles incoming IPv4 or IPv6 packets.
- IEEE 802.1p—Handles Layer 2 CoS.

To configure a code-point alias:

1. Specify a CoS marker type (IEEE 802.1 or DSCP).
2. Assign an alias.
3. Specify the code point that corresponds to the alias.

```
[edit class-of-service code-point-aliases]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) alias-name code-point-bits
```

For example, to configure a code-point alias for an IEEE 802.1 CoS marker type that has the alias name `fcoe1` and maps to the code-point bits 011:

```
[edit class-of-service code-point-aliases]
user@switch# set ieee-802.1 fcoe1 011
```

### Related Documentation

- [Monitoring CoS Value Aliases on page 5920](#)
- [Understanding CoS Code-Point Aliases on page 5453](#)

## Defining CoS Unicast BA Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP or DSCP IPv6) value or IEEE 802.1p CoS value in the packet header to determine the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.

Unicast traffic must use different classifiers than multidestination (multicast, broadcast, and destination lookup fail) traffic.

To configure a unicast BA classifier using the CLI:

1. Create a unicast BA classifier:

- To create a unicast BA classifier based on the default classifier, import the default DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name import default forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

- To create a unicast BA classifier that is not based on the default classifier, create a DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the unicast classifier to a specific 10-Gigabit Ethernet interface or to all 10-Gigabit Ethernet interfaces or to all Fibre Channel interfaces on the switch.

- To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name unit unit classifiers (dscp | ieee-802.1) classifier-name
```

- To apply the classifier to all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and the logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * classifiers (dscp | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)
- [Monitoring CoS Classifiers on page 5915](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

## Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP) value or IEEE 802.1p CoS value in the packet header to determine the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.



**NOTE:** DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.

Unicast and multidestination traffic must use different classifiers.

To configure a multidestination BA classifier using the CLI:

1. Create a DSCP or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class forwarding-class-name
loss-priority level code-points [aliases] [bit-patterns]
```

2. Configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers (dscp | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 5661](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)
- [Monitoring CoS Classifiers on page 5915](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

## Configuring CoS Tail-Drop Profiles

You can configure an interpolated weighted random early detection (WRED) tail-drop profile to control packet drop characteristics for different traffic loss priorities.



**NOTE:** You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

*Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point.

The drop start point is the average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop. Specify the drop start point using the first of two **fill-level** statements.

The drop end point is the average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.) Specify the drop end point using the second of two **fill-level** statements.

The minimum drop rate is always 0. Specify the minimum drop rate using the first of two **drop-probability** statements. The maximum drop rate is the drop probability when the average queue fill level reaches the drop end point. Specify the maximum drop rate using the second of two **drop-probability** statements.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

To configure a tail-drop profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

## Configuring CoS Drop Profile Maps

A drop-profile map associates a tail-drop profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified tail-drop profile to traffic in the forwarding class that matches the specified loss priority.

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

To configure a drop-profile map using the CLI:

- For the desired scheduler, configure the traffic loss priority and specify the drop profile you want to use to control the drop characteristics for traffic of that loss priority:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name drop-profile-map loss-priority level protocol
any drop-profile drop-profile-name
```

### Related Documentation

- [Example: Configuring Drop Profile Maps on page 5666](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

## Defining CoS Forwarding Classes

Forwarding classes allow you to group packets for transmission. The switch supports a total of 12 forwarding classes. In order to forward traffic, you map (assign) the forwarding classes to unicast or multdestination (multicast, broadcast, and destination lookup fail) output queues.

The switch has 12 output queues. Queues 0 through 7 are for unicast traffic and queues 8 through 11 are for multicast traffic. Forwarding classes mapped to unicast queues must carry unicast traffic, and forwarding classes mapped to multdestination queues must carry multdestination traffic. There are four default unicast forwarding classes and one default multdestination forwarding class.

The default unicast forwarding classes are:

- **best-effort**—Best-effort traffic
- **fcoe**—Guaranteed delivery for FCoE traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **network-control**—Network control traffic

The default multidestination forwarding class is:

- **mcast**—Multidestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. In addition, all forwarding classes mapped to a particular queue must be either lossless or lossy. You cannot mix lossless and lossy forwarding classes (traffic) on the same queue. Also, you cannot mix unicast and multicast forwarding classes on the same queue.

[edit [class-of-service forwarding-classes](#)]

```
user@switch# class class-name queue-num queue-number <no-loss>
```

---



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment unless you include the optional **no-loss** packet drop attribute introduced in Junos OS Release 12.3 in the forwarding class configuration..

---



**NOTE:** Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

[edit [class-of-service forwarding-classes](#)]

```
user@switch# queue queue-number class-name
```

The **queue** statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the **queue** statement to map forwarding classes to queues, edit the configuration to replace the **queue** statement with the **class** statement.

---

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Monitoring CoS Forwarding Classes on page 5916](#)
- [Understanding CoS Forwarding Classes on page 5469](#)



## Defining CoS Forwarding Class Sets

A forwarding class set is a priority group for enhanced transmission selection (ETS) traffic control. Each forwarding class set consists of one or more forwarding classes (priorities, which can also be considered as output queues).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set.

To configure a forwarding class set using the CLI:

1. Assign one or more forwarding classes to the forwarding class set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets forwarding-class-set-name class
forwarding-class-name
```

2. Map the forwarding class set to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set forwarding-class-set-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5474](#)

## Defining CoS Queue Schedulers

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, the tail-drop profiles associated with the queue, and the queue buffer size.

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR), set as a percentage rate or as an absolute value in bits per second. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.



**NOTE:** You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR), set as a percentage rate or as an absolute value in bits per second.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:
  - **low**—The scheduler has low priority.
  - **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.
- **drop-profile-map**—Drop profile mapping to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.



**NOTE:** Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

To apply scheduling properties to traffic, map schedulers to forwarding classes using a scheduler map, and then associate the scheduler map with the interfaces. This applies the configured scheduling to the traffic in the specified forwarding class on the associated interface. Using different scheduler maps, you can map different schedulers to the same traffic (the same forwarding class) to apply different scheduling to that traffic on different interfaces.

To configure a scheduler using the CLI:

1. Name the scheduler and define the minimum guaranteed bandwidth for the queue:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name transmit-rate (rate | percent percentage)
```

2. Define the maximum bandwidth for the queue:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Define the queue priority:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set priority level
```

4. Define the drop profile using a drop profile map:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set drop-profile-map loss-priority (low | medium-high | high) protocol protocol
drop-profile drop-profile-name
```

5. Configure the size of the port dedicated buffer space for the queue:

```
[edit class-of-service schedulers scheduler-name]
```

```
user@switch# set buffer-size percent 20
```

6. Configure a scheduler map to map the scheduler to a forwarding class, which applies the scheduler's properties to the traffic in that forwarding class:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

7. Assign the scheduler map to one or more interfaces to apply the scheduling to the traffic that belongs to the specified forwarding class on those interfaces:

```
[edit class-of-service]
user@switch# set interfaces interface-name scheduler-map scheduler-map-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Defining CoS Queue Scheduling Priority on page 5792](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Monitoring CoS Scheduler Maps on page 5919](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)
- [Understanding CoS Priority Group Scheduling on page 5493](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## Defining CoS Queue Scheduling Priority

You can configure the scheduling priority of individual queues by specifying the priority in a scheduler, and then associating the scheduler with a queue by using a scheduler map. Queues can have one of two bandwidth priorities:

- **strict-high** —The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.
- **low** —Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.
- To configure queue priority using the CLI:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name priority level
```

#### Related Documentation

- [Example: Configuring Queue Scheduling Priority on page 5679](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Monitoring CoS Scheduler Maps on page 5919](#)

- [Understanding CoS Output Queue Schedulers on page 5486](#)

## Changing the Host Outbound Traffic Default Queue Mapping

If you do not want to use the default mapping of host Routing Engine and CPU outbound traffic to queues, you can change the default output queue. You can also change the default DSCP bits used in the type of service (ToS) field of packets generated by the Routing Engine.

Configuring a queue for host outbound traffic maps all traffic that the host generates to one forwarding class (queue). The configuration is global and applies to all host-generated traffic on the switch. Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.



**NOTE:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (`fabric`) and the egress Node device.

To change the host outbound traffic egress queue by including the **host-outbound-traffic** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
host-outbound-traffic {
 forwarding-class class-name;
 dscp-code-point code-point;
}
```

For example, to map host outbound traffic to queue 7 (the network control forwarding class) and set the DSCP code point value to 101010:

```
[edit class-of-service]
host-outbound-traffic {
 forwarding-class network-control;
 dscp-code-point 101010
}
```

### Related Documentation

- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5451](#)

## Defining CoS Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) contained in a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR).
- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

To configure a traffic control profile using the CLI:

1. Name the traffic control profile and define the minimum guaranteed bandwidth for the priority group:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name guaranteed-rate (rate
| percent percentage)
```

2. Define the maximum bandwidth for the priority group:

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Attach a scheduler map to the traffic control profile:

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set scheduler-map scheduler-map-name
```

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Understanding CoS Traffic Control Profiles on page 5496](#)

**Configuring CoS PFC (Congestion Notification Profiles)**

A congestion notification profile (CNP) enables priority-based flow control (PFC) on specified IEEE 802.1p priorities (code points). A CNP has two components:

- Input CNP:
  - Enable PFC on a specified priority.
  - Configure the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority (optional).
  - Specify the length of the attached cable on the ingress interface (optional)
- Output CNP (optional): Configure flow control to enable PFC pause on specific output queues for specified priorities.



**NOTE:** By default, output queues 3 and 4 (which are mapped to default lossless forwarding classes *fcoe* and *no-loss*, respectively) are configured to respond to PFC pause messages received from the connected peer on priorities 3 and 4 (code points 011 and 100, respectively). If you explicitly configure flow control on any output queue, you must configure flow control on every output queue that you want to respond to pause messages. (The explicit configuration overrides the default configuration.)

To achieve lossless behavior, the output queue priorities on which you enable PFC flow control must match the PFC priorities on which you enable PFC on the input interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101) in the output component of the CNP, then you must also enable pause on priorities 3 and 5 on the input component of the CNP. (In addition, the forwarding classes mapped to the paused output queues must be lossless forwarding classes.)

Associating a CNP with an interface enables PFC on the ingress traffic that matches the priority specified in the input CNP, and programs the queues listed in the output CNP to pause when the interface receives a PFC pause message from the connected peer. Configure PFC on a priority end to end along the entire data path to create a lossless lane of traffic on the network.



**NOTE:** You must enable PFC on the priority used by FCoE traffic on ingress interfaces (input CNP). Enable PFC on the FCoE priority on every interface that carries FCoE traffic. By convention, FCoE traffic uses priority 3 (code point 011), which maps to queue 3. If your network uses priority 3 for FCoE traffic, the default forwarding class and classifier configuration support lossless transport, but you must still configure a CNP and apply it to the correct ingress interfaces to enable PFC and achieve lossless transport.

If your network does not use priority 3 for FCoE traffic, you need to configure a classifier that classifies FCoE traffic into a lossless forwarding class, based on the priority your network uses for FCoE traffic. If you are not using the default lossless forwarding class configuration, then you also need to ensure that the output queue mapped to the lossless FCoE forwarding class is programmed to pause.

---

You can attach only one CNP to an interface. There is no limit to the total number of CNPs you can create.

Configuring a CNP consists of:

- Naming the CNP.
- Specifying the IEEE 802.1 code point (priority) on which you want to enable PFC on ingress interfaces (input CNP).
- Optionally, specifying the MRU and the length of the attached cable on ingress interfaces (input CNP).
- Optionally, configuring flow control (PFC pause) on specified output queues if you want queues other than queues 3 and 4 to respond to pause messages received from the connected peer (output CNP).
- Mapping the CNP to an interface.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

---



1. Enable PFC on the desired priority in the input CNP and optionally configure the interface MRU for traffic on that priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input ieee-802.1 code-point
code-point bits pfc mru mru-value
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC on IEEE 802.1 code point **011** and configures an MRU value of **2240**:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011
pfc mru 2240
```

2. Configure the length of the cable attached to the ingress interface (optional):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input cable-length
cable-length-value
```

For example, to configure a CNP named **fcoe-cnp** that sets the length of the ingress interface cable to **100** meters:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input cable-length 100
```

3. (Optional) Configure flow control on output queues:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
code-point-bits flow-control-queue [queue | list-of-queues]
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC pause flow control on output queues 3 and 5 for FCoE traffic that uses priority 3 (code point **011**) and on output queue 4 for traffic that uses priority 4 (code point **100**):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
011 flow-control-queue [3 5]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
100 flow-control-queue 4
```

4. Map the CNP to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface congestion-notification-profile cnp-name
```

For example, to map the CNP **fcoe-cnp** to the interface **xe-0/0/7**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 congestion-notification-profile fcoe-cnp
```

#### Related Documentation

- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Assigning CoS Components to Interfaces on page 5807](#)
- [Monitoring Interfaces That Have CoS Components on page 5917](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

## Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to Ethernet link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Symmetric flow control means that an interface has the same PAUSE configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled, or they are both disabled.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. If you attempt to configure both features, the switch returns a commit error. Ethernet PAUSE and PFC are also mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

By default, all flow control features are disabled. You enable symmetric flow control on the interfaces on which you want to PAUSE all of the traffic on a link.

- To enable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set flow-control
```

- To disable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set no-flow-control
```

### Related Documentation

- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## Configuring CoS Asymmetric Ethernet PAUSE Flow Control

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction.

Symmetric flow control means that the interface has the same configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled or they are both disabled. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

Asymmetric flow control provides the ability to configure the receive buffer and transmit buffer Ethernet PAUSE actions independently on an interface. The buffers perform the following actions:

- The receive buffers generate and send PAUSE messages to the connected peer to ask the peer to stop sending traffic for a time period specified in the PAUSE frame. The peer interface's buffers may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.
- The transmit buffers respond to PAUSE messages received from the connected peer to stop sending traffic to the peer. The transmit buffer may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to PAUSE messages:

- Receive buffers on—Enable PAUSE transmission (generate and send PAUSE frames)
- Transmit buffers on—Enable PAUSE reception (respond to received PAUSE frames)

You must explicitly set both the receive buffer and the transmit buffer to configure asymmetric flow control.

- To configure asymmetric flow control on an interface:  
[edit **interfaces** *interface-name* **ether-options**]

```
user@switch# set configured-flow-control rx-buffers (on | off) tx-buffers (on | off)
```

For example, to configure interface **xe-0/0/24** to generate and send PAUSE messages but not to respond to received PAUSE messages:

```
set interfaces xe-0/0/24 ether-options configured-flow-control rx-buffers on tx-buffers off
```

For example, to configure interface **xe-0/0/30** to respond to received PAUSE messages but not to generate and send PAUSE messages:

```
set interfaces xe-0/0/30 ether-options configured-flow-control rx-buffers off tx-buffers on
```



**NOTE:** If you configure both buffers to be on, that is equivalent to symmetric flow control. If you configure both buffers to be off, there is no flow control (flow control is disabled).

---

**Related  
Documentation**

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP\_Ports)

Fibre Channel over Ethernet (FCoE) traffic typically uses IEEE 802.1p priority 3 (code point 011). When Fibre Channel (FC) traffic arrives on a native FC interface (NP\_Port) on an FCoE-FC gateway, the interface encapsulates the FC traffic in Ethernet to create FCoE frames. By default, the native FC interface assigns priority 3 to the FCoE traffic. The traffic is then forwarded internally to the gateway Ethernet interfaces, and then forwarded to the FCoE network.

If your FCoE network uses priority 3 for FCoE traffic, you do not need to use a rewrite value to remap the FCoE priority on native FC interfaces, because the default configuration maps priority 3 to the FCoE forwarding class.

However, if the FCoE network uses a different priority than priority 3 for FCoE traffic, then you can configure a rewrite value to remap incoming traffic from the FC SAN to that priority after the interface encapsulates the FC packets in Ethernet. Setting a rewrite value for the IEEE 802.1p code point (priority) configures the gateway native FC interface to assign the rewrite value to the encapsulated FCoE frames before forwarding the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

Traffic coming from the FC SAN is classified into a lossless forwarding class, and that lossless forwarding class is mapped to the rewrite value (the priority used for FCoE traffic on the converged Ethernet network). You specify the lossless forwarding class used for FCoE traffic on a native FC interface by configuring a fixed classifier and applying it to the native FC interface. (The same forwarding class must also be mapped to the rewrite value priority in the ingress classifier applied to the FCoE Ethernet interfaces.) All traffic received from the FC SAN on that FC interface is encapsulated in Ethernet, classified into the forwarding class specified in the fixed classifier, and assigned the rewrite value priority.

Configuring a rewrite value consists of:

- Configuring a fixed classifier on the native FC interface. The fixed classifier assigns all the traffic that arrives at the interface from the connected peer in the FC SAN to one fixed forwarding class. The forwarding class must be a lossless forwarding class and must be classified to the rewrite value priority in the ingress classifier configuration on the FCoE Ethernet interfaces.
- Specifying an IEEE 802.1p rewrite value for the native FC interface. The traffic mapped to the forwarding class in the fixed classifier is marked with the priority you specify in the rewrite value when the traffic is encapsulated in Ethernet. The rewrite value must be the IEEE 802.1p priority used for FCoE traffic in your converged Ethernet network.

You can configure one rewrite value per native FC interface. Each native FC interface can use a different rewrite value.

1. Configure a fixed classifier on the native FC interface:

```
[edit class-of-service]
user@switch# set interfaces fc-interface-name forwarding-class
lossless-forwarding-class-name
```

For example, to configure a fixed classifier on native FC interface **fc-0/0/2** that specifies the lossless forwarding class **fcoe1**:

```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 forwarding-class fcoe1
```

2. Configure a rewrite value for the traffic classified into the fixed classifier (this must be the IEEE 802.1p priority used for the traffic on your converged Ethernet network):

```
[edit class-of-service]
user@switch# set interfaces fc-interface-name rewrite-value input ieee-802.1 code-point
code-point-bits
```

For example, to configure a rewrite value on native FC interface **fc-0/0/2** that specifies an IEEE 802.1p priority of **101** (the lossless forwarding class specified in the fixed classifier must be classified to this priority in the ingress classifier configuration on the FCoE Ethernet interfaces):

```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 rewrite-value input ieee-802.1 code-point 101
```

In the example, all traffic from the FC SAN that arrives at FCoE-FC gateway interface **fc-0/0/2** is encapsulated in Ethernet, classified into the lossless **fcoe1** forwarding class, and tagged with the IEEE 802.1p priority 5 (code point 101). In this example, we assume that the converged Ethernet network uses priority 5 for FCoE traffic, and that the **fcoe1** forwarding class is mapped to priority 5 in the ingress classifier configuration on the Ethernet interfaces. To achieve lossless transport, you must also enable PFC on priority 5 on the Ethernet interfaces that connect the FCoE traffic to the Ethernet network.

**Related  
Documentation**

- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)

## Configuring Global Ingress and Egress Shared Buffers

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of ingress and egress buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best-effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from one of the recommended configurations, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more

than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffer allocation and partitioning using the CLI:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent percent
```

2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent percent
user@switch# set ingress buffer-partition lossless-headroom percent percent
user@switch# set ingress buffer-partition lossy percent percent
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set egress percent percent
```

4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent percent
user@switch# set egress buffer-partition lossy percent percent
user@switch# set egress buffer-partition multicast percent percent
```

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 5754](#)



- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## Defining CoS Rewrite Rules

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on an interface. You can also apply an existing rewrite rule on an interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the rewrite rule and then apply the new rule.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named **customup-rw** in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point
000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point
001
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point
010
user@switch# set ieee-802.1 customup-rw forwarding-class fcfe loss-priority low code-point
011
```

```
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority low
code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority high
code-point 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low code-point
110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high code-point
111
```

- To enable an 802.1p rewrite rule named **customup-rw** on a Layer 2 interface:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/7 unit 0 rewrite-rules ieee-802.1
customup-rw
```



**NOTE:** All forwarding classes assigned to port xe-0/0/7 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

- To enable an 802.1p rewrite rule named **customup-rw** on all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical interface (unit) number:

```
[edit]
user@switch# set class-of-service interfaces xe-* unit * rewrite-rules customup-rw
```



**NOTE:** In this case, *all* forwarding classes assigned to *all* 10-Gigabit Ethernet ports must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

#### Related Documentation

- [Monitoring CoS Rewrite Rules on page 5918](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

## Assigning CoS Components to Interfaces

After you define the following CoS components, you assign them to physical or logical interfaces. Components that you assign to physical interfaces are valid for all of the logical interfaces configured on the physical interface. Components that you assign to a logical interface are valid only for that logical interface.

- Classifiers—Assign only to logical interfaces.
- Congestion notification profiles—Assign only to physical interfaces.
- Forwarding classes—Assign to interfaces by mapping to forwarding class sets.
- Forwarding class sets—Assign only to physical interfaces.
- Output traffic control profiles—Assign only to physical interfaces (with a forwarding class set).
- Rewrite rules—Assign only to logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wildcards. You can also assign a congestion notification profile or a forwarding class set globally to all interfaces.

To assign CoS components to interfaces:

Assign CoS components to a single interface by associating a CoS component (for example a forwarding class set named **san-priority-group**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set san-priority-group
```

Assign a CoS component to multiple interfaces by associating a CoS component (for example, a rewrite rule named **customup-rw**) to all 10-Gigabit Ethernet interfaces on the switch, use wildcard characters for the interface name and logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * rewrite-rules ieee-802.1 customup-rw
```

Assign a congestion notification profile or a forwarding class set globally to all interfaces using the **set class-of-service interfaces all** statement. For example, to assign a forwarding class set named **be\_fcset** to all interfaces:

```
[edit class-of-service interfaces]
user@switch# set all forwarding-class-set be_fcset
```



**NOTE:** If there is an existing CoS configuration of any type on an interface, the global configuration is not applied to that particular interface. The global configuration is applied to all interfaces that do not have an existing CoS configuration.

For example, if you configure a rewrite rule, assign it to interfaces **xe-0/0/20.0** and **xe-0/0/22.0**, and then configure a congestion notification profile and apply it to all interfaces, the congestion notification profile is applied to every interface except **xe-0/0/20** and **xe-0/0/22**.

- Related Documentation**
- [Monitoring Interfaces That Have CoS Components on page 5917](#)
  - [Understanding Junos CoS Components on page 5436](#)
  - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. QFX Systems support three DCBX modes:

- Autonegotiation—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- IEEE DCBX—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- DCBX Version 1.01—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric Node devices come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** QFX Systems do not support pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00. If a QFX Series interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

- Related Documentation**
- [Configuring DCBX Autonegotiation on page 5204](#)
  - [Disabling the ETS Recommendation TLV on page 5207](#)
  - [Understanding DCBX on page 5051](#)
  - [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
  - [show dcbx neighbors on page 5300](#)

## Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to

the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

---

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE-FC gateway, it does not send or receive FCoE Initialization Protocol (FIP) packets.
- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

#### Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Disabling the ETS Recommendation TLV on page 5207](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Using DCBX Protocol to Lower Costs](#)

## Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



**NOTE:** Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

### Related Documentation

- [Configuring the DCBX Mode on page 5203](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Understanding DCBX on page 5051](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)





**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

#### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [show dcbx neighbors on page 5300](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [aliases] [bit-patterns]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[001 101]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5300](#)

## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

```
[edit protocols]
```

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209](#)
- [Configuring DCBX Autonegotiation on page 5204](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5300](#)

## Configuration Statements

- [application \(Application Maps\) on page 5818](#)
- [application \(Applications\) on page 5819](#)
- [application-map on page 5820](#)
- [application-maps on page 5821](#)
- [applications \(Applications\) on page 5822](#)
- [applications \(DCBX\) on page 5823](#)

- [buffer-partition \(Egress\) on page 5824](#)
- [buffer-partition \(Ingress\) on page 5826](#)
- [buffer-size on page 5828](#)
- [cable-length \(Congestion Notification\) on page 5830](#)
- [class-of-service on page 5831](#)
- [class \(Forwarding Classes\) on page 5835](#)
- [class \(Forwarding Class Sets\) on page 5836](#)
- [classifiers on page 5837](#)
- [code-point \(Fibre Channel Interfaces\) on page 5838](#)
- [code-point \(Input Congestion Notification\) on page 5839](#)
- [code-point \(Output Congestion Notification\) on page 5840](#)
- [code-point \(Rewrite Rules\) on page 5841](#)
- [code-point-aliases on page 5841](#)
- [code-points \(Application Maps\) on page 5842](#)
- [code-points \(CoS\) on page 5842](#)
- [configured-flow-control on page 5843](#)
- [congestion-notification-profile on page 5844](#)
- [dcbx on page 5846](#)
- [dcbx-version on page 5847](#)
- [destination-port \(Applications\) on page 5848](#)
- [disable \(DCBX\) on page 5849](#)
- [drop-probability on page 5850](#)
- [drop-profile on page 5851](#)
- [drop-profile-map on page 5851](#)
- [drop-profiles on page 5852](#)
- [dscp on page 5853](#)
- [dscp-ipv6 on page 5855](#)
- [dscp-code-point on page 5856](#)
- [egress \(Buffer Configuration\) on page 5857](#)
- [enhanced-transmission-selection on page 5858](#)
- [ether-type on page 5859](#)
- [fill-level on page 5860](#)
- [flow-control on page 5861](#)
- [flow-control-queue \(Output Congestion Notification\) on page 5862](#)
- [forwarding-class on page 5863](#)
- [forwarding-class \(Fibre Channel Interfaces\) on page 5864](#)
- [forwarding-class \(Host Outbound Traffic\) on page 5865](#)

- [forwarding-class-set \(Node Device\) on page 5866](#)
- [forwarding-class-sets on page 5866](#)
- [forwarding-classes on page 5867](#)
- [guaranteed-rate on page 5869](#)
- [host-outbound-traffic on page 5870](#)
- [ieee-802.1 on page 5871](#)
- [ieee-802.1 \(Fibre Channel Interfaces\) on page 5872](#)
- [ieee-802.1 \(Input Congestion Notification\) on page 5873](#)
- [ieee-802.1 \(Output Congestion Notification\) on page 5874](#)
- [import on page 5875](#)
- [ingress \(Buffer Configuration\) on page 5876](#)
- [input \(Congestion Notification\) on page 5877](#)
- [input \(Fibre Channel Interfaces\) on page 5878](#)
- [interface \(DCBX\) on page 5879](#)
- [interfaces \(Class of Service\) on page 5880](#)
- [interpolate on page 5881](#)
- [loss-priority \(Classifiers\) on page 5882](#)
- [loss-priority \(Drop Profiles\) on page 5883](#)
- [loss-priority \(Rewrite Rules\) on page 5884](#)
- [multi-destination on page 5885](#)
- [mru on page 5886](#)
- [output \(Congestion Notification\) on page 5887](#)
- [output-traffic-control-profile on page 5888](#)
- [pfc \(Input Congestion Notification\) on page 5889](#)
- [policy-options on page 5890](#)
- [priority \(Schedulers\) on page 5891](#)
- [priority-flow-control on page 5892](#)
- [protocol \(Applications\) on page 5893](#)
- [protocol \(Drop Profile Map\) on page 5894](#)
- [queue-num on page 5895](#)
- [recommendation-tlv on page 5896](#)
- [rewrite-rules on page 5897](#)
- [rewrite-value \(Fibre Channel Interfaces\) on page 5898](#)
- [rx-buffers on page 5900](#)
- [scheduler \(Node Device\) on page 5901](#)
- [scheduler-map on page 5901](#)
- [scheduler-maps on page 5902](#)

- [schedulers on page 5903](#)
- [shaping-rate on page 5904](#)
- [shared-buffer on page 5906](#)
- [traceoptions \(Class of Service\) on page 5908](#)
- [traffic-control-profiles on page 5910](#)
- [transmit-rate on page 5911](#)
- [tx-buffers on page 5913](#)
- [unit on page 5914](#)

---

## application (Application Maps)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>application <i>application-name</i> {<br/>    code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Hierarchy Level          | [edit policy-options <a href="#">application-maps</a> <i>application-map-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description              | Add an application to an application map and define the application's code points.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Options                  | <i>application-name</i> —Name of the application.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application (Applications)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application <i>application-name</i> {     <i>destination-port</i> <i>port-value</i>;     <i>protocol</i> (tcp   udp);     <i>ether-type</i> <i>type</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit applications]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure properties to define an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5207</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application-map <i>application-map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify an application map to apply to an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>application-map-name</i> —Name of the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5210</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |



## application-maps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application-maps <i>application-map-name</i> {   application <i>application-name</i> {     code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit policy-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Define an application map by specifying the applications that belong to the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><i>application-map-name</i>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## applications (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>  application application-name {<br/>    destination-port port-value;<br/>    protocol (tcp   udp);<br/>    ether-type type;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Define applications that DCBX advertises.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5207</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

---

## applications (DCBX)

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for the EX Series                                                  |
| <b>Description</b>              | Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.                                                                                      |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li></ul> |

## buffer-partition (Egress)

**Syntax** `buffer-partition (lossless | lossy | multicast) {  
percent percent;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer egress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The egress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best-effort, lossless, and multicast traffic (multicast also includes broadcast and destination lookup fail traffic). To better support the mix of traffic on your network, you can optimize the allocation of egress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three egress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ports. All ports on the switch receive the same allocation of egress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default egress buffer partition shown in [Table 548 on page 5824](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 548: Default Egress Shared Buffer Partitioning**

| Lossless Partition | Lossy Partition | Multicast Partition |
|--------------------|-----------------|---------------------|
| 50%                | 31%             | 19%                 |

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages must always total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless egress traffic.

**lossy**—Shared buffer space for best-effort unicast egress traffic.

**multicast**—Shared buffer space reserved for all multicast (including broadcast and destination lookup fail) egress traffic.

**percent percent**—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossy, or multicast buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interfaces—To view this statement in the configuration.       |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 5803</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 5527</a></li></ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## buffer-partition (Ingress)

**Syntax** `buffer-partition (lossless | lossless-headroom | lossy) {  
percent percentage;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer ingress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The ingress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best effort (best-effort unicast and multicast) and lossless traffic. To better support the mix of traffic on your network, you can optimize the allocation of ingress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three ingress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ingress traffic. All ports on the switch receive the same allocation of ingress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default ingress buffer partition shown in [Table 549 on page 5826](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 549: Default Ingress Shared Buffer Partitioning**

| Lossless Partition | Lossless-Headroom Partition | Lossy Partition |
|--------------------|-----------------------------|-----------------|
| 9%                 | 45%                         | 46%             |

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages always must total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless ingress traffic.

**lossless-headroom**—Shared buffer space reserved to store packets received while either an 802.3x Ethernet PAUSE or a priority-based flow control (PFC) pause is asserted. (When an ingress interface pauses traffic, it must have the buffer space to store all of the packets currently in the buffer, and also all of the packets received before the connected peer stops sending traffic and the wire is cleared of packets.)

**lossy**—Shared buffer space for best-effort ingress traffic.

**percent *percent***—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossless-headroom, or lossy buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## buffer-size

---

**Syntax** `buffer-size (percent percent | remainder);`

**Hierarchy Level** [edit `class-of-service schedulers scheduler-name`]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.

However, the **buffer-size** statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two ports and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the **transmit-rate**).

You configure the proportion of port dedicated buffers allocated to a particular output queue using the following process:

1. Configure a scheduler and set the **buffer-size** option to match the scheduler **transmit-rate** value.
2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.

For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the `fcoe` forwarding class, and the `fcoe` forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the **fcoe** forwarding class.

3. Associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic.
4. Associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic.

Queue 3, which is mapped to the `fcoe` forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the **buffer-size** statement.



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

---



**Default** The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the **remainder** option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling, as shown in [Table 550 on page 5829](#). The default buffer size is the same as the default transmit rate for each default queue:

**Table 550: Default Output Queue Buffer Sizes**

| Queue Number | Forwarding Class | Transmit Rate | Buffer Size |
|--------------|------------------|---------------|-------------|
| 0            | best-effort      | 5%            | 5%          |
| 3            | fcoe             | 35%           | 35%         |
| 4            | no-loss          | 35%           | 35%         |
| 7            | network-control  | 5%            | 5%          |
| 8            | mcast            | 20%           | 20%         |

Because the default scheduler includes only five forwarding classes, only the queues mapped to those forwarding classes receive dedicated buffers from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

**Options** **percent percent**—Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.

**remainder**—Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with **remainder** as the buffer size).


**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
  - [Understanding CoS Buffer Configuration on page 5527](#)

---

## cable-length (Congestion Notification)

---

|                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                 | <code>cable-length <i>cable-length-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                        | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i></a> <a href="#">input</a> ]                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                    | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                            | <p>Specify the length of the cable between the interface and its peer interface in meters. The system uses the cable length and the maximum receive unit (MRU) to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The the shorter the cable length and lower the MRU, the less headroom buffer space is required for PFC.</p>                                                                                                                             |
| <div> <b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                | The default cable length value is 100 meters (approximately 328 feet).                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                | <code><i>cable-length-value</i></code> —Length of the cable in meters. (Generally from 1 to 300 meters, but there is no configuration restriction.)                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                               | <code>interfaces</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li></ul> |

## class-of-service

```

Syntax class-of-service {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
 import (classifier-name | default);
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 }
 }
 code-point-aliases {
 (dscp | dscp-ipv6 | ieee-802.1) {
 alias-name bits;
 }
 }
 congestion-notification-profile profile-name {
 input {
 ieee-802.1 {
 code-point [code-point-bits] {
 pfc {
 mru mru-value;
 }
 }
 }
 cable-length cable-length-value;
 }
 output {
 ieee-802.1 {
 code-point [code-point-bits] {
 flow-control-queue [queue | list-of-queues];
 }
 }
 }
 }
 drop-profiles {
 profile-name {
 interpolate {
 fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;
 }
 }
 }
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 forwarding-class class-name {
 scheduler scheduler-name;
 }
 forwarding-class-sets forwarding-class-set-name {
 class class-name;
 }
 }

```

```
}
forwarding-classes {
 class {
 class-name {
 queue-num queue-number <no-loss>;
 }
 }
}
host-outbound-traffic {
 forwarding-class class-name;
 dscp-code-point code-point;
}
interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class lossless-forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1 {
 code-point code-point-bits;
 }
 }
 }
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 }
 }
}
multi-destination {
 classifiers {
 (dscp | ieee-802.1) classifier-name;
 }
}
rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
 import (rewrite-name | default);
 forwarding-class class-name {
 loss-priority priority code-point (alias | bits);
 }
 }
}
scheduler-map-forwarding-class-sets {
 fabric-scheduler-map-name {
 forwarding-class-set fabric-forwarding-class-set-name scheduler scheduler-name;
 }
}
```

```

scheduler-maps {
 map-name {
 forwarding-class class-name scheduler scheduler-name;
 }
}
schedulers {
 scheduler-name {
 buffer-size (percent percentage | remainder);
 drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile
 drop-profile-name;
 priority priority;
 shaping-rate (rate | percent percentage);
 transmit-rate (percent percentage);
 }
}
shared-buffer {
 egress {
 percent percent;
 buffer-partition (lossless | lossy | multicast) {
 percent percent
 }
 }
 ingress {
 percent percent;
 buffer-partition (lossless | lossless-headroom | lossy) {
 percent percent
 }
 }
}
traffic-control-profiles profile-name {
 guaranteed-rate(rate| percent percentage);
 scheduler-map map-name;
 shaping-rate (rate| percent percentage);
}
}

```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure class-of-service parameters on the switch.

The remaining statements are explained separately.

**Default** If you do not configure any CoS features, the default CoS settings are used.


**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Assigning CoS Components to Interfaces on page 5807](#)
- [Defining CoS Unicast BA Classifiers on page 5784](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)
- [Defining CoS Code-Point Aliases on page 5783](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Configuring CoS Drop Profile Maps on page 5787](#)
- [Defining CoS Forwarding Class Sets on page 5789](#)
- [Defining CoS Forwarding Classes on page 5787](#)
- [Defining CoS Rewrite Rules on page 5805](#)
- [Defining CoS Queue Schedulers on page 5789](#)
- [Configuring CoS Fabric Forwarding Class Set Scheduler Maps \(Fabric Scheduler to Fabric FC-Set Mapping\)](#)
- [Configuring CoS Tail-Drop Profiles on page 5786](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 5794](#)
- [Overview of Junos OS CoS for the QFX Series on page 5412](#)

## class (Forwarding Classes)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre> class {   class-name {     queue-num queue-number &lt;no-loss&gt;;   } } </pre>                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit <a href="#">class-of-service forwarding-classes</a> ]                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>No-loss option introduced in Junos OS Release 12.3 for the QFX Series.</p>                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Map one or more forwarding classes to a single queue. You can map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class.</p> |
| <div>  <p><b>NOTE:</b> If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does <i>not</i> receive lossless treatment.</p> <p>If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.</p> </div> |                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p><b>class-name</b> —Name of the forwarding class.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Forwarding Classes on page 5668</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes on page 5469</a></li> </ul>                                                                                                                                         |

## class (Forwarding Class Sets)

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>class <i>class-name</i>;</code>                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service forwarding-class-sets</a> <i>forwarding-class-set-name</i> ]                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.                                                                                      |
| <b>Options</b>                  | <i>class-name</i> —Name of the forwarding class.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 5671</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5474</a></li></ul> |



## classifiers

|                                         |                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                           | <pre> classifiers {     (dscp   dscp-ipv6   ieee-802.1) classifier-name {         import (classifier-name   default);         forwarding-class class-name {             loss-priority level {                 code-points [ aliases ] [ bit-patterns ];             }         }     } } </pre>                        |
| <b>Multidestination BA Classifiers</b>  | <pre> classifiers {     (dscp   ieee-802.1) classifier-name; } </pre>                                                                                                                                                                                                                                                 |
| <b>Interface Classifier Association</b> | <pre> classifiers {     (dscp   dscp-ipv6   ieee-802.1) (default   classifier-name); } </pre>                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                  | <pre> [edit class-of-service], [edit class-of-service multi-destination], [edit class-of-service interfaces interface-name unit logical-unit-number] </pre>                                                                                                                                                           |
| <b>Release Information</b>              | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>                      | Define a unicast or multidestination CoS behavior aggregate (BA) classifier.                                                                                                                                                                                                                                          |
| <b>Options</b>                          | The statements are explained separately.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>         | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                              |
| <b>Related Documentation</b>            | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li> <li>• <a href="#">Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers on page 5661</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li> </ul> |

## code-point (Fibre Channel Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point <i>code-point-bits</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> <a href="#">rewrite-value</a> <a href="#">input ieee-802.1</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure the IEEE 802.1p code point value assigned to all traffic received from the Fibre Channel (FC) network on the specified FC interface (NP_Port). When native FC traffic from the FC SAN arrives at the NP_Port interface, the NP_Port interface encapsulates it in Ethernet to create FCoE packets before forwarding the traffic onto the FCoE network. Instead of using the default value of priority 3 (code point 011) for the FCoE traffic, the interface rewrites the IEEE 802.1p code point to the value specified in the rewrite value code points.</p> <p>After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.</p> <p>FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).</p> <p>The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.</p> |
| <b>Options</b>                  | <i>code-point-bits</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">forwarding-class (Fibre Channel Interfaces) on page 5864</a></li> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## code-point (Input Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | code-point [ <i>code-point-bits</i> ] {<br>pfc {<br>mru <i>mru-value</i> ;<br>}<br>}                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> ]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Enable priority-based flow control (PFC) on an IEEE 802.1p code point (priority).                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>code-point-bits</b>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5113</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> </ul> |

## code-point (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point [ <i>code-point-bits</i> ] {<br/>    <i>flow-control-queue</i> [ <i>queue</i>   <i>list-of-queues</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <i>class-of-service congestion-notification-profile profile-name output ieee-802.1</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the IEEE 802.1p code point bits that identify the traffic you want to enable for priority-based flow control (PFC) pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | <p>By default, IEEE 802.1p priorities 3 and 4 (code points 011 and 100, respectively) are enabled for PFC pause on all Ethernet interfaces. If you explicitly configure priorities to pause and the output queues on which to enable pause, the explicit configuration overrides the default configuration. When you apply an explicit output congestion notification profile to an interface, only the priorities and queues specified in the output congestion notification profile are enabled for pause on that interface.</p> <p>For example, if you configure an output congestion notification profile that specifies priority 2 (code point 010), then traffic with IEEE 802.1p priority 2 is paused on the configured output queue during periods of congestion. However, traffic with priority 3 and priority 4 is not programmed to pause, because the explicit configuration overwrites the default configuration, and the explicit configuration does not pause priority 3 and priority 4. If you configure an explicit output congestion notification profile, all of the priorities you want to enable for PFC and all of the output queues you want to pause must be explicitly configured.</p> |
| <b>Options</b>                  | <p><i>code-point-bits</i>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li></ul>                                                                                                                                                                                                                                                                                             |

## code-point (Rewrite Rules)

|                                 |                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point [ <i>alias</i> ] [ <i>bit-pattern</i> ];</code>                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <a href="#">forwarding-class</a> <a href="#">class-name</a> <a href="#">loss-priority</a> <a href="#">level</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                        |
| <b>Description</b>              | Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.                                                                                                                                                               |
| <b>Options</b>                  | <p><i>alias</i>—Name of the alias.</p> <p><i>bit-pattern</i>—Value of the code-point bits, in decimal form.</p>                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li> </ul>                                                                        |

## code-point-aliases

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-point-aliases {   (<a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>) {     <i>alias-name</i> <i>bits</i>;   } }</pre>                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                |
| <b>Description</b>              | Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.                                                                                                                                         |
| <b>Options</b>                  | <p>(<a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>)—Set the type of classifier for which you are creating an alias.</p> <p><i>alias-name</i>—Name of the code-point alias.</p> <p><i>bits</i> —Value of the code-point bits, in decimal form.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Code-Point Aliases on page 5783</a></li> <li>• <a href="#">Understanding CoS Code-Point Aliases on page 5453</a></li> </ul>                                                                                    |

## code-points (Application Maps)

---


|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Hierarchy Level          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Release Information      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description              | Define one or more code-point aliases or bit sets for an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Options                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5209</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## code-points (CoS)

---

|                          |                                                                                                                                                                                                              |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                         |
| Hierarchy Level          | [edit <b>class-of-service</b> <b>classifiers</b> ( <b>dscp</b>   <b>dscp-ipv6</b>   <b>ieee-802.1</b> ) <i>classifier-name</i> <b>forwarding-class</b> <i>class-name</i> <b>loss-priority</b> <i>level</i> ] |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                            |
| Description              | Configure one or more code-point aliases or bit sets to apply to a forwarding class.                                                                                                                         |
| Options                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                     |
| Required Privilege Level | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                     |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li></ul>                 |

## configured-flow-control

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | configured-flow-control {<br><b>rx-buffers</b> (on   off);<br><b>tx-buffers</b> (on   off);<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <hr/> <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> <hr/> |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">congestion-notification-profile on page 5844</a></li> <li>• <a href="#">flow-control on page 2572</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## congestion-notification-profile

|                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                | <pre> congestion-notification-profile <i>profile-name</i> {   input {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         pfc {           mru <i>mru-value</i>;         }       }     }     cable-length <i>cable-length-value</i>;   }   output {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         flow-control-queue [<i>queue</i>   <i>list-of-queues</i>];       }     }   } } </pre> |
| <b>Interface Congestion Notification Profile Association</b> | <pre> congestion-notification-profile <i>profile-name</i> { </pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                       | <pre> [edit <i>class-of-service</i>], [edit <i>class-of-service interfaces interface-name</i>] </pre>                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                   | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                           | Configure a congestion notification profile to enable priority-based flow control (PFC) on traffic specified by an IEEE 802.1 code point, and apply the profile to an interface.                                                                                                                                                                                                                                               |



**NOTE:** You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually 011).

You can attach a maximum of one congestion notification profile to an interface. There is no limit to the total number of congestion notification profiles you can create.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.



**Options** *profile-name*—Name of the congestion notification profile.

The remaining statements are explained separately.

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 5693](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 5723](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

## dcbx

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>dcbx {<br/>  disable;<br/>  interface (interface-name   all) {<br/>    disable;<br/>    application-map application-map-name;<br/>    applications {<br/>      no-auto-negotiation;<br/>    }<br/>    enhanced-transmission-selection {<br/>      no-auto-negotiation;<br/>      no-recommendation-tlv;<br/>      recommendation-tlv {<br/>        no-auto-negotiation;<br/>      }<br/>    }<br/>  }<br/>  dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);<br/>  priority-flow-control {<br/>    no-auto-negotiation;<br/>  }<br/>}</pre> |
| Hierarchy Level          | [edit <a href="#">protocols</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                      |
| Description              | Configure DCBX properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Options                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li></ul>                                                                                                          |


## dcbx-version

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p> |
| <b>Default</b>                  | The default DCBX mode is autonegotiation.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5300</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li> <li>• <a href="#">Understanding DCBX on page 5051</a></li> </ul>                                                                                                  |

## destination-port (Applications)

---

|                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                             | <code>destination-port <i>port-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                    | [edit applications <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                        | <p>Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers.</p> |
| <hr/>                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <div> <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>.</div> <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                            | <i>port-value</i> —Identifier for the port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                                                                                                                                                                                                           | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5207</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>                            |

## disable (DCBX)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]<br><br>[edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.<br><br>DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li> <li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li> <li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li> </ul> |

## drop-probability

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | drop-probability 0 drop-probability <i>high-value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service drop-profiles profile-name interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>When configuring WRED, map the packet <b>drop-probability</b> to the fullness of a queue (<b>fill-level</b>). You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability.</p> |
| <b>Options</b>                  | <p>0—Probability that packets will drop at the lowest <b>fill-level</b> value. This is always zero, because until the queue reaches the specified low <b>fill-level</b> value, no packets are scheduled to drop.</p> <p><b>high-value</b>—The maximum probability that packets will drop before queue fullness exceeds the high value of the queue <b>fill-level</b>, expressed as a percentage. If the queue fills beyond the high <b>fill-level</b> value, all packets drop.</p> <p><b>Range:</b> 0 through 100</p>                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li><li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## drop-profile

|                                 |                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile <i>profile-name</i>;</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> <a href="#">drop-profile-map <i>loss-priority</i></a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> ]                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                    |
| <b>Description</b>              | Define drop profiles for random early detection (RED). When a packet arrives, RED checks the queue fill level specified in the drop profile. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.          |
| <b>Options</b>                  | <i>profile-name</i> —Name of the drop profile.                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 5666</a></li> <li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li> <li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li> </ul> |

## drop-profile-map

|                                 |                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile-map <a href="#">loss-priority</a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> <a href="#">drop-profile <i>drop-profile-name</i></a>;</code>                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> ]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                    |
| <b>Description</b>              | Map a drop profile to a loss priority and protocol for random early detection (RED). When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.          |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 5666</a></li> <li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li> <li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li> </ul> |

## drop-profiles

---

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>drop-profiles {<br/>    profile-name {<br/>        interpolate {<br/>            fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;<br/>        }<br/>    }<br/>}</pre>                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Define drop profiles for weighted random early detection (WRED).</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.</p> |
| <b>Options</b>                  | <p><i>profile-name</i>—Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li><li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li></ul>                                                                                                                                  |




## dscp

|                                                  |                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                    | <pre> dscp classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                                                                                          |
| <b>Code-Point Alias Configuration</b>            | <code>dscp alias-name bit-pattern;</code>                                                                                                                                                                                                                                                                                                                                    |
| <b>Multidestination Classifier Configuration</b> | <code>dscp classifier-name;</code>                                                                                                                                                                                                                                                                                                                                           |
| <b>Interface Classifier Association</b>          | <code>dscp (classifier-name   default);</code>                                                                                                                                                                                                                                                                                                                               |
| <b>Rewrite Rule Configuration</b>                | <pre> dscp rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                                                                                                 |
| <b>Hierarchy Level</b>                           | <pre> [edit class-of-service classifiers], [edit class-of-service code-point-aliases], [edit class-of-service multi-destination classifiers], [edit class-of-service interfaces interface-name unit logical-unit-number classifiers], [edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules], [edit class-of-service rewrite-rules] </pre> |
| <b>Release Information</b>                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                               | Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                   | <p><b>classifier-name</b>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                  | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                     | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li> <li>• <a href="#">Defining CoS Code-Point Aliases on page 5783</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> </ul>                                                                                                         |

- [Assigning CoS Components to Interfaces on page 5807](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)

## dscp-ipv6

|                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                   | <pre> dscp-ipv6 classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                              |
| <b>Code-Point Alias Configuration</b>                                                                                                                                                                                                                                                                           | <pre> dscp-ipv6 alias-name bit-pattern; </pre>                                                                                                                                                                                                                                                                        |
| <b>Interface Classifier Association</b>                                                                                                                                                                                                                                                                         | <pre> dscp-ipv6 (classifier-name   default); </pre>                                                                                                                                                                                                                                                                   |
| <b>Rewrite Rule Configuration</b>                                                                                                                                                                                                                                                                               | <pre> dscp-ipv6 rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                          | <pre> [edit class-of-service classifiers], [edit class-of-service code-point-aliases], [edit class-of-service interfaces interface-name unit logical-unit-number classifiers], [edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules], [edit class-of-service rewrite-rules] </pre> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                              | Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.                                                                                                                                                                                                                     |
| <div>  <p><b>NOTE:</b> There is no DSCP IPv6 classifier for multdestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.</p> </div> |                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                                                                                                  | The statements are explained separately.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                 | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                              |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li> <li>• <a href="#">Defining CoS Code-Point Aliases on page 5783</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> </ul>                                                  |

- [Assigning CoS Components to Interfaces on page 5807](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Rewrite Rules on page 5549](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5460](#)


---

## dscp-code-point

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dscp-code-point</code> <i>code-point</i> ;                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">host-outbound-traffic</a> ]                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                     |
| <b>Description</b>              | Set the value of the DSCP code point in the type of service (ToS) field of the packet generated by the Routing Engine (host).                                                                                                                         |
| <b>Options</b>                  | <b>code-point</b> —Six-bit DSCP code point value.                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 5793</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5451</a></li></ul> |

## egress (Buffer Configuration)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>egress {   percent <i>percent</i>;   <b>buffer-partition</b> (lossless   lossy   multicast) {     percent <i>percent</i>;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service shared-buffer</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure the global shared buffer pool allocation for egress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for other buffers such as dedicated buffers (buffers allocated permanently to ports).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports. However, on a port, you can configure the portion of dedicated port buffer space allocated to each queue in the scheduler configuration using the <b>buffer-size</b> option.)</p> |
|                                 | <div>  <p><b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|                                 | <p>You can also partition the shared buffer pool to adjust the egress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>percent <i>percent</i></b>—Percentage of available egress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

---

## enhanced-transmission-selection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>enhanced-transmission-selection {<br/>    no-auto-negotiation;<br/>    no-recommendation-tlv;<br/>    recommendation-tlv {<br/>        no-auto-negotiation;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p> |
| <b>Options</b>                  | <p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li></ul>                                                                                                                                                                |

## ether-type

---

|                            |                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>ether-type <i>ether-type</i>;</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                     |
| <b>Description</b>         | Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes. |



**NOTE:** To create a FIP application, use the EtherType 0x8914.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>type</i> —Identifier for the EtherType.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5207</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li> </ul> |


## fill-level

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>fill-level <i>low-value</i> fill-level <i>high-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service drop-profiles <i>profile-name</i> interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>When configuring random early detection (RED), map the fullness of a queue to a packet <a href="#">drop-probability</a> value. You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability.</p> |
| <b>Options</b>                  | <p><b><i>low-value</i></b>—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.</p> <p><b>Range:</b> 0 through 100</p> <p><b><i>high-value</i></b>—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.</p> <p><b>Range:</b> 0 through 100</p>                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li><li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



## flow-control

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> |
|                                 | <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>                                                                                                                                                                                                          |
|                                 | <ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | Flow control is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 5843</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2533</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> <li>• <i>Junos® OS Network Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## flow-control-queue (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>flow-control-queue [ <i>queue</i>   <i>list-of-queues</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> output ieee-802.1 code-point <i>code-point-bits</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify one or more output queues to pause, to support priority-based flow control (PFC). The specified queues pause when the interface receives a PFC frame with a matching IEEE 802.1p code point.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | <p>Queue 3 (mapped to the fcoe forwarding class) and queue 4 (mapped to the no-loss forwarding class) are programmed as flow control queues to pause. No other output queues are programmed to pause by default.</p> <p>If you configure flow control queues explicitly, only the queues that you specify are programmed to pause. The explicit flow control queue to pause configuration overrides the default setting, so the queues paused in the default configuration are no longer paused by default.</p> <p>For example, if you configure queue 2 as a flow control queue, then queue 2 pauses when congestion occurs, but queues 3 and 4 do not pause because they were not explicitly specified. To enable pause on output queues 2, 3, and 4, you must explicitly configure all three of the queues as flow control queues.</p> <p>The same behavior applies to the IEEE 802.1p code points (priorities) on which PFC is enabled. By default, priorities 3 (011) and 4 (100) are enabled for PFC pause. If you explicitly configure flow control queues to pause, you must also explicitly configure pause for each priority (code point) that you want to pause, because the explicit configuration overrides the default configuration.</p> |
| <b>Options</b>                  | <code>[ <i>queue</i>   <i>list-of-queues</i> ]</code> —The output queue or a list of output queues to pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                |

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506](#)

## forwarding-class

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                      | <pre>forwarding-class <i>class-name</i> {     loss-priority <i>level</i> {         code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];     } }</pre>                                                                                                                                                                                                                                                                                                                                                         |
| Rewrite Rule Configuration  | <pre>forwarding-class <i>class-name</i> {     loss-priority <i>level</i> {         code-point [ <i>aliases</i> ] [ <i>bit-patterns</i> ];     } }</pre>                                                                                                                                                                                                                                                                                                                                                          |
| Scheduler Map Configuration | <pre>forwarding-class <i>class-name</i> {     scheduler <i>scheduler-name</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Interface Configuration     | <pre>forwarding-class <i>class-name</i>;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Hierarchy Level             | <pre>[edit class-of-service classifiers (dscp   dscp-ipv6   ieee-802.1) <i>classifier-name</i>], [edit class-of-service rewrite-rules] (dscp   dscp-ipv6   ieee-802.1) <i>rewrite-name</i>], [edit class-of-service scheduler-maps <i>map-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>                                                                                                                                                              |
| Release Information         | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description                 | Configure forwarding class name and option values (classifier configuration), map rewrite rules to forwarding classes (rewrite rules), map forwarding classes to schedulers (scheduler maps), or map forwarding classes to logical interfaces (interfaces).                                                                                                                                                                                                                                                      |
| Options                     | <p><b><i>class-name</i></b>—Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| Required Privilege Level    | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation       | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Forwarding Classes on page 5668</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes on page 5469</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 5549</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul> |

## forwarding-class (Fibre Channel Interfaces)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>forwarding-class <i>lossless-forwarding-class-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | <p>Configure a Layer 3 fixed classifier on a Fibre Channel (FC) interface. The fixed classifier places all traffic received from the FC network into the specified forwarding class. The forwarding class must be lossless. (That is, the forwarding class must be either the default <b>fcoe</b> or <b>no-loss</b> forwarding class, or the forwarding class must be configured with the <b>no-loss</b> drop attribute.) If you attempt to specify a lossy forwarding class, the system returns a commit error.</p> <p>FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. If the IEEE 802.1p code point value that the Ethernet network uses for FCoE traffic is different than code point 3, you can rewrite the code point to the value used in your Ethernet (FCoE) network. The lossless forwarding class specified in the fixed classifier uses the <b>rewrite-value</b> statement as the IEEE 802.1p code point (priority) for FCoE traffic on the FCoE network.</p> <p>To rewrite the code point value, include the <b>rewrite-value input ieee code-point code-point-bits</b> statement at the [edit <a href="#">class-of-service interfaces</a> <i>fc-interface-name</i>] hierarchy level.</p> |



**NOTE:** If you are not using the default configuration (priority 3 for FCoE traffic), the lossless forwarding class specified in the FC interface fixed classifier must be mapped to the IEEE 802.1p code point specified in the rewrite value statement.

In order to avoid fate sharing (separate flows that affect each other's throughput), the code point (priority) used for the lossless forwarding class (the code point specified in the rewrite value statement) should be the only code point classified to that forwarding class (at the [edit [class-of-service classifiers](#)] hierarchy level). For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class `fcoe_fc1`, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the `fcoe_fc1` forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class `fcoe_fc1`, then congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into forwarding class `fcoe_fc1`, so the traffic from both priorities shares the same fate.

**Options** *lossless-forwarding-class-name*—Name of the lossless forwarding class.

|                                 |                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li> </ul> |

## forwarding-class (Host Outbound Traffic)

---

|                                 |                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forwarding-class <i>class-name</i>;</code>                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">host-outbound-traffic</a> ]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                        |
| <b>Description</b>              | Define forwarding class name for outbound host traffic (traffic generated by the Routing Engine).                                                                                                                                                        |
| <b>Options</b>                  | <i>class-name</i> —Name of the forwarding class.                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 5793</a></li> <li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5451</a></li> </ul> |

## forwarding-class-set (Node Device)

---

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forwarding-class-set forwarding-class-set-name {<br/>    output-traffic-control-profile profile-name;<br/>}</code>                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service interfaces interface-name</code> ]                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply a previously defined forwarding class set to an output traffic control profile.                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>forwarding-class-set-name</i> —Name of the forwarding class set.<br><br>The remaining statement is explained separately.                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 5807</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5474</a></li></ul> |

## forwarding-class-sets

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forwarding-class-sets forwarding-class-set-name {<br/>    class class-name;<br/>}</code>                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service</code> ]                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Assign forwarding classes to forwarding class sets (priority groups).                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>forwarding-class-set-name</i> —Name of the forwarding class set.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 5671</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5474</a></li></ul> |

## forwarding-classes

**Syntax**

```
forwarding-classes {
 class {
 class-name {
 queue-num queue-number <no-loss>;
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Map one or more forwarding classes to a single queue. You can configure up to 12 forwarding classes (8 unicast forwarding classes on queues 0 through 7 and 4 multidestination forwarding classes on queues 8 through 11) and map them to queues. You can map multiple forwarding classes to a single queue using the **class** statement. All forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if



you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

.....

|                                 |                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                           |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 5668</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 5469</a></li></ul> |



## guaranteed-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>guaranteed-rate (rate  percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service traffic-control-profiles</code> <i>traffic-control-profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.                                                                                      |
|                                 | <div>  <p><b>NOTE:</b> You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.</p> </div>                                                                                                                                                                                                     |
| <b>Default</b>                  | If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                 | <div>  <p><b>NOTE:</b> If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.</p> </div>                                                                                                 |
| <b>Options</b>                  | <p><b>percent <i>percentage</i></b>—Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Minimum transmission rate allocated to the forwarding class set or logical interface, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 10,000,000,000 bps</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
  - [Understanding CoS Traffic Control Profiles on page 5496](#)
  - [output-traffic-control-profile on page 5888](#)

---

## host-outbound-traffic

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>host-outbound-traffic {<br/>    forwarding-class <i>class-name</i>;<br/>    dscp-code-point <i>code-point</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Allow queue selection for traffic generated by the Routing Engine (host). The selected queue must be configured properly. You can also configure specific DSCP code point bits for the type of service (ToS) field of the generated packets. This configuration does not affect transit packets or incoming packets. This is a global configuration that only affects packets originating on the Routing Engine. If you do not configure an output queue for host outbound traffic, the switch uses the default queue mapping. |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 5793</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5451</a></li></ul>                                                                                                                                                                                                                                                                          |

## ieee-802.1

|                                           |                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                    | <pre> ieee-802.1 classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                                                                                    |
| Code-Point Alias Configuration            | <pre> ieee-802.1 alias-name bit-pattern; </pre>                                                                                                                                                                                                                                                                                                                              |
| Multidestination Classifier Configuration | <pre> ieee-802.1 classifier-name; </pre>                                                                                                                                                                                                                                                                                                                                     |
| Interface Classifier Association          | <pre> ieee-802.1 (classifier-name   default); </pre>                                                                                                                                                                                                                                                                                                                         |
| Rewrite Rule Configuration                | <pre> ieee-802.1 rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } } </pre>                                                                                                                                                           |
| Hierarchy Level                           | <pre> [edit class-of-service classifiers], [edit class-of-service code-point-aliases], [edit class-of-service multi-destination classifiers], [edit class-of-service interfaces interface-name unit logical-unit-number classifiers], [edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules], [edit class-of-service rewrite-rules] </pre> |
| Release Information                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                            |
| Description                               | Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.                                                                                                                                                                                                      |
| Options                                   | <p><b>classifier-name</b>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                              |
| Required Privilege Level                  | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                          |
| Related Documentation                     | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li> <li>• <a href="#">Defining CoS Code-Point Aliases on page 5783</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> </ul>                                                                                                         |

- [Assigning CoS Components to Interfaces on page 5807](#)
- [Understanding CoS Classifiers on page 5455](#)
- [Understanding CoS Rewrite Rules on page 5549](#)

---

## ieee-802.1 (Fibre Channel Interfaces)

---

**Syntax**    `ieee-802.1 {  
                  code-point code-point-bits;  
                  }`

**Hierarchy Level**    [edit [class-of-service interfaces](#) *fibre-channel-interface-name* [rewrite-value input](#)]

**Description**    Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).

The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

The statement is described separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**    • [forwarding-class \(Fibre Channel Interfaces\) on page 5864](#)  
                                  • [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)  
                                  • [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)  
                                  • [Understanding CoS Classifiers on page 5455](#)

## ieee-802.1 (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> ieee-802.1 {   code-point [code-point-bits] {     pfc {       mru mru-value;     }   } } </pre>                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile</a> <i>profile-name</i> <a href="#">input</a> ]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point.                                                                                                                                                                                                               |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5113</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> </ul> |

## ieee-802.1 (Output Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ieee-802.1 {<br/>    code-point [ code-point-bits ] {<br/>        flow-control-queue [ queue   list-of-queues ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>output</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li></ul> |


## import

---

|                                 |                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import (<i>import</i>   default);</code>                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ],<br>[edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ]          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify a default or previously defined classifier.                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>import</i></b>—Name of the classifier mapping configured at the [edit <a href="#">class-of-service classifiers</a>] hierarchy level.</p> <p><b>default</b>—Default classifier mapping.</p> <p>The remaining statements are explained separately.</p>                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 5549</a></li> </ul> |

## ingress (Buffer Configuration)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ingress {<br/>    <b>buffer-partition</b> (lossless   lossless-headroom   lossy) {<br/>        percent <i>percent</i>;<br/>    }<br/>    percent <i>percent</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service shared-buffer</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure the global shared buffer pool allocation for ingress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for buffers such as dedicated buffers (buffers allocated permanently to ports) and headroom buffers (buffers that help prevent packet loss on lossless flows).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared ingress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)</p> |
|                                 | <div> <b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|                                 | <p>You can also partition the shared buffer pool to adjust the ingress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>percent <i>percent</i></b>—Percentage of available ingress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 5803](#)
- [Understanding CoS Buffer Configuration on page 5527](#)

## input (Congestion Notification)

**Syntax**

```
input {
 ieee-802.1 {
 code-point [code-point-bits] {
 pfc {
 mru mru-value;
 }
 }
 }
 cable-length cable-length-value;
}
```

**Hierarchy Level** [edit [class-of-service congestion-notification-profile](#) *profile-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure priority-based flow control (PFC) on incoming traffic.

**Options** The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## input (Fibre Channel Interfaces)

---

**Syntax**    input {  
              ieee-802.1p {  
                  code-point *code-point-bits*;  
              }  
          }

**Hierarchy Level**    [edit [class-of-service interfaces](#) *fibre-channel-interface-name* [rewrite-value](#)]

**Description**    Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).

The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

The statements are described separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [forwarding-class \(Fibre Channel Interfaces\) on page 5864](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)
- [Understanding CoS Classifiers on page 5455](#)

## interface (DCBX)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface (<i>interface-name</i>   all) {   disable;   application-map <i>application-map-name</i>;   applications {     no-auto-negotiation;   }   enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {       no-auto-negotiation;     }   }   dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);   priority-flow-control {     no-auto-negotiation;   } } </pre>                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure DCBX properties on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5300</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## interfaces (Class of Service)

```
Syntax interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class lossless-forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1{
 code-point code-point-bits;
 }
 }
 }
 }
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure interface-specific CoS properties for incoming packets.

**Options** *interface-name*—Name of the interface.

The statements are explained separately.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Unicast Classifiers on page 5658](#)
- [Example: Configuring Forwarding Classes on page 5668](#)
- [Example: Configuring Forwarding Class Sets on page 5671](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 5682](#)
- [Assigning CoS Components to Interfaces on page 5807](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Defining CoS Rewrite Rules on page 5805](#)

- [Interfaces Overview on page 2415](#)

## interpolate

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interpolate {   fill-level <i>low-value</i> fill-level <i>high-value</i>;   drop-probability 0 drop-probability <i>high-value</i>; }</pre>                                                       |
| <b>Hierarchy Level</b>          | [edit class-of-service <b>drop-profiles</b> <i>profile-name</i> ]                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| <b>Description</b>              | <p>Specify values for interpolating the relationship between queue fill level and drop probability.</p> <p>The statements are explained separately.</p>                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li> <li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li> </ul> |

## loss-priority (Classifiers)

---

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority level {<br/>    code-points [ aliases ] [ bit-patterns ];<br/>}</code>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i><br><a href="#">forwarding-class</a> <i>class-name</i> ]                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure packet loss priority value for a specific set of code-point aliases and bit patterns.                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 5658</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li></ul>                                                                                                                                |

---

## loss-priority (Drop Profiles)

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority <i>level</i> <i>protocol</i> <i>protocol</i> <b>drop-profile</b> <i>profile-name</i>;</code>                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers scheduler-name drop-profile-map</a> ]                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure packet loss priority value for a drop profile mapped to a system drop profile.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Drop Profile Maps on page 5666</a></li><li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li><li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li></ul>                                              |

## loss-priority (Rewrite Rules)

---

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority <i>level</i> {<br/>    <i>code-point</i> (<i>alias</i>   <i>bit-pattern</i>);<br/>}</code>                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service rewrite-rules</code> ( <code>dscp</code>   <code>dscp-ipv6</code>   <code>ieee-802.1</code> ) <i>rewrite-name</i><br><i>forwarding-class</i> <i>class-name</i> ]                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern.                        |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5549</a></li></ul>                                                                                                                                            |




---

## multi-destination

---

|                                 |                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multi-destination {<br/>  classifiers {<br/>    (dscp   ieee-802.1) classifier-name;<br/>  }<br/>}</pre>                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Define a multicast CoS behavior aggregate (BA) classifier.                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers on page 5661</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 5807</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5455</a></li></ul> |

## mrp

|                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                        | <code>mrp <i>mrp-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                               | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> input ieee-802.1 code-point <i>code-point-bits</i> pfc</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                           | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                   | Configure the maximum receive unit (MRU) of the interface in bytes (incoming packet sizes must be less than or equal to the MRU, or the packets are dropped). The system uses the MRU and the cable length to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The lower the MRU and the shorter the cable length, the less headroom buffer space is required for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <div>  <p><b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                       | <p>For priority 3 traffic, the default MRU value is 2500 bytes.</p> <p>For priority 4 traffic, the default MRU value is 9612 bytes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                       | <b><i>mrp-value</i></b> —Value of the maximum packet receive unit size in bytes (generally from 1500 to 9216 bytes, but there is no configuration restriction).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                      | <p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li> </ul> |

## output (Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {   ieee-802.1 {     code-point [code-point-bits] {       flow-control-queue [queue   list-of-queues];     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure priority-based flow control (PFC) on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li> </ul> |

## output-traffic-control-profile

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>output-traffic-control-profile <i>profile-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">forwarding-class-set</a> <i>forwarding-class-set-name</i> ]                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Apply an output traffic scheduling and shaping profile to a forwarding class set (priority group).                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>profile-name</i> —Name of the traffic-control profile to apply to the specified forwarding class set.                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <code>interfaces</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 5807</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5496</a></li></ul> |

## pfc (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pfc {<br/>    <b>mru</b> <i>mru-value</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> <b>code-point</b> <i>code-point-bits</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Enable and configure ingress interface priority-based flow control (PFC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 5710</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 5701</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 5693</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 5723</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5064</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5506</a></li> </ul> |

## policy-options

---

**Syntax**    `policy-options`  
              `application-maps` *application-map-name* {  
                  `application` *application-name* {  
                      `code-points` [ *aliases* ] [ *bit-patterns* ];  
                  }  
              }  
              `policy-statement` *policy-name* {  
                  `term` *term-name* {  
                      `from` {  
                          `family` *family-name*;  
                          `match-conditions`;  
                          `policy` *subroutine-policy-name*;  
                          `prefix-list` *prefix-list-name*;  
                          `prefix-list-filter` *prefix-list-name* *match-type* <*actions*>;  
                          `route-filter` *destination-prefix* *match-type* <*actions*>;  
                          `source-address-filter` *source-prefix* *match-type* <*actions*>;  
                      }  
                      `to` {  
                          `match-conditions`;  
                          `policy` *subroutine-policy-name*;  
                      }  
                      `then` *actions*;  
                  }  
              }

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 12.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 12.1 for the EX Series.

**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    `storage`—To view this statement in the configuration.  
                                  `storage-control`—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5207](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5144](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5060](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

## priority (Schedulers)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the packet-scheduling drop priority value.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>priority</i></b>—It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Scheduler has low priority.</li> <li>• <b>strict-high</b>—Scheduler has strict high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul>                                                                                                                                                                                               |


## priority-flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>priority-flow-control {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 5795</a></li><li>• <i>Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</i></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5113</a></li><li>• <i>Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</i></li><li>• <i>Understanding Priority-Based Flow Control</i></li><li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li></ul> |



## protocol (Applications)

|                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                            | <code>protocol (tcp   udp);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                   | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                               | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                       | Networking protocol type, which combines with <b>destination-port</b> to identify an application type.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number 3260. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                           | <p><code>tcp</code>—Transmission Control Protocol</p> <p><code>udp</code>—User Datagram Protocol</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                          | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5207</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5060</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## protocol (Drop Profile Map)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol protocol <b>drop-profile</b> profile-name;</code>                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service schedulers scheduler-name drop-profile-map loss-priority</b> (low   medium-high   high)]                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the protocol type for the specified drop profile.                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>protocol</b>—Type of protocol. The protocol can be:</p> <ul style="list-style-type: none"><li>• <b>any</b>—Accept any protocol type.</li></ul> <p>The remaining statement is explained separately.</p>                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Drop Profile Maps on page 5666</a></li><li>• <a href="#">Example: Configuring Tail-Drop Profiles on page 5663</a></li><li>• <a href="#">Understanding CoS Tail-Drop Profiles on page 5545</a></li></ul> |

## queue-num

|                            |                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>queue-num <i>queue-number</i> &lt;no-loss&gt;;</code>                                                                                 |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service forwarding-classes class <i>class-name</i></code> ]                                                            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>No-loss option introduced in Junos OS Release 12.3 for the QFX Series. |
| <b>Description</b>         | Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class.                |

You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) or some or all of the four multdestination (multicast, broadcast, destination lookup fail) forwarding classes to the same multdestination queue (8 through 11), providing that you do not map one forwarding class to more than one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.

|                                 |                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b><i>queue-number</i></b> —Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).<br><br><b><i>no-loss</i></b> —Optional packet drop attribute keyword to configure the forwarding class as lossless. |
| <b>Required Privilege Level</b> | <b>interfaces</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 5668</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 5469</a></li></ul>                                              |

---

## recommendation-tlv

---

|                                 |                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>recommendation-tlv {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable or enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b><i>no-auto-negotiation</i></b> —Disable sending of the ETS recommendation TLV.                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5300</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li></ul>                                                                                                                                                                  |

## rewrite-rules

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1) rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {       loss-priority priority code-point (alias   bits);     }   } }</pre> |
| <b>Interface Association</b>    | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1) rewrite-name; }</pre>                                                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit class-of-service], [edit class-of-service interfaces interface-name unit logical-unit-number]</pre>                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                      |
| <b>Description</b>              | Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.                                                                            |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <pre>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</pre>                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 5805</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 5549</a></li> </ul>                                    |

## rewrite-value (Fibre Channel Interfaces)

```
Syntax rewrite-value {
 input {
 ieee-802.1p {
 code-point code-point-bits;
 }
 }
 }
```

**Hierarchy Level** [edit [class-of-service interfaces](#) *fibre-channel-interface-name*]

**Description** Configure the IEEE 802.1p code point value (priority) for all traffic received from the Fibre Channel (FC) network on the specified FC interface. Instead of using the default priority 3 (011) for FCoE traffic, the priority is rewritten to the specified priority before being forwarded. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011) for FCoE traffic.

The system supports only one IEEE 802.1p code point value per FC interface, so you cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.



**NOTE:** In order to avoid fate sharing (separate flows that affect each others' throughput), the code point used for the rewrite value should be the only code point used for the lossless FCoE forwarding class (the forwarding class used for the fixed classifier on the Fibre Channel interface). When you configure classifiers for ingress Ethernet interfaces, map only the rewrite value code point to the forwarding class.

For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class `fcoe_fc1`, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the `fcoe_fc1` forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class `fcoe_fc1`, then congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into

forwarding class `fcoe_fc1`, so the traffic from both priorities shares the same fate.


.....  
The remaining statements are described separately.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [forwarding-class \(Fibre Channel Interfaces\) on page 5864](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5739](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 5524](#)
- [Understanding CoS Classifiers on page 5455](#)

## rx-buffers

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | rx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit <a href="#">interfaces interface-name ether-options configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
| <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2572</a></li> <li>• <a href="#">tx-buffers on page 5913</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## scheduler (Node Device)

|                                 |                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>scheduler <i>scheduler-name</i>;</code>                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i></a> ]                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                            |
| <b>Description</b>              | Map a scheduler to a forwarding class using a scheduler map.                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>scheduler-name</i> —Name of the scheduler to map to the forwarding class.                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul> |

## scheduler-map

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>scheduler-map <i>map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service traffic-control-profiles <i>traffic-control-profile-name</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Associate a scheduler map with a traffic control profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>map-name</i> —Name of the scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5496</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul> |

## scheduler-maps


---

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>scheduler-maps {<br/>  map-name {<br/>    forwarding-class class-name scheduler scheduler-name;<br/>  }<br/>}</pre>                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify a scheduler map name to map a scheduler configuration to a forwarding class.                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><i>map-name</i>—Name of the scheduler map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li><li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li></ul> |

## schedulers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> schedulers {   scheduler-name {     buffer-size (percent <i>percentage</i>   remainder);     drop-profile-map loss-priority (low   medium-high   high) protocol <i>protocol</i> drop-profile       drop-profile-name;     priority <i>priority</i>;     shaping-rate (<i>rate</i>   percent <i>percentage</i>);     transmit-rate (percent <i>percentage</i>);   } } </pre>           |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify scheduler name and parameter values such minimum bandwidth ( <b>transmit-rate</b> ), maximum bandwidth ( <b>shaping-rate</b> ), and priority ( <b>priority</b> ).                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>scheduler-name</b> —Name of the scheduler.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 5666</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul> |

## shaping-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>shaping-rate (rate   percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers scheduler-name</a> ],<br>[edit <a href="#">class-of-service traffic-control-profiles profile-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate. You specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile, and then applying the traffic control profile and a forwarding class set to an interface.</p> <p>We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).</p> |
|                                 | <div>  <p><b>NOTE:</b> When you set the maximum bandwidth (<a href="#">shaping-rate</a> value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured <a href="#">shaping-rate</a> value.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>percent <i>percentage</i></b>—Shaping rate as a percentage of the available interface bandwidth.<br/> <b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Peak (maximum) rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).<br/> <b>Range:</b> 1000 through 10,000,000,000 bps</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

- [Understanding CoS Traffic Control Profiles on page 5496](#)

## shared-buffer

```
Syntax shared-buffer {
 egress {
 buffer-partition (lossless | lossy | multicast) {
 percent percent
 }
 percent percent;
 }
 ingress {
 percent percent;
 buffer-partition (lossless | lossless-headroom | lossy) {
 percent percent
 }
 }
 }
```

Hierarchy Level [edit [class-of-service](#)]

Release Information Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Configure the global shared buffer pool allocation to ports. Shared buffers are a pool of buffer space that the system can allocate dynamically across all of its ports as memory space is needed. Some buffer space is reserved for dedicated buffers (buffers allocated permanently to ports), headroom buffers (buffers that help prevent packet loss on lossless flows), and other buffers.

Configure the way the system uses the available (user-configurable) buffer space by setting the **shared-buffer** percentage for the ingress buffer pool and for the egress buffer pool.

The percentage you specify is the percentage of available buffer space allocated to the global shared ingress buffer pool or to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until the buffer reprogramming is complete.

You can also partition the ingress shared buffer pool and the egress shared buffer pool to adjust the buffer allocations for different mixes of network traffic (best-effort, lossless, multicast) using the **buffer-partition** statement.

**Options** The statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 5803</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 5527</a></li></ul> |

## traceoptions (Class of Service)

---

|                            |                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt;<br/>    &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt;;<br/>    no-remote-trace<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>         | Set class-of-service (CoS) tracing options.                                                                                                                                                                                                   |



**NOTE:** The `traceoptions` statement is not supported on the QFabric system.

**Default** Traceoptions is disabled.

**Options** **file *filename***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***. The traceoption output continues in a second trace file named ***trace-file.1***. When ***trace-file.1*** reaches its maximum size, output continues in a third file named ***trace-file.2***, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

**Range:** 2 through 1000 files

**Default:** 1 trace file

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all operations.
- **asynch**—Trace asynchronous configuration processing.
- **chassis-scheduler**—Trace chassis stream scheduler processing.
- **cos-adjustment**—Trace CoS rate adjustments.
- **dynamic**—Trace dynamic CoS functions.
- **hardware-database**—Trace the chassis hardware database related processing.
- **init**—Trace initialization events.



- **performance-monitor**—Trace performance monitor counters.
- **process**—Trace configuration processing.
- **restart**—Trace restart processing.
- **route-socket**—Trace route-socket events.
- **show**—Trace show command servicing.
- **snmp**—Trace SNMP-related processing.
- **util**—Trace utilities.

The following are the global tracing options:

- **all**—Perform all tracing operations
- **parse**—Trace parser processing.

**no-remote-trace**—(Optional) Disable remote tracing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.




|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

## traffic-control-profiles

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traffic-control-profiles <i>profile-name</i> {<br/>    <b>guaranteed-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>    <b>scheduler-map</b> <i>map-name</i>;<br/>    <b>shaping-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>}</pre>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>profile-name</b>—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 5671</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 5807</a></li><li>• <a href="#">output-traffic-control-profile on page 5888</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5496</a></li></ul> |

## transmit-rate

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>transmit-rate (rate   percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers <i>scheduler-name</i></code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | Specify the minimum transmission rate or percentage for a queue (forwarding class) scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|                            | <p> <b>NOTE:</b> The <code>transmit-rate</code> setting works only if you also configure the <code>guaranteed-rate</code> in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the <code>transmit-rate</code> statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.</p> <p> <b>NOTE:</b> You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.</p> <p> <b>NOTE:</b> For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.</p> |
| <b>Default</b>             | If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Queue Number        | Default Minimum Guaranteed Bandwidth |
|---------------------|--------------------------------------|
| 0 (best-effort)     | 5 %                                  |
| 1                   | 0                                    |
| 2                   | 0                                    |
| 3 (fcoe)            | 35 %                                 |
| 4 (no-loss)         | 35 %                                 |
| 5                   | 0                                    |
| 6                   | 0                                    |
| 7 (network control) | 5 %                                  |
| 8 (mcast)           | 20 %                                 |
| 9                   | 0                                    |
| 10                  | 0                                    |
| 11                  | 0                                    |

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

**Options** **rate**—Minimum transmission rate for the queue, in bps. You can specify a value in bits-per-second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 1000 through 10,000,000,000 bps

**percent** **percentage**—Minimum percentage of transmission capacity allocated to the queue. A percentage of zero means that there is no minimum bandwidth guarantee for the queue.


**Range:** 0 through 100 percent

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

## tx-buffers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | tx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>on   off</b> —Enable or disable an interface to respond to an Ethernet PAUSE message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2572</a></li> <li>• <a href="#">rx-buffers on page 5900</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 5798](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 5799](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

---

## unit

---

**Syntax**    `unit logical-unit-number {  
              classifiers {  
                  (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);  
              }  
              forwarding-class class-name;  
              rewrite-rules {  
                  (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);  
              }  
          }`

**Hierarchy Level**    [edit **class-of-service interfaces** *interface-name*]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.

**Options**    *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,385

The remaining statements are explained separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**    • [Assigning CoS Components to Interfaces on page 5807](#)

## CHAPTER 65

# Administration

- [Routine Monitoring on page 5915](#)
- [Operational Commands on page 5921](#)

### Routine Monitoring

---

- [Monitoring CoS Classifiers on page 5915](#)
- [Monitoring CoS Forwarding Classes on page 5916](#)
- [Monitoring Interfaces That Have CoS Components on page 5917](#)
- [Monitoring CoS Rewrite Rules on page 5918](#)
- [Monitoring CoS Scheduler Maps on page 5919](#)
- [Monitoring CoS Value Aliases on page 5920](#)

### Monitoring CoS Classifiers

**Purpose** Display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

**Action** To monitor CoS classifiers in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier
```

To monitor a particular classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier name classifier-name
```

To monitor a particular type of classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier type classifier-type
```

**Meaning** [Table 551 on page 5915](#) summarizes key output fields for CoS classifiers.

**Table 551: Summary of Key CoS Classifier Output Fields**

| Field      | Values                |
|------------|-----------------------|
| Classifier | Name of a classifier. |

Table 551: Summary of Key CoS Classifier Output Fields (*continued*)

| Field            | Values                                                                                                                                                                                                                                                                    |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Code point type  | Type of classifier: <ul style="list-style-type: none"> <li>• <b>dscp</b>—All classifiers of the DSCP type.</li> <li>• <b>ieee-802.1</b>—All classifiers of the IEEE 802.1 type.</li> <li>• <b>ieee-mcast</b>—All classifiers of the IEEE 802.1 multicast type.</li> </ul> |
| Index            | Internal index of the classifier.                                                                                                                                                                                                                                         |
| Code point       | DSCP or IEEE 802.1 code point value of the incoming packets, in bits. These values are used for classification.                                                                                                                                                           |
| Forwarding Class | Name of the forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.                                                                       |
| Loss Priority    | Loss priority value that the classifier assigns to the incoming packet based on its code point value.                                                                                                                                                                     |

- Related Documentation**
- [Defining CoS Unicast BA Classifiers on page 5784](#)
  - [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 5785](#)

## Monitoring CoS Forwarding Classes

**Purpose** Use the monitoring functionality to view the current assignment of CoS forwarding classes to queue numbers on the system.

**Action** To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
user@switch> show class-of-service forwarding-class
```

**Meaning** [Table 552 on page 5917](#) summarizes key output fields for CoS forwarding classes.



Table 552: Summary of Key CoS Forwarding Class Output Fields

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class | <p>Names of forwarding classes assigned to queue numbers. By default, the following unicast forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li>• <b>fcoe</b>—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.</li> <li>• <b>no-loss</b>—Provides guaranteed delivery for TCP lossless traffic</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> <p>By default, the following multideestination forwarding class is assigned to queue 8:</p> <ul style="list-style-type: none"> <li>• <b>mcast</b>—Provides no special CoS handling of packets.</li> </ul> |
| Queue            | <p>Queue number corresponding to the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to unicast forwarding classes and one queue (8) is assigned to a multideestination forwarding class.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| No-Loss          | <p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> <li>• Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion)</li> <li>• Enabled—The forwarding class is configured for lossless transport</li> </ul> <p><b>NOTE:</b> To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p>                                                                                                                                                                                                                        |

- Related Documentation**
- [Defining CoS Forwarding Classes on page 5787](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)

## Monitoring Interfaces That Have CoS Components

**Purpose** Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

**Action** To monitor interfaces that have CoS components in the CLI, enter the command:

```
user@switch> show class-of-service interface
```

To monitor a specific interface in the CLI, enter the command:

```
user@switch> show class-of-service interface interface-name
```

**Meaning** [Table 553 on page 5918](#) summarizes key output fields for CoS interfaces.

Table 553: Summary of Key CoS Interfaces Output Fields

| Field                         | Values                                                                                                                                                                                                                                        |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface            | Name of a physical interface to which CoS components are assigned.                                                                                                                                                                            |
| Index                         | Index of this interface or the internal index of a specific object.                                                                                                                                                                           |
| Queues supported              | Number of queues you can configure on the interface.                                                                                                                                                                                          |
| Queues in use                 | Number of queues currently configured.                                                                                                                                                                                                        |
| Scheduler map                 | Name of the scheduler map associated with this interface.                                                                                                                                                                                     |
| Congestion-notification       | Status of congestion notification (enabled or disabled).                                                                                                                                                                                      |
| Rewrite Input IEEE Code-point | (Fibre Channel NP_Port interfaces only) IEEE 802.1p code point (priority) the interface assigns to incoming Fibre Channel (FC) traffic when the interface encapsulates the FC traffic in Ethernet before forwarding it onto the FCoE network. |
| Logical Interface             | Name of a logical interface on the physical interface to which CoS components are assigned.                                                                                                                                                   |
| Object                        | Category of an object—for example, <b>classifier</b> , <b>scheduler-map</b> , or <b>rewrite</b> .                                                                                                                                             |
| Name                          | Name of the object—for example, <b>ba-classifier</b> .                                                                                                                                                                                        |
| Type                          | Type of the object—for example, <b>ieee8021p</b> for a classifier.                                                                                                                                                                            |

**Related Documentation** • [Assigning CoS Components to Interfaces on page 5807](#)

## Monitoring CoS Rewrite Rules

**Purpose** Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

**Action** To monitor CoS rewrite rules in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule
```

To monitor a particular rewrite rule in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule name rewrite-rule-name
```

To monitor a particular type of rewrite rule (for example, DSCP, DSCP IPv6, or IEEE-802.1) in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule type rewrite-rule-type
```

**Meaning** [Table 554 on page 5919](#) summarizes key output fields for CoS rewrite rules.

**Table 554: Summary of Key CoS Rewrite Rule Output Fields**

| Field            | Values                                                                                                                                                                                                                                     |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rewrite rule     | Name of the rewrite rule.                                                                                                                                                                                                                  |
| Code point type  | Rewrite rule type: <ul style="list-style-type: none"> <li>• <b>dscp</b>—For IPv4 DiffServ traffic.</li> <li>• <b>dscp-ipv6</b>—For IPv6 Diffserv traffic.</li> <li>• <b>ieee-802.1</b>—For Layer 2 traffic.</li> </ul>                     |
| Index            | Internal index for the rewrite rule.                                                                                                                                                                                                       |
| Forwarding class | Name of the forwarding class that is used to determine CoS values for rewriting in combination with loss priority.<br><br>Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting. |
| Loss priority    | Level of loss priority that is used to determine CoS values for rewriting in combination with forwarding class.                                                                                                                            |
| Code point       | Rewrite code point value.                                                                                                                                                                                                                  |

**Related Documentation** • [Defining CoS Rewrite Rules on page 5805](#)

## Monitoring CoS Scheduler Maps

**Purpose** Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

**Action** To monitor CoS scheduler maps in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map
```

To monitor a specific scheduler map in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

**Meaning** [Table 555 on page 5919](#) summarizes key output fields for CoS scheduler maps.

**Table 555: Summary of Key CoS Scheduler Maps Output Fields**

| Field         | Values                                                                   |
|---------------|--------------------------------------------------------------------------|
| Scheduler map | Name of the scheduler map.                                               |
| Index         | Index of a specific object—scheduler maps, schedulers, or drop profiles. |

Table 555: Summary of Key CoS Scheduler Maps Output Fields (*continued*)

| Field            | Values                                                                                                                                                                                                                                                                                                                                                      |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scheduler        | Name of the scheduler.                                                                                                                                                                                                                                                                                                                                      |
| Forwarding class | Names of the forwarding classes to which the scheduler is assigned.                                                                                                                                                                                                                                                                                         |
| Transmit rate    | Configured transmit rate of the scheduler as a percentage of the total interface bandwidth.                                                                                                                                                                                                                                                                 |
| Priority         | <p>Scheduling priority of a queue:</p> <ul style="list-style-type: none"> <li>• <b>strict-high</b> or <b>high</b>—Packets in this queue are transmitted first. Only one queue can be configured as <b>strict-high</b> or <b>high</b>.</li> <li>• <b>low</b>—Packets in this queue are transmitted after packets in the <b>strict-high</b> queue.</li> </ul> |
| Drop Profiles    | Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.                                                                                                                                                                                                                                                            |
| Loss Priority    | Drop profile associated with each packet loss priority. You can configure different drop profiles for <b>low</b> , <b>medium-high</b> , and <b>high</b> loss priority traffic.                                                                                                                                                                              |
| Protocol         | Transport protocol of the drop profile for the particular priority.                                                                                                                                                                                                                                                                                         |
| Name             | Name of the drop profile.                                                                                                                                                                                                                                                                                                                                   |

**Related Documentation** • [Defining CoS Queue Schedulers on page 5789](#)

## Monitoring CoS Value Aliases

**Purpose** Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP and IEEE 802.1p code point bits.

**Action** To monitor CoS value aliases in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases
```

To monitor a specific type of code-point alias (for example, DSCP or IEEE 802.1) in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases ieee-802.1
```

**Meaning** [Table 556 on page 5921](#) summarizes key output fields for CoS value aliases.

Table 556: Summary of Key CoS Value Alias Output Fields

| Field           | Values                                                                                                                                                                                                                                       |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Code point type | Type of the CoS value: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Examines Layer 3 packet headers for IP packet classification.</li> <li>• <b>ieee-802.1</b>—Examines Layer 2 packet headers for packet classification.</li> </ul> |
| Alias           | Name given to a set of bits—for example, <b>af11</b> is a name for bits <b>001010</b> .                                                                                                                                                      |
| Bit pattern     | Set of bits associated with the alias.                                                                                                                                                                                                       |

**Related Documentation**

- [Defining CoS Code-Point Aliases on page 5783](#)

## Operational Commands

- [show class-of-service](#)
- [show class-of-service classifier](#)
- [show class-of-service code-point-aliases](#)
- [show class-of-service congestion-notification](#)
- [show class-of-service drop-profile](#)
- [show class-of-service forwarding-class](#)
- [show class-of-service forwarding-class-set](#)
- [show class-of-service forwarding-table](#)
- [show class-of-service forwarding-table classifier](#)
- [show class-of-service forwarding-table classifier mapping](#)
- [show class-of-service forwarding-table drop-profile](#)
- [show class-of-service forwarding-table rewrite-rule](#)
- [show class-of-service forwarding-table rewrite-rule mapping](#)
- [show class-of-service forwarding-table scheduler-map](#)
- [show class-of-service interface](#)
- [show class-of-service multi-destination](#)
- [show class-of-service rewrite-rule](#)
- [show class-of-service scheduler-map](#)
- [show class-of-service shared-buffer](#)
- [show class-of-service traffic-control-profile](#)
- [show dcbx](#)
- [show dcbx neighbors](#)
- [show interfaces queue](#)
- [show pfe next-hop](#)

- `show pfe route`
- `show pfe terse`
- `show pfe version`

## show class-of-service

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show class-of-service</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display the class-of-service (CoS) information.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring CoS Value Aliases on page 5920</a></li> <li>• <a href="#">Monitoring CoS Classifiers on page 5915</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 5916</a></li> <li>• <a href="#">Monitoring Interfaces That Have CoS Components on page 5917</a></li> <li>• <a href="#">Monitoring CoS Scheduler Maps on page 5919</a></li> <li>• <a href="#">Monitoring CoS Rewrite Rules on page 5918</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of- service on page 5924</a>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 557 on page 5923</a> lists the output fields for the <b>show class-of-service</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                         |

**Table 557: show class-of-service Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                             | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Forwarding class</b> | The forwarding class configuration: <ul style="list-style-type: none"> <li>• <b>Forwarding class</b>—Name of the forwarding class.</li> <li>• <b>ID</b>—Forwarding class ID.</li> <li>• <b>Queue</b>—Queue number.</li> </ul> | All levels      |
| <b>Code point type</b>  | The type of code-point alias: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Aliases for DiffServ code point (DSCP) values.</li> <li>• <b>ieee-802.1</b>—Aliases for IEEE 802.1p values.</li> </ul>                     | All levels      |
| <b>Alias</b>            | Names given to CoS values.                                                                                                                                                                                                    | All levels      |
| <b>Bit pattern</b>      | Set of bits associated with an alias.                                                                                                                                                                                         | All levels      |
| <b>Classifier</b>       | Name of the classifier.                                                                                                                                                                                                       | All levels      |
| <b>Code point</b>       | Code-point values.                                                                                                                                                                                                            | All levels      |
| <b>Loss priority</b>    | Loss priority assigned to specific CoS values and aliases of the classifier.                                                                                                                                                  | All levels      |
| <b>Rewrite rule</b>     | Name of the rewrite rule if one has been configured.                                                                                                                                                                          | All levels      |

Table 557: show class-of-service Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                     | Level of Output |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Drop profile</b>            | Name of the drop profile.                                                                                                                             | All levels      |
| <b>Type</b>                    | Type of drop profile. QFX Series supports only the <b>discrete</b> type of drop-profile.                                                              | All levels      |
| <b>Fill level</b>              | Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.                                    | All levels      |
| <b>Scheduler map</b>           | Name of the scheduler map.                                                                                                                            | All levels      |
| <b>Scheduler</b>               | Name of the scheduler.                                                                                                                                | All levels      |
| <b>Transmit rate</b>           | Transmission rate of the scheduler.                                                                                                                   | All levels      |
| <b>Buffer size</b>             | Delay buffer size in the queue.                                                                                                                       | All levels      |
| <b>Drop profiles</b>           | Drop profiles configured for the specified scheduler.                                                                                                 | All levels      |
| <b>Protocol</b>                | Transport protocol corresponding to the drop profile.                                                                                                 | All levels      |
| <b>Name</b>                    | Name of the drop profile.                                                                                                                             | All levels      |
| <b>Queues supported</b>        | Number of queues that can be configured on the interface.                                                                                             | All levels      |
| <b>Queues in use</b>           | Number of queues currently configured.                                                                                                                | All levels      |
| <b>Physical interface</b>      | Name of the physical interface.                                                                                                                       | All levels      |
| <b>Scheduler map</b>           | Name of the scheduler map.                                                                                                                            | All levels      |
| <b>Congestion-notification</b> | Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface. | All levels      |
| <b>Forwarding class set</b>    | Name of the forwarding class set (priority group).                                                                                                    |                 |
| <b>Index</b>                   | Internal index of an object.                                                                                                                          | All levels      |

## Sample Output

### show class-of- service

```

user@switch> show class-of-service
Forwarding class ID Queue
best-effort 0 0
fcoe 1 3
no-loss 2 4
network-control 3 7
mcast 8 8

Code point type: dscp

```



```

Alias Bit pattern
af11 001010
af12 001100
... ...

Code point type: ieee-802.1
Alias Bit pattern
af11 100
... ...

Classifier: dscp-default, Code point type: dscp, Index: 7
Code point Forwarding class Loss priority
000000 best-effort low
000001 best-effort low
...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point Forwarding class Loss priority
000 best-effort low
001 best-effort low
010 best-effort low
011 fcoe low
100 no-loss low
101 best-effort low
110 network-control low
111 network-control low

Drop profile:<default-drop-profile>, Type: discrete, Index: 1
Fill level
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 21
Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>

Scheduler: <default-fcoe>, Forwarding class: fcoe, Index: 50
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>

Scheduler: <default-noloss>, Forwarding class: no-loss, Index: 51
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low

```

```
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>

Scheduler: <default-nc>, Forwarding class: network-control, Index: 23
 Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
 Priority: low
 Excess Priority: low
 drop-profile-map-set-type: mark
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>

Scheduler: <default-mcast>, Forwarding class: mcast, Index: 49
 Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent, Buffer
Limit: none,
 Priority: low
 Excess Priority: low
 drop-profile-map-set-type: mark
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 <default-drop-profile>
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>

Physical interface: xe-0/0/0, Index: 129
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Physical interface: xe-0/0/1, Index: 130
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

...

Forwarding class set: lan-fcset, Type: normal-type, Forwarding class set index:
7
 Forwarding class Index
 best-effort 0
```

## show class-of-service classifier

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service classifier<br><name <i>name</i> ><br><type dscp   type dscp-ipv6   type exp   type ieee-802.1   type inet-precedence>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>none</b>—Display all classifiers.</p> <p><b>name <i>name</i></b>—(Optional) Display named classifier.</p> <p><b>type dscp</b>—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p><b>type dscp-ipv6</b>—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p><b>type exp</b>—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p><b>type ieee-802.1</b>—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p><b>type inet-precedence</b>—(Optional) Display all classifiers of the inet-precedence type.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service classifier type ieee-802.1 on page 5928</a><br><a href="#">show class-of-service classifier type ieee-802.1 (QFX Series) on page 5928</a>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 558 on page 5927</a> describes the output fields for the <b>show class-of-service classifier</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 558: show class-of-service classifier Output Fields**

| Field Name              | Field Description                                                                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Classifier</b>       | Name of the classifier.                                                                                                                                                 |
| <b>Code point type</b>  | Type of the classifier: <b>exp</b> (not on EX Series switch), <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>ieee-802.1</b> , or <b>inet-precedence</b> . |
| <b>Index</b>            | Internal index of the classifier.                                                                                                                                       |
| <b>Code point</b>       | Code point value used for classification                                                                                                                                |
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.                                        |

Table 558: show class-of-service classifier Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Loss priority</b> | Loss priority value used for classification. For most platforms, the value is <b>high</b> or <b>low</b> . For some platforms, the value is <b>high</b> , <b>medium-high</b> , <b>medium-low</b> , or <b>low</b> . |

## Sample Output

### show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point Forwarding Class Loss priority
000 best-effort low
001 best-effort high
010 expedited-forwarding low
011 expedited-forwarding high
100 assured-forwarding low
101 assured-forwarding medium-high
110 network-control low
111 network-control high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point Forwarding class Loss priority
100 expedited-forwarding low

```

### show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point Forwarding class Loss priority
000 best-effort low
001 best-effort low
010 best-effort low
011 fcoe low
100 no-loss low
101 best-effort low
110 network-control low
111 network-control low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point Forwarding class Loss priority
000 mcast low
001 mcast low
010 mcast low
011 mcast low
100 mcast low
101 mcast low
110 mcast low
111 mcast low

```

## show class-of-service code-point-aliases

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service code-point-aliases</code><br><code>&lt;dscp   dscp-ipv6   exp   ieee-802.1   inet-precedence&gt;</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display code point aliases of all code point types.</p> <p><b>dscp</b>—(Optional) Display Differentiated Services code point (DSCP) aliases.</p> <p><b>dscp-ipv6</b>—(Optional) Display IPv6 DSCP aliases.</p> <p><b>exp</b>—(Optional) Display MPLS EXP code point aliases.</p> <p><b>ieee-802.1</b>—(Optional) Display IEEE-802.1 code point aliases.</p> <p><b>inet-precedence</b>—(Optional) Display IPv4 precedence code point aliases.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service code-point-aliases exp on page 5930</a>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 559 on page 5929</a> describes the output fields for the <b>show class-of-service code-point-aliases</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                     |

**Table 559: show class-of-service code-point-aliases Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Code point type</b> | Type of the code points displayed: <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>exp</b> (not on EX Series switch or the QFX Series), <b>ieee-802.1</b> , or <b>inet-precedence</b> (not on the QFX Series). |
| <b>Alias</b>           | Alias for a bit pattern.                                                                                                                                                                                                    |
| <b>Bit pattern</b>     | Bit pattern for which the alias is displayed.                                                                                                                                                                               |

## Sample Output

`show class-of-service code-point-aliases exp`

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
Alias Bit pattern
af11 100
af12 101
be 000
be1 001
cs6 110
cs7 111
ef 010
ef1 011
nc1 110
nc2 111
```

## show class-of-service congestion-notification

|                                 |                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service congestion-notification                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                  |
| <b>Description</b>              | Display whether priority-based flow control (PFC) is enabled for each IEEE 802.1p code point.                                                                                                                    |
| <b>Options</b>                  | <b>none</b> —Display the PFC state for all IEEE 802.1p code points.                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 560 on page 5931</a> describes the output fields for the <b>show class-of-service congestion-notification</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 560: show class-of-service congestion-notification Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>                | Type of interfaces on which congestion notification is applied. Congestion notification is applied on input interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Index</b>               | Index of this congestion notification profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Name</b>                | Name of the congestion notification profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Cable Length</b>        | Length of the attached physical cable in meters. The default value is 100 meters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Priority</b>            | IEEE 802.1p code point.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>PFC</b>                 | State of PFC for the corresponding code point, either <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>MRU</b>                 | <p>Maximum receive unit of the interface in bytes. (Incoming traffic that exceeds the MRU size of an interface is dropped.) The default values are:</p> <ul style="list-style-type: none"> <li>2500 bytes for priority 3 traffic</li> <li>9216 bytes for priority 4 traffic</li> </ul> <p><b>NOTE:</b> If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.</p> |
| <b>Flow-Control-Queues</b> | Output queue mapping to IEEE 802.1p code points (priorities). Explicit output queue to priority mapping overwrites the default configuration, and only explicitly mapped queues are displayed in the output. Flow control is only enabled on a queue when you enable PFC on the corresponding priority in the input stanza of the congestion notification profile.                                                                                                                                                                                                                       |

## Sample Output

### show class-of-service congestion-notification

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Enabled  | 9216 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |
| 100      |                     |
|          | 4                   |
| 101      |                     |
|          | 5                   |
| 110      |                     |
|          | 6                   |
| 111      |                     |
|          | 7                   |

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2240 |
| 100      | Disabled |      |
| 101      | Enabled  | 2240 |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 011      |                     |
|          | 3                   |
| 101      |                     |
|          | 5                   |



## show class-of-service drop-profile

|                                 |                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service drop-profile<br><profile-name <i>profile-name</i> >                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Display data points for each class-of-service (CoS) random early detection (RED) drop profile.                                                                                                                                        |
| <b>Options</b>                  | <b>none</b> —Display all drop profiles.<br><br><b>profile-name <i>profile-name</i></b> —(Optional) Display the specified profile only.                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service drop-profile on page 5934</a><br><a href="#">show class-of-service drop-profile (EX4200 Switch) on page 5934</a><br><a href="#">show class-of-service drop-profile (EX8200 Switch) on page 5934</a> |
| <b>Output Fields</b>            | <a href="#">Table 561 on page 5933</a> describes the output fields for the <b>show class-of-service drop-profile</b> command. Output fields are listed in the approximate order in which they appear.                                 |

**Table 561: show class-of-service drop-profile Output Fields**

| Field Name              | Field Description                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profile</b>     | Name of a drop profile.                                                                                                                                   |
| <b>Type</b>             | Type of drop profile: <ul style="list-style-type: none"> <li>• <b>discrete</b> (default)</li> <li>• <b>interpolated</b> (EX8200 switches only)</li> </ul> |
| <b>Index</b>            | Internal index of this drop profile.                                                                                                                      |
| <b>Fill Level</b>       | Percentage fullness of a queue.                                                                                                                           |
| <b>Drop probability</b> | Drop probability at this fill level.                                                                                                                      |

## Sample Output

### show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
 Fill level Drop probability
 0 0
 1 1
 2 2
 4 4
 5 5
 6 6
 8 8
 10 10
 12 15
 14 20
 15 23
... 64 entries total
 90 96
 92 96
 94 97
 95 98
 96 98
 98 99
 99 99
 100 100
```

### show class-of-service drop-profile (EX4200 Switch)

```
user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level
 100
Drop profile: dp1, Type: discrete, Index: 40496
 Fill level
 10
```

### show class-of-service drop-profile (EX8200 Switch)

```
user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: dp1, Type: interpolated, Index: 40496
 Fill level Drop probability
 0 0
 1 80
 2 90
 4 90
 5 90
 6 90
 8 90
 10 90
 12 91
 14 91
 15 91
 16 91
```

|                                                 |                  |
|-------------------------------------------------|------------------|
| 18                                              | 91               |
| 20                                              | 91               |
| 22                                              | 92               |
| 24                                              | 92               |
| 25                                              | 92               |
| 26                                              | 92               |
| 28                                              | 92               |
| 30                                              | 92               |
| 32                                              | 93               |
| 34                                              | 93               |
| 35                                              | 93               |
| 36                                              | 93               |
| 38                                              | 93               |
| 40                                              | 93               |
| 42                                              | 94               |
| 44                                              | 94               |
| 45                                              | 94               |
| 46                                              | 94               |
| 48                                              | 94               |
| 49                                              | 94               |
| 51                                              | 95               |
| 52                                              | 95               |
| 54                                              | 95               |
| 55                                              | 95               |
| 56                                              | 95               |
| 58                                              | 95               |
| 60                                              | 95               |
| 62                                              | 96               |
| 64                                              | 96               |
| 65                                              | 96               |
| 66                                              | 96               |
| 68                                              | 96               |
| 70                                              | 96               |
| 72                                              | 97               |
| 74                                              | 97               |
| 75                                              | 97               |
| 76                                              | 97               |
| 78                                              | 97               |
| 80                                              | 97               |
| 82                                              | 98               |
| 84                                              | 98               |
| 85                                              | 98               |
| 86                                              | 98               |
| 88                                              | 98               |
| 90                                              | 98               |
| 92                                              | 99               |
| 94                                              | 99               |
| 95                                              | 99               |
| 96                                              | 99               |
| 98                                              | 99               |
| 99                                              | 99               |
| 100                                             | 100              |
| Drop profile: dp2, Type: discrete, Index: 40499 |                  |
| Fill level                                      | Drop probability |
| 10                                              | 5                |
| 50                                              | 50               |

## show class-of-service forwarding-class

|                                 |                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-class                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring CoS on EX Series Switches</i></li> <li>• <i>Monitoring CoS Forwarding Classes</i></li> <li>• <i>Defining CoS Forwarding Classes (CLI Procedure)</i></li> <li>• <i>Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-class on page 5937</a><br><a href="#">show class-of-service forwarding-class (EX8200 Switch) on page 5937</a><br><a href="#">show class-of-service forwarding-class (QFX Series) on page 5937</a>                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 562 on page 5936</a> describes the output fields for the <b>show class-of-service forwarding-class</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                 |

**Table 562: show class-of-service forwarding-class Output Fields**

| Field Name               | Field Description                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding class</b>  | Name of the forwarding class.                                                                                                                                      |
| <b>ID</b>                | Forwarding class identifier.                                                                                                                                       |
| <b>Queue</b>             | CoS queue mapped to the forwarding class.                                                                                                                          |
| <b>Policing priority</b> | Not supported on EX Series switches or the QFX Series and can be ignored.                                                                                          |
| <b>Fabric priority</b>   | (EX8200 switches only) Fabric priority for the forwarding class, either <b>high</b> or <b>low</b> . Determines the priority of packets entering the switch fabric. |

Table 562: show class-of-service forwarding-class Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>No-Loss</b> | <p>(QFX Series only) Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes:</p> <ul style="list-style-type: none"> <li>Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is <b>drop</b>).</li> <li>Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is <b>no-loss</b>).</li> </ul> |

## Sample Output

### show class-of-service forwarding-class

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Policing priority
best-effort 0 0 normal
expedited-forwarding 1 5 normal
assured-forwarding 2 1 normal
network-control 3 7 normal

```

## Sample Output

### show class-of-service forwarding-class (EX8200 Switch)

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Fabric priority
best-effort 0 0 low
expedited-forwarding 1 5 low
assured-forwarding 2 1 low
network-control 3 7 low
mcast-be 4 2 low
mcast-ef 5 4 low
mcast-af 6 6 low

```

## Sample Output

### show class-of-service forwarding-class (QFX Series)

```

user@switch> show class-of-service forwarding-class
Forwarding class ID Queue Policing priority No-Loss
best-effort 0 0 normal Disabled
fcoe 1 3 normal Enabled
no-loss 2 4 normal Enabled
network-control 3 7 normal Disabled
mcast 8 8 normal Disabled

```

## show class-of-service forwarding-class-set

|                                 |                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-class-set<br><forwarding-class-set-name>                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                               |
| <b>Description</b>              | Display the forwarding classes associated with each forwarding class set.                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display all forwarding class sets.</p> <p><b>forwarding-class-set-name</b>—(Optional) Display the forwarding classes associated with the specified forwarding class set.</p>   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 563 on page 5938 describes the output fields for the <b>show class-of-service forwarding-class-set</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 563: show class-of-service forwarding-class-set Output Fields**

| Field Name                 | Field Description                   |
|----------------------------|-------------------------------------|
| Forwarding class set       | Name of the forwarding class set.   |
| Type                       | Internal Junos OS type.             |
| Forwarding class set index | Index of this forwarding class set. |
| Forwarding class           | Name of a forwarding class.         |
| Index                      | Index of this forwarding class.     |

## Sample Output

### show class-of-service forwarding-class-set

```

user@switch> show class-of-service forwarding-class-set
Forwarding class set: san_fcset, Type: normal-type, Forwarding class set index:
37839
 Forwarding class Index
 fcoe 1

Forwarding class set: lan_fcset, Type: normal-type, Forwarding class set index:
37840
 Forwarding class Index
 best-effort 0

Forwarding class set: multicast_fcset, Type: normal-type, Forwarding class set
index: 37841

```

Forwarding class  
mcast

Index  
8

## show class-of-service forwarding-table

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 5940</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 5940</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Syntax                                       | show class-of-service forwarding-table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Syntax (TX Matrix and TX Matrix Plus Router) | show class-of-service forwarding-table<br><lcc number>   <sfc number>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Description                                  | Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all <b>show class-of-service forwarding-table</b> commands in succession.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Options                                      | <p><b>lcc number</b>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li><li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li><li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li></ul> <p><b>sfc number</b>—(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace <i>number</i> with 0.</p> |
| Required Privilege Level                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| List of Sample Output                        | <a href="#">show class-of-service forwarding-table on page 5941</a><br><a href="#">show class-of-service forwarding-table lcc (TX Matrix Plus Router) on page 5942</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Output Fields                                | See the output field descriptions for <b>show class-of-service forwarding-table</b> commands: <ul style="list-style-type: none"><li>• <a href="#">show class-of-service forwarding-table classifier</a></li><li>• <a href="#">show class-of-service forwarding-table classifier mapping</a></li><li>• <a href="#">show class-of-service forwarding-table drop-profile</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table loss-priority-map*
- *show class-of-service forwarding-table loss-priority-map mapping*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

## Sample Output

### show class-of-service forwarding-table

```

user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry # Code point Forwarding-class # PLP
0 000 0 0
1 001 0 1
2 010 1 0
3 011 1 1
4 100 2 0
5 101 2 1
6 110 3 0
7 111 3 1

Interface Index Table Index/ Q num Table type
sp-0/0/0.1001 66 11 11 IPv4 precedence
sp-0/0/0.2001 67 11 11 IPv4 precedence
sp-0/0/0.16383 68 11 11 IPv4 precedence
fe-0/0/0.0 69 11 11 IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
 Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
 Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
 Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
 Tx rate: 0 Kb (95%), Buffer size: 95 percent
 Priority low

```

```

 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
 Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1
 Drop
Entry Fullness(%) Probability(%)
 0 100 100

```

### show class-of-service forwarding-table lcc (TX Matrix Plus Router)

```

user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:

```

```

Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP
Entry # Code point Forwarding-class # PLP
 0 000000 0 0
 1 000001 0 0
 2 000010 0 0
 3 000011 0 0
 4 000100 0 0
 5 000101 0 0
 6 000110 0 0
 7 000111 0 0
 8 001000 0 0
 9 001001 0 0
 10 001010 0 0
 11 001011 0 0
 12 001100 0 0
 13 001101 0 0
 14 001110 0 0
 15 001111 0 0
 16 010000 0 0
 17 010001 0 0
 18 010010 0 0
 19 010011 0 0
 20 010100 0 0
 21 010101 0 0
 22 010110 0 0
 23 010111 0 0
 24 011000 0 0
 25 011001 0 0
 26 011010 0 0
 27 011011 0 0
 28 011100 0 0
 29 011101 0 0
 30 011110 0 0
 31 011111 0 0
 32 100000 0 0
 33 100001 0 0
 34 100010 0 0
 35 100011 0 0
 36 100100 0 0
 37 100101 0 0
 38 100110 0 0
 39 100111 0 0

```

|     |        |   |   |
|-----|--------|---|---|
| 40  | 101000 | 0 | 0 |
| 41  | 101001 | 0 | 0 |
| 42  | 101010 | 0 | 0 |
| 43  | 101011 | 0 | 0 |
| 44  | 101100 | 0 | 0 |
| 45  | 101101 | 0 | 0 |
| 46  | 101110 | 0 | 0 |
| ... |        |   |   |

## show class-of-service forwarding-table classifier

|                                 |                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table classifier                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                   |
| <b>Description</b>              | Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.                                                                                  |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table classifier on page 5944</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 564 on page 5944</a> describes the output fields for the <b>show class-of-service forwarding-table classifier</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 564: show class-of-service forwarding-table classifier Output Fields**

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Classifier table index</b> | Index of the classifier table.                                                                                                                                                                                                                                                                                                                                                   |
| <b>entries</b>                | Total number of entries.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Table type</b>             | Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> .                                                                                                                                                                                                    |
| <b>Entry #</b>                | Entry number.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Code point</b>             | Code point value used for classification.                                                                                                                                                                                                                                                                                                                                        |
| <b>Forwarding-class #</b>     | Forwarding class to which the code point is assigned.                                                                                                                                                                                                                                                                                                                            |
| <b>PLP</b>                    | Packet loss priority value set by classification. For most platforms, the value can be <b>0</b> or <b>1</b> . For some platforms, the value is <b>0</b> , <b>1</b> , <b>2</b> , or <b>3</b> . The value <b>0</b> represents low PLP. The value <b>1</b> represents <b>high</b> PLP. The value <b>2</b> represents medium-low PLP. The value <b>3</b> represents medium-high PLP. |

## Sample Output

### show class-of-service forwarding-table classifier

```
user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP
```

| Entry # | Code point | Forwarding-class # | PLP |
|---------|------------|--------------------|-----|
| 0       | 000000     | 0                  | 0   |
| 1       | 000001     | 0                  | 0   |
| 2       | 000010     | 0                  | 0   |
| 3       | 000011     | 0                  | 0   |
| 4       | 000100     | 0                  | 0   |
| 5       | 000101     | 0                  | 0   |
| 6       | 000110     | 0                  | 0   |
| 7       | 000111     | 0                  | 0   |
| 8       | 001000     | 0                  | 0   |
| 9       | 001001     | 0                  | 0   |
| 10      | 001010     | 1                  | 1   |
| 11      | 001011     | 0                  | 0   |
| ...     |            |                    |     |
| 60      | 111100     | 0                  | 0   |
| 61      | 111101     | 0                  | 0   |
| 62      | 111110     | 0                  | 0   |
| 63      | 111111     | 0                  | 0   |

## show class-of-service forwarding-table classifier mapping

|                                 |                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table classifier mapping                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                           |
| <b>Description</b>              | For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.                                      |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table classifier mapping on page 5946</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 565 on page 5946</a> describes the output fields for the <b>show class-of-service forwarding-table classifier mapping</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 565: show class-of-service forwarding-table classifier mapping Output Fields**

| Field Name         | Field Description                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Table index/ Q num | If the table type is <b>Fixed</b> , the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.                                |
| Interface          | Name of the logical interface. This field can also show the physical interface (QFX Series).                                                                                                 |
| Index              | Logical interface index.                                                                                                                                                                     |
| Table type         | Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>Fixed</b> , <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> . |

## Sample Output

### show class-of-service forwarding-table classifier mapping

```

user@host> show class-of-service forwarding-table classifier mapping
Table index/
Interface Index Q num Table type
so-5/0/0.0 10 62436 DSCP
so-0/1/0.0 11 62436 DSCP
so-0/2/0.0 12 1 Fixed
so-0/2/1.0 13 62436 DSCP
so-0/2/1.0 13 62437 IEEE 802.1
so-0/2/2.0 14 62436 DSCP
so-0/2/2.0 14 62438 IPv4 precedence

```



## show class-of-service forwarding-table drop-profile

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table drop-profile                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                     |
| <b>Description</b>              | Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.                                                                                                       |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table drop-profile on page 5948</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 566 on page 5948</a> describes the output fields for the <b>show class-of-service forwarding-table drop-profile</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 566: show class-of-service forwarding-table drop-profile Output Fields**

| Field Name             | Field Description                                         |
|------------------------|-----------------------------------------------------------|
| RED drop profile index | Index of this drop profile.                               |
| # entries              | Number of entries in a particular RED drop profile index. |
| Entry                  | Drop profile entry number.                                |
| Fullness(%)            | Percentage fullness of a queue.                           |
| Drop probability(%)    | Drop probability at this fill level.                      |

## Sample Output

### show class-of-service forwarding-table drop-profile

```

user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
 Drop
Entry Fullness(%) Probability(%)
 0 100 100

RED drop profile index: 8742, # entries: 3
 Drop
Entry Fullness(%) Probability(%)
 0 10 10
 1 20 20
 2 30 30

```



RED drop profile index: 24627, # entries: 64

| Entry | Fullness(%) | Drop           |  |
|-------|-------------|----------------|--|
|       |             | Probability(%) |  |
| 0     | 0           | 0              |  |
| 1     | 1           | 1              |  |
| 2     | 2           | 2              |  |
| 3     | 4           | 4              |  |
| ...   |             |                |  |
| 61    | 98          | 99             |  |
| 62    | 99          | 99             |  |
| 63    | 100         | 100            |  |

RED drop profile index: 25393, # entries: 64

| Entry | Fullness(%) | Drop           |  |
|-------|-------------|----------------|--|
|       |             | Probability(%) |  |
| 0     | 0           | 0              |  |
| 1     | 1           | 1              |  |
| 2     | 2           | 2              |  |
| 3     | 4           | 4              |  |
| ...   |             |                |  |
| 61    | 98          | 98             |  |
| 62    | 99          | 99             |  |
| 63    | 100         | 100            |  |

## show class-of-service forwarding-table rewrite-rule

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table rewrite-rule                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                     |
| <b>Description</b>              | Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.                                                                                      |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table rewrite-rule on page 5950</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 567 on page 5950</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 567: show class-of-service forwarding-table rewrite-rule Output Fields**

| Field Name          | Field Description                                                                                                                                                                                                                                                      |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rewrite table index | Index for this rewrite rule.                                                                                                                                                                                                                                           |
| # entries           | Number of entries in this rewrite rule.                                                                                                                                                                                                                                |
| Table type          | Type of table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>EXP-SWAP-PUSH-2</b> , (J Series routers only), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> . |
| Q#                  | Queue number to which this entry is assigned.                                                                                                                                                                                                                          |
| Low bits            | Code point value for low-priority loss profile.                                                                                                                                                                                                                        |
| State               | State of this code point: <b>enabled</b> , <b>rewritten</b> , or <b>disabled</b> .                                                                                                                                                                                     |
| High bits           | Code point value for high-priority loss profile.                                                                                                                                                                                                                       |

## Sample Output

### show class-of-service forwarding-table rewrite-rule

```

user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q# Low bits State High bits State
0 000111 Enabled 001010 Enabled
2 000000 Disabled 001100 Enabled

```

|   |        |         |        |         |
|---|--------|---------|--------|---------|
| 1 | 101110 | Enabled | 110111 | Enabled |
| 3 | 110000 | Enabled | 111000 | Enabled |

## show class-of-service forwarding-table rewrite-rule mapping

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table rewrite-rule mapping                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                             |
| <b>Description</b>              | For each logical interface, display the table identifier of the rewrite rule map for each code point type.                                                                                                                     |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table rewrite-rule mapping on page 5952</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 568 on page 5952</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule mapping</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 568: show class-of-service forwarding-table rewrite-rule mapping Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>   | Name of the logical interface. This field can also show the physical interface (QFX Series).                                                                                                                                                                                                                             |
| <b>Index</b>       | Logical interface index.                                                                                                                                                                                                                                                                                                 |
| <b>Table index</b> | Rewrite table index.                                                                                                                                                                                                                                                                                                     |
| <b>Type</b>        | Type of classifier: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>EXP-SWAP-PUSH-2</b> (not on the QFX Series), <b>Frame-Relay DE</b> (J Series routers only), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> . |

### Sample Output

#### show class-of-service forwarding-table rewrite-rule mapping

```

user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface Index Table index Type
so-5/0/0.0 10 3753 DSCP
so-0/1/0.0 11 3753 DSCP
so-0/2/0.0 12 3753 DSCP
so-0/2/1.0 13 3753 DSCP
so-0/2/2.0 14 3753 DSCP
so-0/2/3.0 15 3753 DSCP

```

## show class-of-service forwarding-table scheduler-map

|                                 |                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-table scheduler-map                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                      |
| <b>Description</b>              | For each physical interface, display the scheduler map information as it exists in the forwarding table.                                                                                                                |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-table scheduler-map on page 5954</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 569 on page 5953</a> describes the output fields for the <b>show class-of-service forwarding-table scheduler-map</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 569: show class-of-service forwarding-table scheduler-map Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface          | Name of the physical interface.                                                                                                                                                                                                                                                             |
| Index              | Physical interface index.                                                                                                                                                                                                                                                                   |
| Map index          | Scheduler map index.                                                                                                                                                                                                                                                                        |
| Num of queues      | Number of queues defined in this scheduler map.                                                                                                                                                                                                                                             |
| Entry              | Number of this entry in the scheduler map.                                                                                                                                                                                                                                                  |
| Scheduler index    | Scheduler policy index.                                                                                                                                                                                                                                                                     |
| Forwarding-class # | Forwarding class number to which this entry is applied.                                                                                                                                                                                                                                     |
| Tx rate            | Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.                                                      |
| Max buffer delay   | Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword <b>remainder</b> , which indicates that the buffer is sized according to what remains after other scheduler buffer allocations. |
| Priority           | <ul style="list-style-type: none"> <li><b>high</b>—Queue priority is high.</li> <li><b>low</b>—Queue priority is low.</li> </ul>                                                                                                                                                            |
| PLP high           | Drop profile index for a high packet loss priority profile.                                                                                                                                                                                                                                 |

Table 569: show class-of-service forwarding-table scheduler-map Output Fields (*continued*)

| Field Name      | Field Description                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------------|
| PLP low         | Drop profile index for a low packet loss priority profile.                                                  |
| PLP medium-high | Drop profile index for a medium-high packet loss priority profile.                                          |
| PLP medium-low  | Drop profile index for a medium-low packet loss priority profile.                                           |
| TCP PLP high    | Drop profile index for a high TCP packet loss priority profile.                                             |
| TCP PLP low     | Drop profile index for a low TCP packet loss priority profile.                                              |
| Policy is exact | If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled. |

## Sample Output

### show class-of-service forwarding-table scheduler-map

```

user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
 Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
 Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
 Priority low
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
 Policy is exact
 Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
 Priority high
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
 Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
 Priority high
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
 Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
 Traffic chunk: Max = 0 bytes, Min = 0 bytes
 Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
 Priority low
 PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

```

## show class-of-service interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service interface<br><detail   comprehensive> <interface-name>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Forwarding class map information added in Junos OS Release 9.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.</p> <p>Options <b>detail</b> and <b>comprehensive</b> introduced in Junos OS Release 11.4.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>comprehensive</b>—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.</p> <p><b>detail</b>—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.</p> <p>If the <b>interface</b> <i>interface-name</i> is a physical interface, the output includes:</p> <ul style="list-style-type: none"> <li>• Brief QoS information about the physical interface</li> <li>• Brief QoS information about the logical interface</li> <li>• CoS information about the physical interface</li> <li>• Brief information about filters or policers of the logical interface</li> <li>• Brief CoS information about the logical interface</li> </ul> <p>If the <b>interface</b> <i>interface-name</i> is a logical interface, the output includes:</p> <ul style="list-style-type: none"> <li>• Brief QoS information about the logical interface</li> <li>• Information about filters or policers for the logical interface</li> <li>• CoS information about the logical interface</li> </ul> <p><b>interface-name</b>—(Optional) Display class-of-service (CoS) associations for the specified interface.</p> <p><b>none</b>—Display CoS associations for all physical and logical interfaces.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service interface (Physical) on page 5966</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

[show class-of-service interface \(Logical\) on page 5967](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 5967](#)  
[show class-of-service interface \(PPPoE Interface\) on page 5967](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 5967](#)  
[show class-of-service interface detail on page 5968](#)  
[show class-of-service interface comprehensive on page 5968](#)  
[show class-of-service interface \(ACX Series Routers\) on page 5978](#)

**Output Fields** Table 570 on page 5956 describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 570: show class-of-service interface Output Fields**

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                             |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>                  | Name of a physical interface.                                                                                                                                                                                                                                                                 |
| <b>Index</b>                               | Index of this interface or the internal index of this object.                                                                                                                                                                                                                                 |
| <b>Dedicated Queues</b>                    | Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.                                                                                                                                                                        |
| <b>Queues supported</b>                    | Number of queues you can configure on the interface.                                                                                                                                                                                                                                          |
| <b>Queues in use</b>                       | Number of queues currently configured.                                                                                                                                                                                                                                                        |
| <b>Total non-default queues created</b>    | Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.                                                                                                                                                                   |
| <b>Rewrite Input IEEE Code-point</b>       | (QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value. |
| <b>Shaping rate</b>                        | Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.            |
| <b>Scheduler map</b>                       | Name of the output scheduler map associated with this interface.                                                                                                                                                                                                                              |
| <b>Scheduler map forwarding class sets</b> | (QFX Series only) Name of the fabric forwarding class set scheduler map associated with a QFabric system Interconnect device interface.                                                                                                                                                       |
| <b>Input shaping rate</b>                  | For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.                                                                                                                                                                                                              |
| <b>Input scheduler map</b>                 | For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.                                                                                                                                                                                                |
| <b>Chassis scheduler map</b>               | Name of the scheduler map associated with the packet forwarding component queues.                                                                                                                                                                                                             |
| <b>Rewrite</b>                             | Name and type of the rewrite rules associated with this interface.                                                                                                                                                                                                                            |
| <b>Classifier</b>                          | Name and type of classifiers associated with this interface.                                                                                                                                                                                                                                  |



Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                   |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding-class-map</b>    | Name of the forwarding map associated with this interface.                                                                                                                                                                                          |
| <b>Congestion-notification</b> | (QFX Series only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                |
| <b>Logical interface</b>       | Name of a logical interface.                                                                                                                                                                                                                        |
| <b>Object</b>                  | Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , or <b>Translation Table</b> (for IQE PICs only).                                                             |
| <b>Name</b>                    | Name of an object.                                                                                                                                                                                                                                  |
| <b>Type</b>                    | Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , or <b>inet-precedence</b> .                                                                                                                        |
| <b>Link-level type</b>         | Encapsulation on the physical interface.                                                                                                                                                                                                            |
| <b>MTU</b>                     | MTU size on the physical interface.                                                                                                                                                                                                                 |
| <b>Speed</b>                   | Speed at which the interface is running.                                                                                                                                                                                                            |
| <b>Loopback</b>                | Whether loopback is enabled and the type of loopback.                                                                                                                                                                                               |
| <b>Source filtering</b>        | Whether source filtering is enabled or disabled.                                                                                                                                                                                                    |
| <b>Flow control</b>            | Whether flow control is enabled or disabled.                                                                                                                                                                                                        |
| <b>Auto-negotiation</b>        | (Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.                                                                                                                                                                       |
| <b>Remote-fault</b>            | (Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Device flags</b>    | <p>The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Interface flags</b> | <p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul> |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flags</b>         | <p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>                                 |
| <b>Encapsulation</b> | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Admin</b>         | Administrative state of the interface ( <b>Up</b> or <b>Down</b> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Link</b>          | Status of physical link ( <b>Up</b> or <b>Down</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Proto</b>         | Protocol configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Input Filter</b>  | Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Filter</b> | Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Link flags</b>    | <p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul> |
| <b>Hold-times</b>    | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>CoS queues</b>    | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Statistics last cleared</b> | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Input errors</b>            | Input errors on the interface. The labels are explained in the following list: <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul> |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output errors</b>                        | <p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>HS link FIFO underflows</b>—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeds the MTU of the interface.</li> </ul> |
| <b>Egress queues</b>                        | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Queue counters</b>                       | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>SONET alarms</b><br><b>SONET defects</b> | <p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SONET PHY</b>                            | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SONET section</b> | <p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SONET line</b>    | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul> |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name                                                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SONET path</b>                                                       | <p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>ES-PFE</b>—Errored seconds (far-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> |
| <b>Received SONET overhead</b><br><br><b>Transmitted SONET overhead</b> | <p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> <li>• <b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li>• <b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li>• <b>K1</b> and <b>K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li>• <b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li>• <b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal.</li> <li>• <b>Z3</b> and <b>Z4</b>—Allocated for future use.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Received path trace</b><br><br><b>Transmitted path trace</b>         | <p>SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>HDLC configuration</b>                                               | <p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> <li>• <b>Policing bucket</b>—Configured state of the receiving policer.</li> <li>• <b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li>• <b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li>• <b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packet Forwarding Engine configuration</b> | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> <li>• <b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>CoS information</b>                        | Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> |
| <b>Forwarding classes</b>                     | Total number of forwarding classes supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Egress queues</b>                          | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Queue</b>                                  | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Forwarding classes</b>                     | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Queued Packets</b>                         | Number of packets queued to this queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Queued Bytes</b>                           | Number of bytes queued to this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Transmitted Packets</b>                    | Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Transmitted Bytes</b>                      | Number of bytes transmitted by this queue. The byte counts vary by PIC type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Tail-dropped packets</b>                   | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RED-dropped packets | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> |
| RED-dropped bytes   | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Transmit rate       | Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Rate Limit          | <p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Buffer size         | Delay buffer size in the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Priority            | Scheduling priority configured as <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Excess Priority     | Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 570: show class-of-service interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Drop profiles          | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Excess Priority        | Priority of the excess bandwidth traffic on a scheduler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Drop profiles          | <p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Adjustment information | <p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>anccp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>anccpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> </ul> |

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4

```

```

Total non-default queues created: 4
Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

| Object               | Name                 | Type | Index |
|----------------------|----------------------|------|-------|
| Scheduler-map        | <default>            |      | 27    |
| Rewrite              | exp-default          | exp  | 21    |
| Classifier           | exp-default          | exp  | 5     |
| Classifier           | ipprec-compatibility | ip   | 8     |
| Forwarding-class-map | exp-default          | exp  | 5     |

#### show class-of-service interface (Logical)

```

user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

| Object               | Name                 | Type | Index |
|----------------------|----------------------|------|-------|
| Scheduler-map        | <default>            |      | 27    |
| Rewrite              | exp-default          | exp  | 21    |
| Classifier           | exp-default          | exp  | 5     |
| Classifier           | ipprec-compatibility | ip   | 8     |
| Forwarding-class-map | exp-default          | exp  | 5     |

#### show class-of-service interface (Gigabit Ethernet)

```

user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

#### show class-of-service interface (PPPoE Interface)

```

user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85

```

| Object                  | Name                 | Type   | Index      |
|-------------------------|----------------------|--------|------------|
| Traffic-control-profile | tcp-pppoe.o.pp0.1    | Output | 2726446535 |
| Classifier              | ipprec-compatibility | ip     | 13         |

```

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

#### show class-of-service interface (T4000 Routers with Type 5 FPCs)

```

user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
Queues supported: 8, Queues in use: 4
Shaping rate: 5000000000 bps
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77

```

| Index | Object     | Name                 | Type |
|-------|------------|----------------------|------|
| 13    | Classifier | ipprec-compatibility | ip   |

**show class-of-service interface detail**

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
```

```
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
```

```
Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled
```

```
Logical interface ge-0/3/0.0
```

```
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
```

```
inet
```

```
mpls
```

| Interface  | Admin | Link | Proto | Input Filter | Output Filter |
|------------|-------|------|-------|--------------|---------------|
| ge-0/3/0.0 | up    | up   | inet  |              |               |

```
mpls
```

| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.0 | up    | up   | inet  |               |                |

```
mpls
```

```
Logical interface: ge-0/3/0.0, Index: 68
```

| Object     | Name                 | Type           | Index |
|------------|----------------------|----------------|-------|
| Rewrite    | exp-default          | exp (mpls-any) | 33    |
| Classifier | exp-default          | exp            | 10    |
| Classifier | ipprec-compatibility | ip             | 13    |

```
Logical interface ge-0/3/0.1
```

```
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2
```

```
inet
```

| Interface  | Admin | Link | Proto | Input Filter | Output Filter |
|------------|-------|------|-------|--------------|---------------|
| ge-0/3/0.1 | up    | up   | inet  |              |               |

```
Interface Admin Link Proto Input Policer
```

| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.1 | up    | up   | inet  |               |                |

```
ge-0/3/0.1 up up inet
```

```
Logical interface: ge-0/3/0.1, Index: 69
```

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

**show class-of-service interface comprehensive**

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
```

```
Interface index: 138, SNMP ifIndex: 601, Generation: 141
```

```
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled,
```

```
Auto-negotiation: Enabled, Remote fault: Online
```

```
Device flags : Present Running
```

```
Interface flags: SNMP-Traps Internal: 0x4000
```

```

CoS queues : 4 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
Last flapped : 2010-09-07 06:35:22 PDT (15:14:42 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 total statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters:

```

|       | Queued packets | Transmitted packets | Dropped packets |
|-------|----------------|---------------------|-----------------|
| 0 af3 | 0              | 0                   | 0               |
| 1 af2 | 0              | 0                   | 0               |
| 2 ef2 | 0              | 0                   | 0               |
| 3 ef1 | 0              | 0                   | 0               |

```

Egress queues: 4 supported, 5 in use
Queue counters:

```

|       | Queued packets | Transmitted packets | Dropped packets |
|-------|----------------|---------------------|-----------------|
| 0 af3 | 0              | 0                   | 0               |
| 1 af2 | 0              | 0                   | 0               |
| 2 ef2 | 0              | 0                   | 0               |
| 3 ef1 | 0              | 0                   | 0               |

```

Active alarms : None
Active defects : None
MAC statistics:

```

|                   | Receive | Transmit |
|-------------------|---------|----------|
| Total octets      | 0       | 0        |
| Total packets     | 0       | 0        |
| Unicast packets   | 0       | 0        |
| Broadcast packets | 0       | 0        |
| Multicast packets | 0       | 0        |

```

CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec high
 2 ef2 39 19500 0 120
none
 Direction : Input
 CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec low
 0 af3 30 3000 45 0
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
 Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps

```

```

RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available
Queue: 1, Forwarding classes: af2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available
Queue: 2, Forwarding classes: ef2
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available
Queue: 3, Forwarding classes: ef1
Queued:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps
Transmitted:
 Packets : 108546 0 pps
 Bytes : 12754752 376 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes : Not Available

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```



Scheduler map: interface-scheduler-map, Index: 58414

Scheduler: ef2, Forwarding class: ef2, Index: 39155

Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer Limit: none, Priority: high

Excess Priority: unspecified

Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Input shaping rate: 10000 bps

Input scheduler map: scheduler-map

Scheduler map: scheduler-map, Index: 15103

Scheduler: af3, Forwarding class: af3, Index: 35058

Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer Limit: none, Priority: low

Excess Priority: unspecified

Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 40582 | green                   |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 18928 | yellow                  |

Drop profile: green, Type: discrete, Index: 40582

| Fill level | Drop probability |
|------------|------------------|
| 50         | 0                |
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: yellow, Type: discrete, Index: 18928

| Fill level | Drop probability |
|------------|------------------|
| 50         | 0                |
| 100        | 100              |

Chassis scheduler map: < default-drop-profile>

Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25

Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low

Excess Priority: low

Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

| Loss priority | Protocol | Index | Name                    |
|---------------|----------|-------|-------------------------|
| Low           | any      | 1     | < default-drop-profile> |
| Medium low    | any      | 1     | < default-drop-profile> |
| Medium high   | any      | 1     | < default-drop-profile> |
| High          | any      | 1     | < default-drop-profile> |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

| Fill level | Drop probability |
|------------|------------------|
| 100        | 100              |

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

```

Fill level Drop probability
 100 100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:
 Loss priority Protocol Index Name
 Low any 1 < default-drop-profile>
 Medium low any 1 < default-drop-profile>
 Medium high any 1 < default-drop-profile>
 High any 1 < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
 Fill level Drop probability
 100 100
 Congestion-notification: Disabled
Forwarding class
priority Policing priority
af3 normal
af2 normal
ef2 normal
ef1 normal
af1 normal
ID Queue Restricted queue Fabric
0 0 0 low
1 1 1 low
2 2 2 high
3 3 3 high
4 4 0 low

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Protocol inet, MTU: 1500, Generation: 172, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Input Filters: filter-in-ge-0/3/0.0-i,
 Policer: Input: p1-ge-0/3/0.0-inet-i
 Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0

```

Flags: Is-Primary  
Output Filters: exp-filter,,,,,

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)  
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2  
Input packets : 0  
Output packets: 0

|            |       |      |       |                        |                |
|------------|-------|------|-------|------------------------|----------------|
| Interface  | Admin | Link | Proto | Input Filter           | Output Filter  |
| ge-0/3/0.0 | up    | up   | inet  | filter-in-ge-0/3/0.0-i |                |
|            |       |      | mpls  |                        | exp-filter     |
| Interface  | Admin | Link | Proto | Input Policer          | Output Policer |
| ge-0/3/0.0 | up    | up   | inet  | p1-ge-0/3/0.0-inet-i   |                |
|            |       |      | mpls  |                        |                |

Filter: filter-in-ge-0/3/0.0-i

Counters:

| Name                         | Bytes | Packets |
|------------------------------|-------|---------|
| count-filter-in-ge-0/3/0.0-i | 0     | 0       |

Filter: exp-filter

Counters:

| Name                  | Bytes | Packets |
|-----------------------|-------|---------|
| count-exp-seven-match | 0     | 0       |
| count-exp-zero-match  | 0     | 0       |

Policers:

| Name                 | Packets |
|----------------------|---------|
| p1-ge-0/3/0.0-inet-i | 0       |

Logical interface: ge-0/3/0.0, Index: 68

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | exp-default | exp (mpls-any) | 33    |

Rewrite rule: exp-default, Code point type: exp, Index: 33

| Forwarding class | Loss priority | Code point |
|------------------|---------------|------------|
| af3              | low           | 000        |
| af3              | high          | 001        |
| af2              | low           | 010        |
| af2              | high          | 011        |
| ef2              | low           | 100        |
| ef2              | high          | 101        |
| ef1              | low           | 110        |
| ef1              | high          | 111        |

| Object     | Name        | Type | Index |
|------------|-------------|------|-------|
| Classifier | exp-default | exp  | 10    |

Classifier: exp-default, Code point type: exp, Index: 10

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | af3              | low           |
| 001        | af3              | high          |
| 010        | af2              | low           |
| 011        | af2              | high          |
| 100        | ef2              | low           |
| 101        | ef2              | high          |
| 110        | ef1              | low           |
| 111        | ef1              | high          |

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | af3              | low           |
| 001        | af3              | high          |
| 010        | af3              | low           |
| 011        | af3              | high          |
| 100        | af3              | low           |
| 101        | af3              | high          |
| 110        | ef1              | low           |
| 111        | ef1              | high          |

| Forwarding class | ID | Queue | Restricted queue | Fabric |
|------------------|----|-------|------------------|--------|
| priority         |    |       |                  |        |
| af3              | 0  | 0     | 0                | low    |
| af2              | 1  | 1     | 1                | low    |
| ef2              | 2  | 2     | 2                | high   |
| ef1              | 3  | 3     | 3                | high   |
| af1              | 4  | 4     | 0                | low    |
| normal           |    |       |                  |        |

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Traffic statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Local statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Transit statistics:

|                 |   |       |
|-----------------|---|-------|
| Input bytes :   | 0 | 0 bps |
| Output bytes :  | 0 | 0 bps |
| Input packets:  | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

Protocol inet, MTU: 1500, Generation: 174, Route table: 0

Flags: Sendbcst-pkt-to-re

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Input packets : 0

Output packets: 0

| Interface  | Admin | Link | Proto | Input Filter  | Output Filter  |
|------------|-------|------|-------|---------------|----------------|
| ge-0/3/0.1 | up    | up   | mpls  |               |                |
| Interface  | Admin | Link | Proto | Input Policer | Output Policer |
| ge-0/3/0.1 | up    | up   | mpls  |               |                |

Logical interface: ge-0/3/0.1, Index: 69

| Object | Name | Type | Index |
|--------|------|------|-------|
|--------|------|------|-------|

| Classifier                                                                    |                   | ipprec-compatibility |       | ip               |        | 13 |
|-------------------------------------------------------------------------------|-------------------|----------------------|-------|------------------|--------|----|
| Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13 |                   |                      |       |                  |        |    |
| Code point                                                                    |                   | Forwarding class     |       | Loss priority    |        |    |
| 000                                                                           |                   | af3                  |       | low              |        |    |
| 001                                                                           |                   | af3                  |       | high             |        |    |
| 010                                                                           |                   | af3                  |       | low              |        |    |
| 011                                                                           |                   | af3                  |       | high             |        |    |
| 100                                                                           |                   | af3                  |       | low              |        |    |
| 101                                                                           |                   | af3                  |       | high             |        |    |
| 110                                                                           |                   | ef1                  |       | low              |        |    |
| 111                                                                           |                   | ef1                  |       | high             |        |    |
| Forwarding class                                                              |                   | ID                   | Queue | Restricted queue | Fabric |    |
| priority                                                                      | Policing priority |                      |       |                  |        |    |
| af3                                                                           |                   | 0                    | 0     | 0                | low    |    |
|                                                                               | normal            |                      |       |                  |        |    |
| af2                                                                           |                   | 1                    | 1     | 1                | low    |    |
|                                                                               | normal            |                      |       |                  |        |    |
| ef2                                                                           |                   | 2                    | 2     | 2                | high   |    |
|                                                                               | normal            |                      |       |                  |        |    |
| ef1                                                                           |                   | 3                    | 3     | 3                | high   |    |
|                                                                               | normal            |                      |       |                  |        |    |
| af1                                                                           |                   | 4                    | 4     | 0                | low    |    |
|                                                                               | normal            |                      |       |                  |        |    |

### show class-of-service interface (ACX Series Routers)

```
user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
```

```
Logical interface: at-0/0/0.0, Index: 69
```

```
Logical interface: at-0/0/0.32767, Index: 70
```

```
Physical interface: at-0/0/1, Index: 133
```

```
Queues supported: 4, Queues in use: 4
```

```
Scheduler map: <default>, Index: 2
```

```
Congestion-notification: Disabled
```

```
Logical interface: at-0/0/1.0, Index: 71
```

```
Logical interface: at-0/0/1.32767, Index: 72
```

```
Physical interface: ge-0/1/0, Index: 146
```

```
Queues supported: 8, Queues in use: 5
```

```
Scheduler map: <default>, Index: 2
```

```
Congestion-notification: Disabled
```

| Object     | Name         | Type      | Index |
|------------|--------------|-----------|-------|
| Rewrite    | dscp-default | dscp      | 31    |
| Classifier | d1           | dscp      | 11331 |
| Classifier | ci           | ieee8021p | 583   |

```
Logical interface: ge-0/1/0.0, Index: 73
```

| Object  | Name       | Type           | Index |
|---------|------------|----------------|-------|
| Rewrite | custom-exp | exp (mpls-any) | 46413 |

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name | Type              | Index |
|------------|------|-------------------|-------|
| Rewrite    | ri   | ieee8021p (outer) | 35392 |
| Classifier | ci   | ieee8021p         | 583   |

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/1/3.0, Index: 77

| Object  | Name        | Type           | Index |
|---------|-------------|----------------|-------|
| Rewrite | custom-exp2 | exp (mpls-any) | 53581 |

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Physical interface: ge-0/1/5, Index: 151

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Physical interface: ge-0/1/6, Index: 152

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Physical interface: ge-0/1/7, Index: 153

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name | Type | Index |
|------------|------|------|-------|
| Classifier | d1   | dscp | 11331 |

Physical interface: ge-0/2/0, Index: 154

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Physical interface: ge-0/2/1, Index: 155

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: ge-0/2/1.0, Index: 78

Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

| Object     | Name                 | Type | Index |
|------------|----------------------|------|-------|
| Classifier | ipprec-compatibility | ip   | 13    |

Logical interface: xe-0/3/1.0, Index: 81

[edit]

user@host-g11#



## show class-of-service multi-destination

|                                 |                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service multi-destination                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                         |
| <b>Description</b>              | For each class-of-service (CoS) multideestination classifier, display the classifier type.                                                                                                                 |
| <b>Options</b>                  | <b>none</b> —Display all multideestination classifiers.                                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 571 on page 5981</a> describes the output fields for the <b>show class-of-service multi-destination</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 571: show class-of-service multi-destination Output Fields**

| Field Name              | Field Description                                          |
|-------------------------|------------------------------------------------------------|
| <b>Family ethernet</b>  | Family to which the classifier belongs.                    |
| <b>Classifier Name</b>  | Name of the classifier.                                    |
| <b>Classifier Type</b>  | Type of the classifier: <b>dscp</b> or <b>ieee-802.1</b> . |
| <b>Classifier Index</b> | Internal index of the classifier.                          |

## Sample Output

### show class-of-service multi-destination

```
user@switch> show class-of-service multi-destination
```

```
Family ethernet:
Classifier Name Classifier Type Classifier Index
ba-mcast-classifier ieee-802.1 62376
```

## show class-of-service rewrite-rule

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service rewrite-rule<br><name <i>name</i> ><br><type <i>type</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display the mapping of forwarding classes and loss priority to code point values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Display all rewrite rules.</p> <p><b>name <i>name</i></b>—(Optional) Display the specified rewrite rule.</p> <p><b>type <i>type</i></b>—(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:</p> <ul style="list-style-type: none"> <li><b>dscp</b>—For IPv4 traffic.</li> <li><b>dscp-ipv6</b>—For IPv6 traffic.</li> <li><b>exp</b>—For MPLS traffic.</li> <li><b>frame-relay-de</b>—(J Series routers only) For Frame Relay traffic.</li> <li><b>ieee-802.1</b>—For Layer 2 traffic.</li> <li><b>inet-precedence</b>—For IPv4 traffic.</li> </ul> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service rewrite-rule type dscp on page 5983</a><br><a href="#">show class-of-service rewrite-rule type dscp (QFX Series) on page 5983</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 572 on page 5982</a> describes the output fields for the <b>show class-of-service rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 572: show class-of-service rewrite-rule Output Fields**

| Field Name              | Field Description                                                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rewrite rule</b>     | Name of the rewrite rule.                                                                                                                  |
| <b>Code point type</b>  | Type of rewrite rule: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>frame-relay-de</b> , or <b>inet-precedence</b> .                    |
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch. |
| <b>Index</b>            | Internal index for this particular rewrite rule.                                                                                           |
| <b>Loss priority</b>    | Loss priority for rewriting.                                                                                                               |

Table 572: show class-of-service rewrite-rule Output Fields (*continued*)

| Field Name | Field Description            |
|------------|------------------------------|
| Code point | Code point value to rewrite. |

## Sample Output

### show class-of-service rewrite-rule type dscp

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
 Forwarding class Loss priority Code point
 gold high 000000
 silver low 110000
 silver high 111000
 bronze low 001010
 bronze high 001100
 lead high 101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
 Forwarding class Loss priority Code point
 gold low 000111
 gold high 001010
 silver low 110000
 silver high 111000
 bronze high 001100
 lead low 101110
 lead high 110111

```

## Sample Output

### show class-of-service rewrite-rule type dscp (QFX Series)

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp, Index: 31
 Forwarding class Loss priority Code point
 best-effort low 000000
 best-effort high 000000
 fcoe low 101110
 fcoe high 101110
 no-loss low 001010
 no-loss high 001100
 network-control low 110000
 network-control high 111000

```

## show class-of-service scheduler-map

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service scheduler-map</code><br><code>&lt;name&gt;</code>                                                                                                                          |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                     |
| <b>Description</b>              | Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.                                                                                          |
| <b>Options</b>                  | <b>none</b> —Display all scheduler maps.<br><br><b>name</b> —(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service scheduler-map on page 5985</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 573 on page 5984</a> describes the output fields for the <b>show class-of-service scheduler-map</b> command. Output fields are listed in the approximate order in which they appear. |

Table 573: show class-of-service scheduler-map Output Fields

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                          |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Scheduler map</b>        | Name of the scheduler map.                                                                                                                                                                                                                                                                                 |
| <b>Index</b>                | Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.                                                                                                                                                                                |
| <b>Scheduler</b>            | Name of the scheduler.                                                                                                                                                                                                                                                                                     |
| <b>Forwarding class</b>     | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.                                                                                                                                                                           |
| <b>Transmit rate</b>        | Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.                                                                     |
| <b>Rate Limit</b>           | Rate limiting configuration of the queue. Possible values are <b>none</b> , meaning no rate limiting, and <b>exact</b> , meaning the queue only transmits at the configured rate.                                                                                                                          |
| <b>Maximum buffer delay</b> | Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword <b>remainder</b> to indicate that the buffer is sized according to what remains after other scheduler buffer allocations. |
| <b>Priority</b>             | Scheduling priority: <b>low</b> or <b>high</b> .                                                                                                                                                                                                                                                           |

Table 573: show class-of-service scheduler-map Output Fields (*continued*)

| Field Name      | Field Description                                                                                                  |
|-----------------|--------------------------------------------------------------------------------------------------------------------|
| Excess priority | Priority of excess bandwidth: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> . |
| Adjust minimum  | Minimum shaping rate for an adjusted queue, in bps.                                                                |
| Adjust percent  | Bandwidth adjustment applied to a queue, in percent.                                                               |
| Drop profiles   | Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.     |
| Loss priority   | Packet loss priority for drop profile assignment.                                                                  |
| Protocol        | Transport protocol for drop profile assignment.                                                                    |
| Name            | Name of the drop profile.                                                                                          |

## Sample Output

### show class-of-service scheduler-map

```

user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
Drop profiles:
 Loss priority Protocol Index Name
 Low non-TCP 8724 aa-drop-profile
 Low TCP 9874 bb-drop-profile
 High non-TCP 8833 cc-drop-profile
 High TCP 8484 dd-drop-profile

Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
Priority: high
Drop profiles:
 Loss priority Protocol Index Name
 Low non-TCP 8724 aa-drop-profile
 Low TCP 9874 bb-drop-profile
 High non-TCP 8833 cc-drop-profile
 High TCP 8484 dd-drop-profile

```

## show class-of-service shared-buffer

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service shared-buffer<br><egress   ingress>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display the shared buffer allocation and partitioning configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display ingress and egress shared buffer settings.</p> <p><b>egress</b>—(Optional) Display the egress shared buffer settings.</p> <p><b>ingress</b>—(Optional) Display the ingress shared buffer settings.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 5748</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 5759</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 5764</a></li> <li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 5803</a></li> <li>• <a href="#">Understanding CoS Buffer Configuration on page 5527</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service shared-buffer on page 5987</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | Table 574 on page 5986 describes the output fields for the <b>show class-of-service shared-buffer</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 574: show class-of-service shared-buffer Output Fields**

| Field Name       | Field Description                                                                                                   |
|------------------|---------------------------------------------------------------------------------------------------------------------|
| Ingress          | Ingress shared buffer configuration.                                                                                |
| Total Buffer     | Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool. |
| Dedicated Buffer | Buffer space allocated to the dedicated buffer pool in KB.                                                          |
| Shared Buffer    | Buffer space allocated to the shared buffer pool in KB.                                                             |
| Lossless         | Buffer space allocated to the lossless traffic buffer pool in KB.                                                   |

Table 574: show class-of-service shared-buffer Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                           |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Lossless Headroom</b>             | Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.)                                            |
| <b>Lossy</b>                         | Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.                                                                                                                                |
| <b>Lossless Headroom Utilization</b> | Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.) |
| <b>Node Device</b>                   | Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).                                                                                                    |
| <b>Total</b>                         | Size of the lossless headroom ingress buffer pool in KB.                                                                                                                                                    |
| <b>Used</b>                          | Amount in KB of lossless headroom ingress buffer used.                                                                                                                                                      |
| <b>Free</b>                          | Amount in KB of lossless headroom ingress buffer free (unused).                                                                                                                                             |
| <b>Egress</b>                        | Egress shared buffer configuration.                                                                                                                                                                         |
| <b>Multicast</b>                     | Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)                                                                                                                     |

## Sample Output

### show class-of-service shared-buffer

```
user@switch> show class-of-service shared-buffer
```

```
Ingress:
```

```
Total Buffer : 9360.00 KB
Dedicated Buffer : 2158.00 KB
Shared Buffer : 7202.00 KB
 Lossless : 648.18 KB
 Lossless Headroom : 3240.90 KB
 Lossy : 3312.92 KB
```

```
Lossless Headroom Utilization:
```

```
Node Device Total Used Free
0 3240.90 KB 0.00 KB 3240.90 KB
```

```
Egress:
```

```
Total Buffer : 9360.00 KB
Dedicated Buffer : 2704.00 KB
Shared Buffer : 6656.00 KB
 Lossless : 3328.00 KB
 Multicast : 1264.64 KB
 Lossy : 2063.36 KB
```

## show class-of-service traffic-control-profile

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service traffic-control-profile</code><br><code>&lt;profile-name&gt;</code>                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 12.2 for ACX Series Routers.                                                                                                                                                               |
| <b>Description</b>              | For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and Trio MPC/MIC interfaces only, display traffic shaping and scheduling profiles.<br><br>(ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.                                                                                          |
| <b>Options</b>                  | <b>none</b> —Display all profiles.<br><br><b>profile-name</b> —(Optional) Display information about a single profile.                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service traffic-control-profile on page 5990</a><br><a href="#">show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 5990</a><br><a href="#">show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 5990</a> |
| <b>Output Fields</b>            | <a href="#">Table 575 on page 5988</a> describes the output fields for the <b>show class-of-service traffic-control-profile</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                        |

**Table 575: show class-of-service traffic-control-profile Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                            |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Traffic control profile</b> | Name of the traffic control profile.                                                                                                                                                                                                                                                                         |
| <b>Index</b>                   | Index number of the traffic control profile.                                                                                                                                                                                                                                                                 |
| <b>ATM Service</b>             | (MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: <ul style="list-style-type: none"> <li>cbr—Constant bit rate.</li> <li>rtvbr—Real time variable bit rate.</li> <li>nrtvbr—Non real time variable bit rate.</li> <li>ubr—Unspecified bit rate.</li> </ul> |
| <b>Maximum Burst Size</b>      | Configured maximum burst size, in cells.                                                                                                                                                                                                                                                                     |
| <b>Peak rate</b>               | Configured peak rate, in cps.                                                                                                                                                                                                                                                                                |



**Table 575: show class-of-service traffic-control-profile Output Fields (*continued*)**

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sustained rate</b>               | Configured sustained rate, in cps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Shaping rate</b>                 | Configured shaping rate, in bps.<br><br><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Shaping rate burst</b>           | Configured burst size for the shaping rate, in bytes.<br><br><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.                                                                                                                                                                                                                                                                                                                                      |
| <b>Shaping rate priority high</b>   | Configured shaping rate for high-priority traffic, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Shaping rate priority medium</b> | Configured shaping rate for medium-priority traffic, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Shaping rate priority low</b>    | Configured shaping rate for low-priority traffic, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Shaping rate excess high</b>     | Configured shaping rate for high-priority excess traffic, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Shaping rate excess low</b>      | Configured shaping rate for low-priority excess traffic, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Scheduler map</b>                | Name of the associated scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Delay Buffer rate</b>            | Configured delay buffer rate, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Excess rate</b>                  | Configured excess rate, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Excess rate high</b>             | Configured excess rate for high priority traffic, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Excess rate low</b>              | Configured excess rate for low priority traffic, in percent or proportion.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Guaranteed rate</b>              | Configured guaranteed rate, in bps or cps.<br><br><b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: <ul style="list-style-type: none"> <li>• <b>cbr</b>—Guaranteed rate is equal to the configured peak rate in cps.</li> <li>• <b>rtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> <li>• <b>nrtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> </ul> |
| <b>Guaranteed rate burst</b>        | Configured burst size for the guaranteed rate, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>adjust-minimum</b>               | Configured minimum shaping rate for an adjusted queue, in bps.                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 575: show class-of-service traffic-control-profile Output Fields (*continued*)

| Field Name               | Field Description                                                |
|--------------------------|------------------------------------------------------------------|
| overhead accounting mode | Configured shaping mode: <b>Frame Mode</b> or <b>Cell Mode</b> . |
| Overhead bytes           | Configured byte adjustment value.                                |

## Sample Output

### show class-of-service traffic-control-profile

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
 Scheduler map: m1
 Delay Buffer rate: 500000
 Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
 Scheduler map: m2
 Delay Buffer rate: 600000
 Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
 Scheduler map: m3
 Delay Buffer rate: 800000
 Guaranteed rate: 3000000
 .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
 Scheduler map: m4
 Delay Buffer rate: 750000
 Guaranteed rate: 4000000
 ..adjust-minimum 20000000

```

### show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
 ATM Service: RTVBR
 Scheduler map: m3
 overhead accounting mode: Frame Mode
 Shaping rate: 1000 cps
 Shaping rate burst: 500 cells
 Delay Buffer rate: 2000 cps
 Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
 ATM Service: UBR
 Scheduler map: m3
 overhead accounting mode: Frame Mode

```

### show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
 ATM Service: RTVBR
 Shaping rate: 2000 cps

```

```
Shaping rate burst: 200 cells
Scheduler map: <default>
Delay Buffer rate: 1000 cps
Guaranteed rate: 1700 cps
```

## show dcbx

|                                 |                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show dcbx                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                               |
| <b>Description</b>              | List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5300</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 416 on page 5299</a> lists the output fields for the <b>show dcbx</b> command. Output fields are listed in the approximate order in which they appear.      |

Table 576: show dcbx output fields

| Field Name | Field Description                                                                                                                                                                                                                                                     |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX       | Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> <li>• Enabled—DCBX is enabled on the switch or on the specified interface</li> <li>• Disabled—DCBX is disabled on the switch or on the specified interface</li> </ul> |
| Interface  | Name of the interface                                                                                                                                                                                                                                                 |

## Sample Output

### show dcbx

```

user@switch> show dcbx
DCBX : Enabled
Interface DCBX
xe-0/0/9.0 enabled
xe-0/0/32.0 enabled
xe-0/0/36.0 enabled

```

## show dcbx neighbors

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dcbx neighbors</b><br><b>&lt;interface interface-name&gt;</b><br><b>&lt;terse&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display information about all DCBX neighbor interfaces.<br><br><b>interface-name</b> —(Optional) Display information for the specified interface.<br><br><b>terse</b> —Display the specified level of output.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5204</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5144</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 4965</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">dcbx on page 5221</a></li> </ul>                   |
| <b>List of Sample Output</b>    | <a href="#">show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 6006</a><br><a href="#">show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 6008</a><br><a href="#">show dcbx neighbors terse (QFX Series) on page 6010</a><br><a href="#">show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 6010</a><br><a href="#">show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 6011</a><br><a href="#">show dcbx neighbors (EX4500 Switch: Includes ETS) on page 6012</a> |
| <b>Output Fields</b>            | <a href="#">Table 417 on page 5300</a> lists the output fields for the <b>show dcbx neighbors</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 577: show dcbx neighbors Output Fields

| Field Name | Field Description      |
|------------|------------------------|
| Interface  | Name of the interface. |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Parent Interface       | Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Active-application-map | Name of the application map applied to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Protocol-Mode          | <p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> <li>IEEE DCBX Version—The interface uses IEEE DCBX mode.</li> <li>DCBX Version 1.01—The interface uses DCBX version 1.01.</li> </ul> <p><b>NOTE:</b> On interfaces that use the IEEE DCBX mode, the <b>show dcbx neighbors interface <i>interface-name</i></b> operational command does not include application, PFC, or ETS operational state in the output.</p>                                        |
| Protocol-State         | <p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li><b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> <li><b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> </ul> |
| Local-Advertisement    | <p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| Operational version    | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| sequence-number        | <p>Number of state change messages sent to the peer.</p> <p>If the interface <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p> <p>If the interface <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p>                                                                                   |
| acknowledge-id         | <p>Number of acknowledge messages received from the peer.</p> <p>If the <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p> <p>If the <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p>                                                                                                  |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer-Advertisement</b>  | (DCBX Version 1.01 only)<br><br>Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Operational version</b> | Version of the DCBX standard used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>sequence-number</b>     | <p>Number of state change messages the peer sent to the local interface.</p> <p>If this number matches the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.</p> <p>If this number does not match the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.</p>            |
| <b>acknowledge-id</b>      | <p>Number of acknowledge messages the peer has received from the local interface.</p> <p>If this number matches the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.</p> <p>If this number does not match the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.</p> |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: PFC</b>               | Priority-based flow control (PFC) feature DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol-State</b>             | (DCBX Version 1.01 only)<br><br>DCBX protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—PFC autonegotiation is disabled.</li> </ul> |
| <b>Operational State</b>          | (DCBX Version 1.01 only)<br><br>Operational state of the feature: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Local-Advertisement</b>        | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                     | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                    | Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the PFC configuration from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the PFC configuration from the peer.</li> </ul>                                                                                                                                                                                                                                        |
| <b>Mac auth Bypass Capability</b> | (IEEE DCBX only)<br><br>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is <b>no</b> .                                                                                                                                                                                                                                                                                                                               |
| <b>Error</b>                      | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                       |



Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                        |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>                                                                                                                                                    |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                    |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                     |
| <b>Operational Mode</b>                               | <p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> <li>• <b>Enable</b>—PFC is enabled on the code point.</li> <li>• <b>Disable</b>—PFC is disabled on the code point.</li> </ul>                                                                                                                       |
| <b>Peer-Advertisement</b>                             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                          |
| <b>Enable</b>                                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                   |
| <b>Willing</b>                                        | <p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the PFC configuration from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the PFC configuration from the local interface.</li> </ul> |
| <b>Error</b>                                          | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>                                                                                                                                                                                                                                                                                            |
| <b>Mac auth Bypass Capability</b>                     | <p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The connected peer supports MAC authentication bypass.</li> <li>• <b>No</b>—The connected peer does not support MAC authentication bypass.</li> </ul> |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>                                                                                                                                                                                                                                                                                                                   |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>                                                                                                                                                                                                                                                                    |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: Application</b> | State information for the DCBX application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Protocol-State</b>       | <p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—The local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.</li> </ul> |
| <b>Local-Advertisement</b>  | <p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Enable</b>               | <p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Willing</b>              | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the FCoE interface state from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the FCoE interface state from the peer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Error</b>                | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. The local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. The local and peer configuration are not compatible.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Appl-Name</b>            | Name of the application:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet-Type</b>                  | <p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                            |
| <b>Socket-Number</b>                  | <p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Priority-Field or Priority-Map</b> | <p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>                                                                                                                                                                                                                                                            |
| <b>Status</b>                         | <p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> <p><b>NOTE:</b> If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p> |
| <b>Peer-Advertisement</b>             | <p>Status of advertisements that the peer sends to the local interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Enable</b>                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                          |
| <b>Willing</b>                        | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the FCoE interface state from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the FCoE interface state from the local interface.</li> </ul>                                                                                                                                                                               |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Error</b>                          | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>                                                                           |
| <b>Appl-Name</b>                      | Name of the application: <ul style="list-style-type: none"> <li>• <b>FCoE</b>—Fibre Channel over Ethernet</li> </ul>                                                                                                                                                                                                                                                     |
| <b>Ethernet-Type</b>                  | Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                |
| <b>Socket-Number</b>                  | Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.                                                                                                                                                         |
| <b>Priority-Field or Priority-Map</b> | Priority assigned to the application.                                                                                                                                                                                                                                                                                                                                    |
| <b>Status</b>                         | (DCBX Version 1.01 only)<br><br>Peer interface status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Feature: ETS</b>        | Enhanced Transmission Selection (ETS) DCBX state information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Protocol-State</b>      | (DCBX Version 1.01 only)<br><br>ETS protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> </ul>                                                          |
| <b>Operational State</b>   | (DCBX Version 1.01 only)<br><br>Operational state of the feature, <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Local-Advertisement</b> | Status of advertisements that the local interface sends to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Enable</b>              | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                                  |
| <b>TLV Type</b>            | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Recommendation-or-Configuration</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>             | Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise <b>No</b> for this field): <ul style="list-style-type: none"> <li>• <b>Yes</b>—Local interface is willing to learn the ETS state from the peer.</li> <li>• <b>No</b>—Local interface is not willing to learn the ETS state from the peer.</li> </ul>                                                                                                                                                                                                 |
| <b>Credit Based Shaper</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                  |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error. This should always be the switch ETS error state.</li> <li>• <b>Yes</b>—Error detected.</li> </ul>                                                                                     |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                         |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                          |
| <b>Code Point</b>                                     | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                      |
| <b>Priority-Group</b>                                 | Class-of-service (CoS) priority group (forwarding class set) identification number.                                                                                                                                                                                                                             |
| <b>Percentage B/W</b>                                 | Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.) |
| <b>Transmission Selection Algorithm</b>               | (IEEE DCBX only)<br><br>The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                             |
| <b>Peer-Advertisement</b>                             | Status of advertisements that the peer sends to the local interface.                                                                                                                                                                                                                                            |
| <b>Enable</b>                                         |                                                                                                                                                                                                                                                                                                                 |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                       | (DCBX Version 1.01 only)<br><br>State that the peer advertises to the local interface: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                                                               |
| <b>TLV Type</b>                                       | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Configuration/Recommendation</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>                                        | Willingness of the peer to learn the ETS state from the local interface using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—Peer is willing to learn the ETS state from the local interface.</li> <li>• <b>No</b>—Peer is not willing to learn the ETS state from the local interface.</li> </ul>                                                                                                                                                                                                                                                             |
| <b>Credit Based Shaper</b>                            | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Error</b>                                          | (DCBX Version 1.01 only)<br><br>Configuration error status of the peer: <ul style="list-style-type: none"> <li>• <b>No</b>—No error in peer ETS TLV.</li> <li>• <b>Yes</b>—Error in peer ETS TLV.</li> </ul>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)                                                                                                                                                                                                                                                                                                                                                    |
| <b>Code Point</b>                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                  | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.                                                                                                                                                                                                                                                                                                                         |
| Priority-Group                   | CoS priority group (forwarding class set) identification number.                                                                                                                                                                                                                                                                                                                                                   |
| Percentage B/W                   | Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)                                                                                                                                                                                                                                      |
| Transmission Selection Algorithm | (IEEE DCBX only)<br><br>Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .                                                                                                                                                                                    |
| PFC                              | (QFX Series, <b>terse</b> option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> <li>• Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled</li> <li>• Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled</li> <li>• Not Advt—Interface does not advertise PFC to the connected peer</li> </ul> |
| ETS                              | ( <b>terse</b> option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> <li>• Advt—Interface advertises ETS TLVs</li> <li>• Disabled—ETS is disabled on the interface (interface does not advertise ETS)</li> </ul>                                                                                                                                                             |
| ETS Rec                          | ( <b>terse</b> option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> <li>• Advt—Peer interface advertises ETS TLVs</li> <li>• Not Advt—Peer interface does not advertise ETS</li> </ul> <p><b>NOTE:</b> When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>                                          |

Table 577: show dcbx neighbors Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version    | <p>(<b>terse</b> option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> <li>• <b>IEEE</b>—The interface uses IEEE DCBX.</li> <li>• <b>1.01</b>—The interface uses DCBX version 1.01.</li> </ul> <p>When the DCBX version used is the result of autonegotiation, the term (<b>Auto</b>) appears next to the version. For example, <b>IEEE (Auto)</b> indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p> |

## Sample Output

### show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
 Operational version: 1
 sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
 Operational version: 1
 sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode Operational Mode
000 Disabled Disable
001 Disabled Disable
010 Disabled Disable
011 Enabled Enable
100 Enabled Enable
101 Disabled Disable
110 Disabled Disable
111 Disabled Disable

Peer-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 8

Code Point Admin Mode
000 Disabled

```

|     |          |
|-----|----------|
| 001 | Disabled |
| 010 | Disabled |
| 011 | Enabled  |
| 100 | Enabled  |
| 101 | Disabled |
| 110 | Disabled |
| 111 | Disabled |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001110     | Enabled |
| iSCSI     |               | 3260          | 10000000     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        | N/A           | 00001110     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |

|                |                |
|----------------|----------------|
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 0              | 40%            |
| 1              | 5%             |

### show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

```
user@switch> show dcbx neighbors interface xe-0/0/0
```

```
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
```

```
Active-application-map: app-map-1
```

```
Protocol-Mode: IEEE-DCBX Version
```

```
Feature: PFC
```

```
Local-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

```
Peer-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

```
Feature: Application
```

```
Local-Advertisement:
```

| Appl-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
| FCoE      | 0x8906        |               | 00001110       |
| iSCSI     |               | 3260          | 10000000       |

```
Peer-Advertisement:
```

| Appl-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
|-----------|---------------|---------------|----------------|

|      |        |     |          |
|------|--------|-----|----------|
| FCoE | 0x8906 | N/A | 00001110 |
|------|--------|-----|----------|

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Recommendation

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |

|                |                                  |
|----------------|----------------------------------|
| 101            | 1                                |
| 110            | 1                                |
| 111            | 7                                |
| Priority-Group | Percentage B/W                   |
| 0              | 40%                              |
| 1              | 5%                               |
| Priority-Group | Transmission Selection Algorithm |
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

**show dcbx neighbors terse (QFX Series)**

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface Rec
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

**show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)**

```

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
 Operational version: 0
 sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
 Enable: Yes, Willing: No, Error: No
 Maximum Traffic Classes capable to support PFC: 6

Code Point Admin Mode
000 Disabled
001 Disabled
010 Disabled
011 Enabled
100 Disabled
101 Disabled
110 Disabled
111 Disabled

```

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Status  | Appl-Name | Ethernet-Type | Socket-Number | Priority-Map |
|---------|-----------|---------------|---------------|--------------|
| Enabled | FCoE      | 0x8906        |               | 00001000     |

### show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)

user@switch&gt; show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

## Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

## Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

**show dcbx neighbors (EX4500 Switch: Includes ETS)**

user@switch&gt; show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0  
Protocol-State: in-sync  
Active-application-map: map\_iscsi

## Local-Advertisement:

Operational version: 0



sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Enabled    |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00000001     | Enabled |
| iscsi     |               | 3260          | 00000010     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00010000     | Enabled |
| iscsi     |               | 3260          | 00010000     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes supported : 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 7              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 7              |
| 101        | 7              |
| 110        | 7              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 7              | 100%           |

## Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes supported : 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 1              |
| 010        | 0              |
| 011        | 0              |
| 100        | 2              |
| 101        | 0              |
| 110        | 0              |
| 111        | 0              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 30%            |
| 1              | 40%            |
| 2              | 30%            |

## show interfaces queue

**Syntax** show interfaces queue  
 <aggregate | remaining-traffic>  
 <both-ingress-egress>  
 <egress>  
 <forwarding-class *forwarding-class*>  
 <ingress>  
 <interface-name *interface-name*>  
 <l2-statistics>  
 <remaining-traffic>

**Release Information** Command introduced before Junos OS Release 7.4.  
**both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
**l2-statistics** option introduced in Junos OS Release 12.1.

**Description** Display class-of-service (CoS) queue information for physical interfaces.

**Options** **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(optional) Display layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 252 on page 2708](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the layer 3 level. In the case of link fragmentation and interleaving (LFI) for which not fragmentation is applied, corresponding layer 2 overheads are added, as shown in [Table 252 on page 2708](#).

Table 578: Layer 2 Overhead, Transmitted Packets/Bytes

| Protocol       | Fragmentation       |                            | LFI |
|----------------|---------------------|----------------------------|-----|
|                | First fragmentation | Second to n fragmentations |     |
|                | Bytes               | Bytes                      |     |
| MLPPP (Long)   | 13                  | 12                         | 8   |
| MLPPP (short)  | 11                  | 10                         | 8   |
| MLFR (FRF15)   | 12                  | 10                         | 8   |
| MFR (FRF16)    | 10                  | 8                          | -   |
| MCMLPPP(Long)  | 13                  | 12                         | -   |
| MCMLPPP(Short) | 11                  | 10                         | -   |

## Layer 2 Statistics - Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header : 1 byte
HDLC flag and FCS bytes : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
Inner proto : 2 bytes
HDLC flag and FCS : 4 bytes

```

## Fragments 2 ...n :

```

Framereelay header : 2 bytes
Control,NLPID : 2 bytes
Fragmentaion header : 2 bytes
HDLC flag and FCS : 4 bytes

```

## MFR (FRF16) Overhead details:

=====

```

Fragment 1:
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 Inner proto : 2 bytes
 HDLC flag and FCS : 4 bytes

Fragments 2 ...n :
 Fragmentation header : 2 bytes
 Framereelay header : 2 bytes
 HDLC flag and FCS : 4 bytes

```

## Overhead with LFI

```

MLPPP(Long & short sequence):
=====
 Outer PPP header : 4 bytes
 HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
 Framereelay header : 2 bytes
 Control,NLPID : 2 bytes
 HDLC flag and FCS : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

**Additional Information** On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular

physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos® OS Network Interfaces*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level** view

**List of Sample Output**

- [show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 6022](#)
- [show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 6024](#)
- [show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 6025](#)
- [show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 6025](#)
- [show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 6029](#)
- [show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 6032](#)
- [show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 6034](#)
- [show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 6035](#)
- [show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 6036](#)
- [show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 6039](#)
- [show interfaces queue \(QFX Series\) on page 6049](#)
- [show interfaces queue l2-statistics \(lsq interface\) on page 6050](#)

**Output Fields** [Table 253 on page 2710](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 579: show interfaces queue Output Fields**

| Field Name                          | Field Description                                                                                                                    |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>           | Name of the physical interface.                                                                                                      |
| <b>Enabled</b>                      | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . |
| <b>Interface index</b>              | Physical interface's index number, which reflects its initialization sequence.                                                       |
| <b>SNMP ifIndex</b>                 | SNMP index number for the interface.                                                                                                 |
| <b>Forwarding classes supported</b> | Total number of forwarding classes supported on the specified interface.                                                             |

Table 579: show interfaces queue Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding classes in use</b> | Total number of forwarding classes in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Ingress queues supported</b>  | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                             |
| <b>Ingress queues in use</b>     | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                |
| <b>Output queues supported</b>   | Total number of output queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output queues in use</b>      | Total number of output queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Egress queues supported</b>   | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Egress queues in use</b>      | Total number of egress queues in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Queue</b>                     | Queue number.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Queue counters (Ingress)</b>  | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                           |
| <b>Burst size</b>                | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.                                                                                                                                                                                                                                                     |
| <b>Forwarding classes</b>        | Forwarding class name.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Queued Packets</b>            | Number of packets queued to this queue. <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p>                               |
| <b>Queued Bytes</b>              | Number of bytes queued to this queue. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a> .                                                                                                                                                                                                                                                                                          |
| <b>Transmitted Packets</b>       | Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values. <p><b>NOTE:</b> For layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2707</a></p> |

Table 579: show interfaces queue Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transmitted Bytes</b>    | <p>Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For layer 2 statistics, see “<a href="#">Overhead for Layer 2 Statistics</a>” on page 2707</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Tail-dropped packets</b> | Number of packets dropped because of tail drop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>RED-dropped packets</b>  | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> |
| <b>RED-dropped bytes</b>    | <p>Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see <a href="#">Table 254 on page 2713</a>.</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> <li>• (J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority bytes dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority bytes dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul> </li> </ul>                           |

Byte counts vary by PIC type. [Table 254 on page 2713](#) shows how the byte counts on the outbound interfaces vary depending on the PIC type. [Table 254 on page 2713](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.



Table 580: Byte Count by PIC Type

| PIC Type                         | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Comments                                                                                                                                                                                                     |
|----------------------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Gigabit Ethernet IQ and IQE PICs | Interface                   | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
|                                  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | —                                                                                                                                                                                                            |
| Non-IQ PIC                       | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Switches:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.</li> </ul> | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>                                                                                                                  |

Table 580: Byte Count by PIC Type (*continued*)

| PIC Type                                             | Output Level                | Byte Count Includes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Comments                                                                                                                           |
|------------------------------------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| IQ and IQE PICs with a SONET/SDH interface           | Interface                   | <p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>                                                                                                                                                                                                                       | The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.                                                   |
|                                                      | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>                                                                                                                                                                                                                                                                                                                              | For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.                      |
| Non-IQ PIC with a SONET/SDH interface                | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul> | For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP). |
| Interfaces configured with Frame Relay Encapsulation | Interface                   | The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                    |
| 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs        | Interface                   | <p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>                                                                                                                                                                                                                                                                                                                        | The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.                                               |
| 4-port 1G IQ2 and IQ2-E PICs                         | Packet forwarding component | Queued: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | —                                                                                                                                  |
| 8-port 1G IQ2 and IQ2-E PICs                         |                             | Transmitted: 478 bytes of Layer 3 packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                    |

## Sample Output

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```

user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 242 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 42603765 595484 pps
 Bytes : 5453281920 609776496 bps
 Transmitted:
 Packets : 42603765 595484 pps
 Bytes : 5453281920 609776496 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Transmitted:
 Packets : 45 0 pps
 Bytes : 3930 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 4, Forwarding classes: af11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 5, Forwarding classes: ef11
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Transmitted:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
Queued:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Transmitted:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
Queue: 1, Forwarding classes: af1
Queued:
Packets : 2480391 1650 pps
Bytes : 1304685666 6945704 bps
Transmitted:
Packets : 2478740 1650 pps
Bytes : 1303817240 6945704 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 1651 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 1651 0 pps

```

```

RED-dropped bytes : 868426 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 868426 0 pps

```

### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Transmitted:
 Packets : 13 0 pps
 Bytes : 622 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af1
 Queued:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Transmitted:
 Packets : 1725947945 372178 pps
 Bytes : 220921336960 381110432 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef1
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Transmitted:
 Packets : 571 0 pps
 Bytes : 49318 336 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use

```

```

Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 148450735 947295 pps
 Bytes : 8016344944 409228848 bps
Transmitted:
 Packets : 76397439 487512 pps
 Bytes : 4125461868 210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets : 72053285 459783 pps
 Low : 72053285 459783 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 3890877444 198626472 bps
 Low : 3890877444 198626472 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 410278257 473940 pps
 Bytes : 22156199518 204742296 bps
Transmitted:
 Packets : 4850003 4033 pps
 Bytes : 261900162 1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets : 405425693 469907 pps
 Low : 405425693 469907 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 21892988124 203000040 bps
 Low : 21892988124 203000040 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 76605230 485376 pps
 Bytes : 5209211400 264044560 bps
 Transmitted:
 Packets : 76444631 484336 pps
 Bytes : 5198235612 263478800 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 160475 1040 pps
 Low : 160475 1040 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 10912300 565760 bps
 Low : 10912300 565760 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 4836136 3912 pps
 Bytes : 333402032 2139056 bps
 Transmitted:
 Packets : 3600866 1459 pps
 Bytes : 244858888 793696 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 1225034 2450 pps
 Low : 1225034 2450 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps

```

```

 High : 0 0 pps
 RED-dropped bytes : 83302312 1333072 bps
 Low : 83302312 1333072 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

 Queued:
 Packets : 77059796 486384 pps
 Bytes : 3544750624 178989576 bps
 Transmitted:
 Packets : 77059797 486381 pps
 Bytes : 3544750670 178988248 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps

```



```

 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Transmitted:
 Packets : 4846580 3934 pps
 Bytes : 222942680 1447768 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
 Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
 use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Transmitted:
 Packets : 418390039 10 pps
 Bytes : 38910269752 7440 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Transmitted:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

## Queue: 3, Forwarding classes: network-control

## Queued:

|         |   |         |          |
|---------|---|---------|----------|
| Packets | : | 77009   | 11 pps   |
| Bytes   | : | 6894286 | 7888 bps |

## Transmitted:

|         |   |         |          |
|---------|---|---------|----------|
| Packets | : | 77009   | 11 pps   |
| Bytes   | : | 6894286 | 7888 bps |

Tail-dropped packets : Not Available

|                    |   |   |       |
|--------------------|---|---|-------|
| RL-dropped packets | : | 0 | 0 pps |
|--------------------|---|---|-------|

|                  |   |   |       |
|------------------|---|---|-------|
| RL-dropped bytes | : | 0 | 0 bps |
|------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

## Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

## Queue: 0, Forwarding classes: best-effort

## Queued:

|         |   |        |       |
|---------|---|--------|-------|
| Packets | : | 1031   | 0 pps |
| Bytes   | : | 147328 | 0 bps |

## Transmitted:

|         |   |        |       |
|---------|---|--------|-------|
| Packets | : | 1031   | 0 pps |
| Bytes   | : | 147328 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 pps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 pps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 pps |
|-----------|---|---|-------|

RED-dropped bytes : 0 0 bps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 bps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 bps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 bps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 bps |
|-----------|---|---|-------|

## Queue: 1, Forwarding classes: expedited-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 pps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 pps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 pps |
|-----------|---|---|-------|

RED-dropped bytes : 0 0 bps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 bps |
|--------------|---|---|-------|

|          |   |   |       |
|----------|---|---|-------|
| Low, TCP | : | 0 | 0 bps |
|----------|---|---|-------|

|               |   |   |       |
|---------------|---|---|-------|
| High, non-TCP | : | 0 | 0 bps |
|---------------|---|---|-------|

|           |   |   |       |
|-----------|---|---|-------|
| High, TCP | : | 0 | 0 bps |
|-----------|---|---|-------|

## Queue: 2, Forwarding classes: assured-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

|              |   |   |       |
|--------------|---|---|-------|
| Low, non-TCP | : | 0 | 0 pps |
|--------------|---|---|-------|

```

 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
Transmitted:
 Packets : 94386 12 pps
 Bytes : 13756799 9568 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 Low, non-TCP : 0 0 pps
 Low, TCP : 0 0 pps
 High, non-TCP : 0 0 pps
 High, TCP : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low, non-TCP : 0 0 bps
 Low, TCP : 0 0 bps
 High, non-TCP : 0 0 bps
 High, TCP : 0 0 bps

```

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
 Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 254 0 pps
 Bytes : 16274 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 3 0 pps
 Bytes : 126 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : Not Available
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps

```

```

Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
Transmitted:
 Packets : 9397 0 pps
 Bytes : 3809052 232 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 288 0 pps
 Bytes : 18450 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

```

Queue: 3, Forwarding classes: network-control
Queued:
 Packets : Not Available
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
Transmitted:
 Packets : 80564692 0 pps
 Bytes : 3383717100 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
Transmitted:
 Packets : 80564685 0 pps
 Bytes : 3383716770 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 9538 0 pps
 Bytes : 3819840 0 bps
Transmitted:
 Packets : 9538 0 pps
 Bytes : 3819840 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps

```

#### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```



```

Queued:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
Transmitted:
 Packets : 110208969 472875 pps
 Bytes : 5951284434 204282000 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 109355853 471736 pps
 Bytes : 7436199152 256627968 bps
 Transmitted:
 Packets : 109355852 471736 pps
 Bytes : 7436198640 256627968 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps

```

```

RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

 Packets : 214886 13449 pps
 Bytes : 9884756 5164536 bps

Transmitted:

 Packets : 214886 13449 pps
 Bytes : 9884756 5164536 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps

```

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

## Queue: 1, Forwarding classes: REALTIME

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 2, Forwarding classes: PRIVATE

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|         |   |      |       |
|---------|---|------|-------|
| Packets | : | 60   | 0 pps |
| Bytes   | : | 4560 | 0 bps |

Transmitted:

|                      |   |      |       |
|----------------------|---|------|-------|
| Packets              | : | 60   | 0 pps |
| Bytes                | : | 4560 | 0 bps |
| Tail-dropped packets | : | 0    | 0 pps |
| RED-dropped packets  | : | 0    | 0 pps |
| Low                  | : | 0    | 0 pps |
| Medium-low           | : | 0    | 0 pps |
| Medium-high          | : | 0    | 0 pps |
| High                 | : | 0    | 0 pps |
| RED-dropped bytes    | : | 0    | 0 bps |
| Low                  | : | 0    | 0 bps |
| Medium-low           | : | 0    | 0 bps |
| Medium-high          | : | 0    | 0 bps |
| High                 | : | 0    | 0 bps |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

|         |   |          |             |
|---------|---|----------|-------------|
| Packets | : | 371365   | 23620 pps   |
| Bytes   | : | 15597330 | 7936368 bps |

Transmitted:

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |
| High                 | : | 0        | 0 pps       |
| RED-dropped bytes    | : | 0        | 0 bps       |
| Low                  | : | 0        | 0 bps       |
| Medium-low           | : | 0        | 0 bps       |
| Medium-high          | : | 0        | 0 bps       |



|                                        |   |   |       |
|----------------------------------------|---|---|-------|
| High                                   | : | 0 | 0 bps |
| Queue: 1, Forwarding classes: REALTIME |   |   |       |
| Queued:                                |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Transmitted:                           |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Tail-dropped packets                   | : | 0 | 0 pps |
| RED-dropped packets                    | : | 0 | 0 pps |
| Low                                    | : | 0 | 0 pps |
| Medium-low                             | : | 0 | 0 pps |
| Medium-high                            | : | 0 | 0 pps |
| High                                   | : | 0 | 0 pps |
| RED-dropped bytes                      | : | 0 | 0 bps |
| Low                                    | : | 0 | 0 bps |
| Medium-low                             | : | 0 | 0 bps |
| Medium-high                            | : | 0 | 0 bps |
| High                                   | : | 0 | 0 bps |
| Queue: 2, Forwarding classes: PRIVATE  |   |   |       |
| Queued:                                |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Transmitted:                           |   |   |       |
| Packets                                | : | 0 | 0 pps |
| Bytes                                  | : | 0 | 0 bps |
| Tail-dropped packets                   | : | 0 | 0 pps |
| RED-dropped packets                    | : | 0 | 0 pps |
| Low                                    | : | 0 | 0 pps |
| Medium-low                             | : | 0 | 0 pps |
| Medium-high                            | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

## Queue: 3, Forwarding classes: CONTROL

## Queued:

|         |   |         |        |
|---------|---|---------|--------|
| Packets | : | 32843   | 0 pps  |
| Bytes   | : | 2641754 | 56 bps |

## Transmitted:

|                      |   |         |        |
|----------------------|---|---------|--------|
| Packets              | : | 32843   | 0 pps  |
| Bytes                | : | 2641754 | 56 bps |
| Tail-dropped packets | : | 0       | 0 pps  |
| RED-dropped packets  | : | 0       | 0 pps  |
| Low                  | : | 0       | 0 pps  |
| Medium-low           | : | 0       | 0 pps  |
| Medium-high          | : | 0       | 0 pps  |
| High                 | : | 0       | 0 pps  |
| RED-dropped bytes    | : | 0       | 0 bps  |
| Low                  | : | 0       | 0 bps  |
| Medium-low           | : | 0       | 0 bps  |
| Medium-high          | : | 0       | 0 bps  |
| High                 | : | 0       | 0 bps  |

## Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: fcoe
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : Not Available
 Total-dropped packets: 0 0 pps
 Total-dropped bytes : 0 0 bps
Queue: 8, Forwarding classes: mcast
 Queued:

```

|                                      |   |   |       |
|--------------------------------------|---|---|-------|
| Packets                              | : | 0 | 0 pps |
| Bytes                                | : | 0 | 0 bps |
| Transmitted:                         |   |   |       |
| Packets                              | : | 0 | 0 pps |
| Bytes                                | : | 0 | 0 bps |
| Tail-dropped packets : Not Available |   |   |       |
| Total-dropped packets:               |   | 0 | 0 pps |
| Total-dropped bytes :                |   | 0 | 0 bps |

### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
 Queued:
 Packets : 1 0 pps
 Bytes : 1001 0 bps
 Transmitted:
 Packets : 5 0 pps
 Bytes : 1062 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: ef
 Queued:
 Packets : 1 0 pps
 Bytes : 1500 0 bps
 Transmitted:
 Packets : 6 0 pps
 Bytes : 1573 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: af
 Queued:
 Packets : 1 0 pps
 Bytes : 512 0 bps
 Transmitted:
 Packets : 3 0 pps
 Bytes : 549 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: nc
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
=====

```

## show pfe next-hop

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 6051</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 6051</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                                        | <pre>show pfe next-hop &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <pre>show pfe next-hop &lt;fpc <i>slot</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;lcc <i>number</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                           | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                   | Display Packet Forwarding Engine next-hop information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                       | <p><b>none</b>—Display all Packet Forwarding Engine next-hop information.</p> <p><b>fpc <i>slot</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> <li>On a TX Matrix router, if you specify the number of a T640 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31.</li> <li>On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31.</li> <li>On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 63.</li> </ul> <p>For example, the following commands have the same result:</p> <pre>user@host&gt; show pfe next-hop fpc 1 lcc 1 user@host&gt; show pfe next-hop fpc 9</pre> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Packet Forwarding Engine next-hop interface.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p> |

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Overview of a Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

- [show pfe next-hop on page 6053](#)
- [show pfe next-hop fpc \(TX Matrix Router\) on page 6053](#)
- [show pfe next-hop fpc \(TX Matrix Plus Router\) on page 6053](#)



## Sample Output

### show pfe next-hop

```
user@host> show pfe next-hop
NextHop Info:
 ID Type Interface Protocol Encap Next Hop Addr MTU
 ---- - - - - - -
 4 Mcast - IPv4 - 0.0.0.0 0
 5 Bcast - IPv4 - - 0
 7 Discard - IPv4 - - 0
 8 MDiscard - IPv4 - - 0
 9 Reject - IPv4 - - 0
 13 Local - IPv4 - 192.168.4.60 0
 14 Resolve fxp0.0 IPv4 Unspecified - 0
 17 Local - IPv4 - 127.0.0.1 0
 18 Unicast fxp0.0 IPv4 Unspecified 192.168.4.254 0
 21 Local - IPv4 - 11.1.0.1 0
 22 Unicast at-0/1/0.0 IPv4 ATM SNAP 11.1.0.2 4482
 ...
```

### show pfe next-hop fpc (TX Matrix Router)

```
user@host> show pfe next-hop fpc 1
Slot 1
NextHop Info:
 ID Type Interface Next Hop Addr Protocol Encap MTU
 ---- - - - - - -
 5 Mcast - default IPv4 - 0
 6 Bcast - - IPv4 - 0
 8 Discard - - IPv4 - 0
 9 MDiscard - - IPv4 - 0
 13 Mcast - default IPV6 - 0
 17 MDiscard - - IPV6 - 0
 18 Reject - - IPV6 - 0
 24 Discard - - None - 0
 68 Local - 192.168.66.113 IPv4 - 0
 69 Resolve fxp0.0 - IPv4 Unspecified 0
 70 Unicast fxp0.0 192.168.71.254 IPv4 Unspecified 0
 256 Local - 10.71.71.1 IPv4 - 0
 257 Local - 127.0.0.1 IPv4 - 0
 258 Mcast.local..1 default IPv4 Unspecified 0
 259 Bcast.local..1 - IPv4 Unspecified 0
 261 Discard.local..1 - IPv4 Unspecified 0
 262 MDiscard.local..1 - IPv4 Unspecified 0
 269 Mcast.local..1 default IPV6 Unspecified 0
 271 Discard.local..1 - IPV6 Unspecified 0
 ...
```


### show pfe next-hop fpc (TX Matrix Plus Router)

```
user@host> show pfe next-hop fpc 0
Slot 0
 ID Type Interface Next Hop Addr Protocol Encap MTU
 ---- - - - - - -
 31 Mcast - default IPv4 - 0
 32 Bcast - - IPv4 - 0
 34 Discard - - IPv4 - 0
 35 MDiscard - - IPv4 - 0
```

|     |          |             |                      |      |             |   |
|-----|----------|-------------|----------------------|------|-------------|---|
| 36  | Reject   | -           | -                    | IPv4 | -           | 0 |
| 39  | Mcast    | -           | default              | IPv6 | -           | 0 |
| 42  | Discard  | -           | -                    | IPv6 | -           | 0 |
| 43  | MDiscard | -           | -                    | IPv6 | -           | 0 |
| 44  | Reject   | -           | -                    | IPv6 | -           | 0 |
| 49  | Receive  | -           | -                    | MPLS | -           | 0 |
| 50  | Discard  | -           | -                    | MPLS | -           | 0 |
| 111 | Mcast    | .local..1   | default              | IPv4 | Unspecified | 0 |
| 112 | Bcast    | .local..1   | -                    | IPv4 | Unspecified | 0 |
| 114 | Discard  | .local..1   | -                    | IPv4 | Unspecified | 0 |
| 115 | MDiscard | .local..1   | -                    | IPv4 | Unspecified | 0 |
| 116 | Reject   | .local..1   | -                    | IPv4 | Unspecified | 0 |
| 119 | Mcast    | .local..1   | default              | IPv6 | Unspecified | 0 |
| 122 | Discard  | .local..1   | -                    | IPv6 | Unspecified | 0 |
| 123 | MDiscard | .local..1   | -                    | IPv6 | Unspecified | 0 |
| 124 | Reject   | .local..1   | -                    | IPv6 | Unspecified | 0 |
| 191 | Mcast    | .local..2   | default              | IPv4 | Unspecified | 0 |
| 192 | Bcast    | .local..2   | -                    | IPv4 | Unspecified | 0 |
| 194 | Discard  | .local..2   | -                    | IPv4 | Unspecified | 0 |
| 195 | MDiscard | .local..2   | -                    | IPv4 | Unspecified | 0 |
| 196 | Reject   | .local..2   | -                    | IPv4 | Unspecified | 0 |
| 322 | Local    | -           | 10.1.0.5             | IPv4 | -           | 0 |
| 323 | Resolve  | bcm0.0      | -                    | IPv4 | Unspecified | 0 |
| 326 | Local    | -           | 129.0.0.5            | IPv4 | -           | 0 |
| 327 | Resolve  | bcm0.0      | -                    | IPv4 | Unspecified | 0 |
| 328 | Local    | -           | fe80::201:ff:fe01:5  | IPv6 | -           | 0 |
| 329 | Receive  | bcm0.0      | ff02::1:ff01:5       | IPv6 | Unspecified | 0 |
| 330 | Receive  | bcm0.0      | fe80::               | IPv6 | Unspecified | 0 |
| 331 | Resolve  | bcm0.0      | -                    | IPv6 | Unspecified | 0 |
| 332 | Local    | -           | fec0::a:1:0:5        | IPv6 | -           | 0 |
| 333 | Receive  | bcm0.0      | ff02::1:ff00:5       | IPv6 | Unspecified | 0 |
| 334 | Receive  | bcm0.0      | fec0::               | IPv6 | Unspecified | 0 |
| 335 | Resolve  | bcm0.0      | -                    | IPv6 | Unspecified | 0 |
| 348 | Local    | -           | 192.168.178.4        | IPv4 | -           | 0 |
| 349 | Resolve  | em0.0       | -                    | IPv4 | Unspecified | 0 |
| 350 | Unicast  | em0.0       | 192.168.178.126      | IPv4 | Unspecified | 0 |
| 357 | Local    | -           | fe80::201:1ff:fe01:5 | IPv6 | -           | 0 |
| 512 | Local    | -           | 10.255.178.11        | IPv4 | -           | 0 |
| 513 | Local    | -           | 127.0.0.1            | IPv4 | -           | 0 |
| 515 | Local    | -           | abcd::10:255:178:11  | IPv6 | -           | 0 |
| 516 | Local    | -           | fe80::200:ff:fe00:0  | IPv6 | -           | 0 |
| 517 | Local    | -           | 127.0.0.1            | IPv4 | -           | 0 |
| 518 | Mcast    | .local..3   | default              | IPv4 | Unspecified | 0 |
| 519 | Bcast    | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 521 | Discard  | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 522 | MDiscard | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 523 | Reject   | .local..3   | -                    | IPv4 | Unspecified | 0 |
| 531 | Mcast    | .local..3   | default              | IPv6 | Unspecified | 0 |
| 533 | Discard  | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 534 | MDiscard | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 535 | Reject   | .local..3   | -                    | IPv6 | Unspecified | 0 |
| 539 | Mgroup   | -           | -                    | IPv4 | -           | 0 |
| 540 | Bcast    | ge-15/0/3.0 | -                    | IPv4 | Ethernet    | 0 |
| 541 | Receive  | ge-15/0/3.0 | 14.2.1.0             | IPv4 | Ethernet    | 0 |
| 542 | Local    | -           | 14.2.1.1             | IPv4 | -           | 0 |
| 543 | Resolve  | ge-15/0/3.0 | -                    | IPv4 | Ethernet    | 0 |
| 544 | Bcast    | ge-31/0/4.0 | -                    | IPv4 | Ethernet    | 0 |

|     |         |             |          |      |          |   |
|-----|---------|-------------|----------|------|----------|---|
| 545 | Receive | ge-31/0/4.0 | 14.1.1.0 | IPv4 | Ethernet | 0 |
| 546 | Local   | -           | 14.1.1.1 | IPv4 | -        | 0 |
| 547 | Resolve | ge-31/0/4.0 | -        | IPv4 | Ethernet | 0 |
| 548 | Unicast | ge-31/0/4.0 | 14.1.1.2 | IPv4 | Ethernet | 0 |
| 549 | Unicast | ge-15/0/3.0 | 14.2.1.2 | IPv4 | Ethernet | 0 |
| 550 | Bcast   | ae1.0       | -        | IPv4 | Ethernet | 0 |
| 551 | Receive | ae1.0       | 11.1.1.0 | IPv4 | Ethernet | 0 |
| 552 | Local   | -           | 11.1.1.1 | IPv4 | -        | 0 |
| 553 | Resolve | ae1.0       | -        | IPv4 | Ethernet | 0 |
| 554 | Aggreg. | ae1.0       | -        | IPv4 | Ethernet | 0 |
| 555 | Unicast | ge-23/0/8.0 | 11.1.1.2 | IPv4 | Ethernet | 0 |
| 556 | Unicast | ge-7/0/9.0  | 11.1.1.2 | IPv4 | Ethernet | 0 |
| 557 | Aggreg. | ae1.0       | -        | MPLS | Ethernet | 0 |
| 558 | Unicast | ge-23/0/8.0 | -        | MPLS | Ethernet | 0 |
| 559 | Unicast | ge-7/0/9.0  | -        | MPLS | Ethernet | 0 |
| 560 | Aggreg. | ae1.0       | -        | MPLS | Ethernet | 0 |
| 561 | Unicast | ge-23/0/8.0 | -        | MPLS | Ethernet | 0 |
| 562 | Unicast | ge-7/0/9.0  | -        | MPLS | Ethernet | 0 |

## show pfe route

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <a href="#">Syntax on page 6056</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 6056</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 6056</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <pre>show pfe route &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and QFX Series)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <pre>show pfe route &lt;fpc slot&gt; &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;lcc number&gt; &lt;mpls&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <div>  <p><b>NOTE:</b> The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the <code>show route forwarding table</code> command. For more information, see the <a href="#">CLI Explorer</a>.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p><b>none</b>—Display all Packet Forwarding Engine forwarding table information.</p> <p><b>fpc slot</b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> <li>On a TX Matrix router, if you specify the number of a T640 router by using the <b>lcc number</b> option (the recommended method), replace <b>slot</b> with a value from <b>0</b> through <b>7</b>. Otherwise, replace <b>slot</b> with a value from <b>0</b> through <b>31</b>.</li> <li>On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <b>lcc number</b> option (the recommended method), replace <b>slot</b> with a value from <b>0</b> through <b>7</b>. Otherwise, replace <b>slot</b> with a value from <b>0</b> through <b>31</b>.</li> </ul> |

- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 63.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

**index index**—(Optional) Display table index.

**inet6**—(Optional) Display Packet Forwarding Engine IPv6 routes.

**ip**—(Optional) Display Packet Forwarding Engine IPv4 routes.

**iso**—(Optional) Display ISO version routing tables.

**lcc number**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**mpls**—(Optional) Display Packet Forwarding Engine MPLS information.

**prefix prefix**—(Optional) IPv4 or IPv6 prefix for which to show table entries.

**summary**—(Optional) Display summary of Packet Forwarding Engine information.

**table <table-name>**—(Optional) Display table information.

**Required Privilege Level**

admin

**Related Documentation**

- Routing Matrix with TXP-T1600 Configuration*
- Routing Matrix with TXP-T1600-3D Configuration*
- Routing Matrix with TXP-T4000-3D Configuration*
- Overview of a Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

- List of Sample Output
- [show pfe route ip on page 6058](#)
  - [show pfe route iso on page 6058](#)
  - [show pfe route lcc summary \(TX Matrix Router\) on page 6058](#)
  - [show pfe route lcc summary \(TX Matrix Plus Router\) on page 6060](#)

## Sample Output

### show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
Destination NH IP Addr Type NH ID Interface

default 127.0.0.1 Discard 8
127.0.0.1 127.0.0.1 Local 256
172.16/12 192.168.71.254 Unicast 68 fxp0.0
192.168.0/18 192.168.71.254 Unicast 68 fxp0.0
192.168.40/22 192.168.71.254 Unicast 68 fxp0.0
192.168.64/18 192.168.71.254 Unicast 68 fxp0.0
192.168.64/21 192.168.71.254 Resolve 67 fxp0.0
192.168.71.249 192.168.71.249 Local 66
192.168.220.0/30 192.168.71.249 Resolve 303 fe-0/0/0.0
192.168.220.0 192.168.220.0 Receive 301 fe-0/0/0.0
224.0.0.1 Mcast 5
255.255.255.255 Bcast 6

...
```

### show pfe route iso

```
user@host# show pfe route iso
```

```
CLNS Route Table 0, CLNP.0, 0x0:
Destination Type NH ID Interface

default Reject 60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152 Local 514
49.0001.00a0.c96b.c491/72 Local 536
```

### show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary
```

```
Slot 0
```

```
IPv4 Route Tables:
Index Routes Size(b)

Default 43 3081
1 4 281
```

```
MPLS Route Tables:
Index Routes Size(b)

Default 1 68
```

```
IPv6 Route Tables:
Index Routes Size(b)

Default 9 717
```

1 5 389

#### Slot 1

##### IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 43     | 3081    |
| 1       | 4      | 281     |

##### MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 1      | 68      |

##### IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 9      | 717     |
| 1       | 5      | 389     |

#### Slot 16

##### IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 41     | 2938    |
| 1       | 4      | 281     |

##### MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 1      | 68      |

##### IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 9      | 717     |
| 1       | 5      | 389     |

#### Slot 17

##### IPv4 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 41     | 2938    |
| 1       | 4      | 281     |

##### MPLS Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 1      | 68      |

##### IPv6 Route Tables:

| Index   | Routes | Size(b) |
|---------|--------|---------|
| Default | 9      | 717     |
| 1       | 5      | 389     |

**show pfe route lcc summary (TX Matrix Plus Router)**

```
user@host> show pfe route lcc 2 summary
```

Slot 0

**IPv4 Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 25     | 2266    |
| 1       | 9      | 815     |
| 2       | 6      | 545     |
| 3       | 5      | 453     |
| 4       | 15     | 1371    |
| 5       | 5      | 453     |
| 6       | 13     | 1187    |

**MPLS Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 88      |
| 4       | 5      | 452     |

**IPv6 Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 7      | 697     |
| 1       | 13     | 1305    |
| 3       | 4      | 385     |
| 4       | 4      | 385     |
| 5       | 4      | 385     |
| 6       | 18     | 1833    |

Slot 6

**IPv4 Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 25     | 2266    |
| 1       | 9      | 815     |
| 2       | 6      | 545     |
| 3       | 5      | 453     |
| 4       | 15     | 1371    |
| 5       | 5      | 453     |
| 6       | 13     | 1187    |

**MPLS Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 1      | 88      |
| 4       | 5      | 452     |

**IPv6 Route Tables:**

| Index   | Routes | Size(b) |
|---------|--------|---------|
| -----   | -----  | -----   |
| Default | 7      | 697     |
| 1       | 13     | 1305    |
| 3       | 4      | 385     |
| 4       | 4      | 385     |



|     |    |      |
|-----|----|------|
| 5   | 4  | 385  |
| 6   | 18 | 1833 |
| ... |    |      |

## show pfe terse

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 6062</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 6062</a><br><a href="#">Syntax (MX Series Router) on page 6062</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                       | show pfe terse                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (TX Matrix and TX Matrix Plus Router)</b> | show pfe terse<br><lcc <i>number</i>   scc><br><sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (MX Series Router)</b>                    | show pfe terse<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Display Packet Forwarding Engine status information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                      | <b>none</b> —Display brief information about the Packet Forwarding Engine.<br><br><b>all-members</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.<br><br><b>lcc <i>number</i></b> —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router.<br>Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li><li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li><li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li></ul><br><b>local</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member. |

**member *member-id***—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level** admin

**List of Sample Output** [show pfe terse \(TX Matrix Router\) on page 6063](#)  
[show pfe terse \(TX Matrix Plus Router\) on page 6063](#)  
[show pfe terse sfc \(TX Matrix Plus Router\) on page 6063](#)

## Sample Output

### show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot State Flags Uptime
0 SFM Present Online 0x0bf 01:25:42
2 SFM Present Online 0x0bf 01:25:40
0 FPC Present Online 0x102 01:25:57
1 FPC Present Online 0x102 01:25:55
2 FPC Present Online 0x102 01:25:53
```

### show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:

Slot Type Slot State Uptime
0 LCC Present Online 2d 05:26

lcc0-re0:

Slot Type Slot State Uptime
0 GFPC Present Online 2d 05:25
1 GFPC Present Online 2d 05:25
```

### show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:

Slot Type Slot State Uptime
0 LCC Present Online 2d 05:25
```

## show pfe version

---

|                                 |                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show pfe version <brief   detail>                                                                                  |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Display Packet Forwarding Engine version information.                                                              |
| <b>Options</b>                  | brief   detail—Display the specified level of output.                                                              |
| <b>Required Privilege Level</b> | admin                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show pfe version brief on page 6064</a><br><a href="#">show pfe version detail on page 6064</a>        |

### Sample Output

#### show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

#### show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC

junos-core01.juniper.net:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

## CHAPTER 66

# Troubleshooting

- [Troubleshooting Procedures on page 6065](#)

## Troubleshooting Procedures

---

- [Troubleshooting Dropped FCoE Traffic on page 6065](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 6068](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 6069](#)
- [Troubleshooting Egress Queue Bandwidth Impacted by Congestion on page 6070](#)
- [Troubleshooting an Unexpected Rewrite Value on page 6071](#)
- [Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic on page 6072](#)

## Troubleshooting Dropped FCoE Traffic

**Problem**    **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

**Cause**    There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.
6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Disabled
 100 Disabled
 101 Enabled 2500
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 101
 5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 5113 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 5931](#)
- [show class-of-service forwarding-class-set on page 5938](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

**Problem**    **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth (shaping rate) configured for the queue.

**Cause**        When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.



The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

**Solution** When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

**Related  
Documentation**

- [shaping-rate on page 5904](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

**Problem** **Description:** The minimum bandwidth of a queue or a priority group when measured at the egress port exceeds the minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

**Cause** When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



**NOTE:** The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

**Solution** When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit

rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

**Related  
Documentation**

- [guaranteed-rate on page 5869](#)
- [transmit-rate on page 5911](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

## Troubleshooting Egress Queue Bandwidth Impacted by Congestion

**Problem**    **Description:** Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

**Cause**       Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

**Solution**    The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a tail-drop profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

**Related  
Documentation**

- [drop-profile on page 5851](#)
- [Example: Configuring Tail-Drop Profiles on page 5663](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
- [Understanding CoS Tail-Drop Profiles on page 5545](#)

## Troubleshooting an Unexpected Rewrite Value

**Problem**    **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

**Cause**    If you configure a rewrite rule for a forwarding class on an egress port but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

**Solution**    If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



**TIP:** If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



**NOTE:** There are no default rewrite rules. You can bind one rewrite rule for each type (DSCP and IEEE 802.1) to a given interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value associations.

1. Assign a rewrite value to a forwarding class. Add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
```

```
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **fcoe** is being randomly rewritten, and you want to use IEEE 802.1 code point **011** for the **fcoe** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class fcoe loss-priority high code-point
011
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

#### Related Documentation

- [interfaces on page 5880](#)
- [rewrite-rules on page 5897](#)
- [Defining CoS Rewrite Rules on page 5805](#)
- [Monitoring CoS Rewrite Rules on page 5918](#)

## Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

**Problem**    **Description:** In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.

**Cause**      Failure of a queue to transmit packets for 12 consecutive seconds may be due to:

- A strict-high priority queue consuming all of the port bandwidth
- Several queues consuming all of the port bandwidth
- Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
- Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

**Solution**    If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map
scheduler-map-name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name
output-traffic-control-profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues.

[Table 581 on page 6073](#) shows the topology for this example:

**Table 581: Components of the Rate Shaping Troubleshooting Example**

| Component                                  | Settings                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Affected interface                         | <b>shpnode:xe-0/0/10</b>                                                                                                                                                                                                                                                                                                                          |
| Scheduler (strict-high priority scheduler) | Name: <b>shp-sched</b><br>Shaping rate: <b>7g</b><br>Priority: <b>strict-high</b><br><br><b>NOTE:</b> This example assumes that the scheduler already exists and has been configured as <b>strict-high</b> priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied. |
| Scheduler map                              | Name: <b>shp-map</b><br>Forwarding class to associate with the <b>shp-sched</b> scheduler: <b>strict-high</b><br><br><b>NOTE:</b> This example assumes that a strict-high priority forwarding class has been configured and assigned the name <b>strict-high</b> .                                                                                |
| Traffic control profile                    | Name: <b>shp-tcp</b><br><br><b>NOTE:</b> This example does not describe how to define a complete traffic control profile.                                                                                                                                                                                                                         |

Table 581: Components of the Rate Shaping Troubleshooting Example (*continued*)

| Component                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class set                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Name: <b>shp-pg</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <p>To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface <b>shpnode:xe-0/0/10</b> using the CLI:</p> <ol style="list-style-type: none"> <li>Specify the scheduler for the strict-high priority queue (<b>shp-sched</b>) with a maximum bandwidth of 7 Gbps: <pre>[edit class-of-service schedulers] user@switch# set shp-sched shaping-rate 7g</pre> </li> <li>Configure a scheduler map (<b>shp-map</b>) that associates the scheduler (<b>shp-sched</b>) with the forwarding class (<b>strict-high</b>): <pre>[edit class-of-service scheduler-maps] user@switch# set shp-map forwarding-class strict-high scheduler shp-sched</pre> </li> <li>Associate the scheduler map <b>shp-map</b> with a traffic control profile (<b>shp-tcp</b>): <pre>[edit class-of-service traffic-control-profiles] user@switch# set shp-tcp scheduler-map shp-map</pre> </li> <li>Associate the traffic control profile <b>shp-tcp</b> with a forwarding class set (<b>shp-pg</b>) and the affected interface (<b>shpnode:xe-0/0/10</b>): <pre>[edit class-of-service] user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg output-traffic-control-profile shp-tcp</pre> </li> </ol> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Related Documentation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li> <li>• <a href="#">Defining CoS Queue Scheduling Priority on page 5792</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 5671</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li> </ul> |

## PART 21

# Network Management and Monitoring

- [Overview on page 6077](#)
- [Configuration on page 6159](#)
- [Administration on page 6385](#)
- [Troubleshooting on page 6471](#)





CHAPTER 67

Overview

- [Network Management on page 6077](#)
- [Automation Scripts on page 6083](#)
- [Fabric OAM on page 6099](#)
- [Junos Space on page 6101](#)
- [sFlow Technology on page 6104](#)
- [SNMP on page 6107](#)
- [System Logging on page 6154](#)

Network Management

- [Understanding Device and Network Management Features on page 6077](#)
- [Understanding Network Management Implementation on the QFabric System on page 6080](#)
- [Understanding Telnet on the QFabric System on page 6081](#)
- [Understanding Tracing and Logging Operations on page 6081](#)

Understanding Device and Network Management Features

After you install a QFX Series product in your network, you need to manage the device. The QFX Series products support features that you use to manage the device within the network, including the management of configuration, system performance, fault monitoring, and remote access.

[Table 582 on page 6077](#) lists the device and network management features on the QFX Series.

Table 582: Device and Network Management Features on the QFX Series

| Feature                                                                                                                                                                                                                                                          | Typical Uses     | Documentation                                                       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|---------------------------------------------------------------------|
| AI-Scripts and Advanced Insight Manager (AIM)—Automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems, and submit problem reports to Juniper Support Systems. | Fault management | <a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a> |

Table 582: Device and Network Management Features on the QFX Series (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Typical Uses                                                                                                                                                                         | Documentation                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarms and LEDs on the switch—Show status of hardware components and indicate warning or error conditions.                                                                                                                                                                                                                                                                                                                                                                                     | Fault management                                                                                                                                                                     | <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 6488</a>                                                                                                                                                                        |
| Firewall filters—Control the packets that are sent to and from the network, balance network traffic, and optimize performance.                                                                                                                                                                                                                                                                                                                                                                 | Performance management                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">Routing Policy Configuration Guide</a></li> <li>• <a href="#">Overview of Firewall Filters on page 4635</a></li> </ul>                                                                      |
| In-band management—Enables connection to the switch using the same interfaces through which customer traffic flows. Communication between the switch and a remote console is typically enabled using SSH and Telnet services. SSH provides secure encrypted communications, whereas Telnet provides unencrypted, and therefore less secure, access to the switch.                                                                                                                              | Remote access management                                                                                                                                                             | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1709</a></li> <li>• <a href="#">Configuring Telnet Service for Remote Access to a Router or Switch</a></li> </ul> |
| Juniper Networks Junos OS automation scripts—Configuration and operations automation tools provided by Junos OS. These tools include commit scripts, operation scripts, event scripts, and event policies. Commit scripts enforce custom configuration rules, whereas operation scripts, event policies, and event scripts automate network troubleshooting and management.                                                                                                                    | <ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>                                           | <a href="#">Junos OS Configuration and Operations Automation Guide</a>                                                                                                                                                                           |
| Junos OS command-line interface (CLI)—CLI configuration statements that enable you to configure the switch based on your networking requirements, such as security, service, and performance.                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• User access management</li> <li>• Remote access management</li> </ul> | <a href="#">CLI User Guide</a>                                                                                                                                                                                                                   |
| Junos Space software—Multipurpose GUI-based network management system that includes a base platform, the Network Application Platform, and other optional applications such as Ethernet Design, Service Now, Service Insight, and Virtual Control.                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>                                           | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Junos Space Support on page 6101</a></li> <li>• <a href="#">Junos Space Network Application Platform User Guide</a></li> </ul>                                                |
| Junos XML API—XML representation of Junos OS configuration statements and operational mode commands. Junos XML configuration tag elements are the content to which the Junos XML protocol operations apply. Junos XML operational tag elements are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device. The Junos XML API also includes tag elements that are the counterpart to Junos CLI configuration statements. | <ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>                                           | <ul style="list-style-type: none"> <li>• <a href="#">Junos XML API Configuration Developer Reference</a></li> <li>• <a href="#">Junos XML API Operational Developer Reference</a></li> </ul>                                                     |

Table 582: Device and Network Management Features on the QFX Series (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Typical Uses                                                                                                                               | Documentation                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NETCONF XML management protocol—XML-based management protocol that client applications use to request and change configuration information on routing, switching, and security platforms running Junos OS. The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations. | <ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul> | <i>NETCONF XML Management Protocol Guide</i>                                                                                                                                                                                                                                                                        |
| Operational mode commands—May be used to do the following: <ul style="list-style-type: none"> <li>• Monitor switch performance. For example, the <b>show chassis routing-engine</b> command shows the CPU utilization of the Routing Engine. High CPU utilization of the Routing Engine can affect performance of the switch.</li> <li>• View current activity and status of the device or network. For example, you can use the <b>ping</b> command to monitor and diagnose connectivity problems, and the <b>traceroute</b> command to locate points of failure on the network.</li> </ul>                          | <ul style="list-style-type: none"> <li>• Performance management</li> <li>• Fault management</li> </ul>                                     | <a href="#">CLI Explorer</a>                                                                                                                                                                                                                                                                                        |
| Out-of-band management—Enables connection to the switch through a management interface. Out-of-band management is supported on two dedicated management Ethernet interfaces as well as on the console and auxiliary ports. The management Ethernet interfaces connect directly to the Routing Engine. No transit traffic is allowed through the interfaces, separating customer and management traffic and ensuring that congestion or failures in the transit network do not affect the management of the switch.                                                                                                    | Remote access management                                                                                                                   | <ul style="list-style-type: none"> <li>• <a href="#">Connecting a QFX3500 Device to a Network for Out-of-Band Management</a></li> <li>• <a href="#">Connecting a QFX Series Device to a Management Console</a></li> <li>• <a href="#">Configuring Console and Auxiliary Port Properties on page 6177</a></li> </ul> |
| SNMP Configuration Management MIB—Provides notification for configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made. A history of the last 32 configuration changes is kept in jnxCmChgEventTable.                                                                                                                                                                                                                                                        | Configuration management                                                                                                                   | <i>SNMP MIBs and Traps Reference</i>                                                                                                                                                                                                                                                                                |

Table 582: Device and Network Management Features on the QFX Series (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                               | Typical Uses                                                                                           | Documentation                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>SNMP MIBs and traps—Enable the monitoring of network devices from a central location. Use SNMP requests such as <b>get</b> and <b>walk</b> to monitor and view system activity.</p> <p>The QFX3500 switch supports SNMP Version 1 (v1), v2, and v3, and both standard and Juniper Networks enterprise-specific MIBs and traps.</p> | Fault management                                                                                       | <ul style="list-style-type: none"> <li>• <i>SNMP MIBs and Traps Reference</i></li> <li>• <a href="#">Understanding the Implementation of SNMP on page 6107</a></li> </ul>                                                                                                          |
| System log messages—Log details of system and user events, including errors. You can specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.                                                                                                     | <ul style="list-style-type: none"> <li>• Fault management</li> <li>• User access management</li> </ul> | <ul style="list-style-type: none"> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">Overview of Junos OS System Log Messages on page 6154</a></li> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6155</a></li> </ul> |

## Understanding Network Management Implementation on the QFabric System

This topic describes network management features on the QFabric system that are implemented differently than on other devices running Junos OS.

The following network management features are supported on the QFabric system:

- **System log messages**—The QFabric system monitors events that occur on its component devices, distributes system log messages about those events to all external system log message servers (hosts) that are configured, and archives the messages. Component devices include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. You configure system log messages at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view messages.
- **Simple Network Management Protocol (SNMP) Version 1 (v1) and v2c**—SNMP monitors network devices from a central location. The SNMP implementation on the QFabric system supports the basic SNMP architecture of Junos OS with some limitations, including a reduced set of MIB objects, read-only access for SNMP communities, and limited support for SNMP requests. You configure SNMP at the **[edit snmp]** hierarchy level. Only the **show snmp statistics** operational mode command is supported, but you can issue SNMP requests using external SNMP client applications.
- **Advanced Insight Solutions (AIS)**—AIS provides tools and processes to automate the delivery of support services for the QFabric system. AIS components include Advanced Insight Scripts (AI-Scripts) and Advanced Insight Manager (AIM). You install AI-Scripts using the **request system scripts add** operational mode command. However, the **jais-activate-scripts.slax** file used during installation is preconfigured for the QFabric system and cannot be changed.



**NOTE:** Do not install Junos Space and AIS on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

#### Related Documentation

- [Advanced Insight Scripts \(AI-Scripts\) Release Notes](#)
- [Understanding Device and Network Management Features on page 6077](#)
- [Overview of Junos OS System Log Messages on page 6154](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
- [SNMP MIBs Support on page 6124](#)

## Understanding Telnet on the QFabric System

This topic describes the support for the Telnet protocol on QFabric systems.

Telnet service is available for devices running Junos OS, including QFX Series devices. However, on QFabric systems, Telnet support is limited and the following conditions apply:

- You can telnet from a QFabric system to external devices that are connected to the QFabric system by way of the network Node group. To connect to these external devices, issue the **telnet** command from the QFabric default partition CLI.
- You cannot use the Telnet protocol to connect from the QFabric system default partition CLI to individual components. To access system components, you must issue the **request component login** command instead.

#### Related Documentation

- [request component login on page 1444](#)
- [telnet](#)

## Understanding Tracing and Logging Operations

Tracing and logging operations enable you to track events that occur in the switch—both normal operations and error conditions—and to track the packets that are generated by or passed through the switch. The results of tracing and logging operations are placed in files in the **/var/log** directory on the switch.

The Junos OS supports remote tracing for the following processes:

- **chassisd**—Chassis-control process
- **eventd**—Event-processing process
- **cosd**—Class-of-service process

You configure remote tracing by using the **tracing** statement at the **[edit system]** hierarchy level.



**NOTE:** The **tracing** statement is not supported on the QFX3000 QFabric system.

If you enabled remote tracing but wish to disable it for specific processes on the switch, use the **no-remote-trace** statement at the **[edit process-name traceoptions]** hierarchy level. This feature does not alter local tracing functionality in any way, and logging files are stored on the switch.

Logging operations use a system logging mechanism similar to the UNIX **syslogd** utility to record systemwide, high-level operations, such as interfaces going up or down and users logging in to or out of the switch. You configure these operations by using the **syslog** statement at the **[edit system]** hierarchy level and by using the **options** statement at the **[edit ethernet-switching-options]** hierarchy level.

Tracing operations record more detailed information about the operations of the switch, including packet forwarding and routing information. To configure tracing operations, use the **traceoptions** statement.



**NOTE:** The **traceoptions** statement is not supported on the QFX3000 QFabric system.

You can define tracing operations in different portions of the switch configuration:

- **SNMP agent activity tracing operations**—Define tracing of the activities of SNMP agents on the switch. You configure SNMP agent activity tracing operations at the **[edit snmp]** hierarchy level.
- **Global switching tracing operations**—Define tracing for all switching operations. You configure global switching tracing operations at the **[edit ethernet-switching-options]** hierarchy level of the configuration.
- **Protocol-specific tracing operations**—Define tracing for a specific routing protocol. You configure protocol-specific tracing operations in the **[edit protocols]** hierarchy when configuring the individual routing protocol. Protocol-specific tracing operations override any equivalent operations that you specify in the global **traceoptions** statement. If there are no equivalent operations, they supplement the global tracing options. If you do not specify any protocol-specific tracing, the routing protocol inherits all the global tracing operations.
- **Tracing operations within individual routing protocol entities**—Some protocols allow you to define more granular tracing operations. For example, in Border Gateway Protocol (BGP), you can configure peer-specific tracing operations. These operations override any equivalent BGP-wide operations or, if there are no equivalents, supplement them. If you do not specify any peer-specific tracing operations, the peers inherit, first, all the BGP-wide tracing operations and, second, the global tracing operations.
- **Interface tracing operations**—Define tracing for individual interfaces and for the interface process itself. You define interface tracing operations at the **[edit interfaces]** hierarchy level of the configuration.
- **Remote tracing**—To enable system-wide remote tracing, configure the **destination-override syslog host** statement at the **[edit system tracing]** hierarchy level. This specifies the remote host running the system log process (syslogd), which collects

the traces. Traces are written to files on the remote host in accordance with the syslogd configuration in `/etc/syslog.conf`. By default, remote tracing is not configured.

To override the system-wide remote tracing configuration for a particular process, include the `no-remote-trace` statement at the `[edit process-name traceoptions]` hierarchy. When `no-remote-trace` is enabled, the process does local tracing.

To collect traces, use the `local0` facility as the selector in the `/etc/syslog.conf` file on the remote host. To separate traces from various processes into different files, include the process name or trace-file name (if it is specified at the `[edit process-name traceoptions file]` hierarchy level) in the Program field in the `/etc/syslog.conf` file. If your system log server supports parsing hostname and program name, then you can separate traces from the various processes.



**NOTE:** During a commit check, warnings about the `traceoptions` configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

#### Related Documentation

- [Overview of Junos OS System Log Messages on page 6154](#)

## Automation Scripts

- [Understanding Automation Scripts Support on page 6083](#)
- [How Commit Scripts Work on page 6084](#)
- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 6089](#)
- [Overview of Generating Persistent or Transient Configuration Changes on page 6090](#)
- [Required Boilerplate for Commit Scripts on page 6095](#)
- [How Op Scripts Work on page 6096](#)
- [Required Boilerplate for Op Scripts on page 6097](#)

## Understanding Automation Scripts Support

This document describes the support for the Junos OS automation scripts on the QFabric system Director devices.

Junos OS automation consists of a suite of tools used to automate operational and configuration tasks on network devices running Junos OS. The automation tools, which leverage the native XML capabilities of the Junos OS, include commit scripts, operation (op) scripts, event policies and event scripts, and macros.



**NOTE:** Event policies and event scripts are not supported on the QFabric system at this time.

The QFabric system supports Junos OS automation scripts that are written in Stylesheet Language Alternative Syntax (SLAX) version 1.0.

Commit scripts automate the commit process and enforce custom configuration rules. You can use commit scripts to generate specific errors and warnings, and customize configurations and configuration templates. When a candidate configuration is committed, it is inspected by each active commit script. If a configuration violates your custom rules and the scripts generate an error, the commit fails. If the commit is successful, any configuration changes (both transient and permanent) are incorporated into the active configuration before it is passed to the Director software, which distributes the configuration to all applicable QFabric system components, including Node devices and Node servers.

Op scripts automate operational and troubleshooting tasks. Op scripts can be executed manually from the Junos OS CLI or NETCONF XML management protocol, or they can be called from another script.

The QFabric system supports the following automation script features:

- Commit scripts and op scripts are supported.
- Scripts written in SLAX version 1 are supported.
- Scripts are configured and deployed from the Director group. Since there is more than one Director device in a Director group, scripts must be deployed by each Director device or deployed in the shared media space.
- Scripts are stored in the shared media at this location:  
`/pbdata/mgd_shared/partition-ip/var/db/scripts`. Under this directory, commit scripts are stored in the **commit** subdirectory, and op scripts are stored in the **op** subdirectory.
- Scripts are not stored in flash memory.

#### Related Documentation

- [How Commit Scripts Work on page 6084](#)
- [How Op Scripts Work on page 6096](#)
- [Required Boilerplate for Commit Scripts on page 6095](#)
- [Required Boilerplate for Op Scripts on page 6097](#)
- [Controlling the Execution of Commit Scripts on page 6181](#)

## How Commit Scripts Work

You enable commit scripts by listing the names of one or more commit script files at the **[edit system scripts commit]** hierarchy level. These scripts contain instructions that enforce custom configuration rules. Commit scripts are invoked during the commit process before the standard Junos OS validity checks are performed.

When you perform a commit operation, Junos OS executes each script in turn, passing the information in the candidate configuration to the scripts. The script inspects the configuration, performs the necessary tests and validations, and generates a set of instructions for performing certain actions. These actions include generating error, warning,




and system log messages. If errors are generated, the commit operation fails and the candidate configuration remains unchanged. This is the same behavior that occurs with standard commit errors.

Commit scripts can also generate changes to the system configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 218 on page 6085 shows the flow of commit script input and output.

**Figure 218: Commit Script Input and Output**



Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

The following sections discuss several important concepts related to the commit script input and output:

- [Commit Script Input on page 6085](#)
- [Commit Script Output on page 6086](#)
- [Commit Scripts and the Junos OS Commit Model on page 6087](#)

### Commit Script Input

The input for a commit script is the postinheritance candidate configuration in Junos XML API format. The term *postinheritance* means that all configuration group values have been inherited by their targets in the candidate configuration and the inactive portions of the configuration have been removed. For more information about configuration groups, see the *CLI User Guide*.

When you issue the **commit** command, Junos OS automatically generates the candidate configuration in XML format and reads it into the management (mgd) process, at which time the input is evaluated by any commit scripts.

To display the XML format of the postinheritance configuration, issue the **show | display commit-scripts view** command:

```
[edit]
user@host# show | display commit-scripts view
```

To display all configuration groups data, including script-generated changes to the groups, issue the **show groups | display commit-scripts** command:

```
[edit]
user@host# show groups | display commit-scripts
```

To save the commit script input to a file, add the **save** command to the command line:

```
[edit]
user@host# show | display commit-scripts view | save filename.xml
```

By default, the file is placed in your home directory on the switch, router, or security device.

---

### Commit Script Output

To specify the desired commit script output—including warning, error, and system log messages, persistent changes, and transient changes—the script can contain tags that appear in any order, in any number. The tags for specifying output are as follows:

- **<xnm:warning>**—Generates a warning message
- **<xnm:error>**—Generates an error message.
- **<syslog><message>**—Generates a system log message.
- **<change>**—Generates a persistent change to the configuration.
- **<transient-change>**—Generates a transient change to the configuration.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="content">**—Generates a persistent change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="tag" select="'transient-change'"/>**
    - <xsl:with-param name="content">**—Generates a transient change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="message">**
    - <xsl:text>**—Generates a warning message in conjunction with a configuration change. You can use this set of tags to generate a notification that the configuration has been changed.

Junos OS processes this output and performs the appropriate actions. Errors and warnings are passed back to the Junos OS CLI or to a Junos XML protocol client application. The presence of an error automatically causes the commit operation to fail. Persistent and transient changes are loaded into the appropriate configuration database.

To test the output of error, warning, and system log messages from commit scripts, issue the **commit check | display xml** command:

```
[edit]
user@host# commit check | display xml
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command:

```
[edit]
user@host# commit check | display detail
```



**NOTE:** System log messages do not appear in the trace output, so you cannot use the commit check operation to test script-generated system log messages. Furthermore, system log messages are written to the system log during a commit operation, but not during a commit check operation.

#### Related Documentation

- *Example: Protecting the Junos OS Configuration from Modification or Deletion.*
- *jcs:emit-change Template*

### Commit Scripts and the Junos OS Commit Model

Junos OS uses a commit model to update the device's configuration. This model allows you to make a series of changes to a candidate configuration without affecting the operation of the device. When the changes are complete, you can commit the configuration. The commit operation saves the candidate configuration changes into the current configuration.

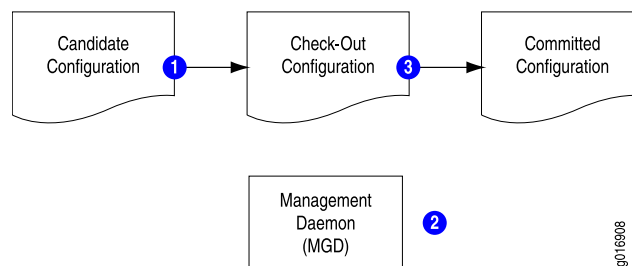
When you commit a set of changes in the candidate configuration, two methods are used to forward these changes to the current configuration:

- Standard commit model—Used when no commit scripts are active on the device.
- Commit script model—Incorporates commit scripts into the commit model.

#### Standard Commit Model

In the standard commit model, the management (mgd) process validates the candidate configuration based on standard Junos validation rules. If the configuration file is valid, it becomes the current active configuration. [Figure 219 on page 6087](#) and the accompanying discussion explain how the standard commit model works:

**Figure 219: Standard Commit Model**



In the standard commit model, the software performs the following steps:

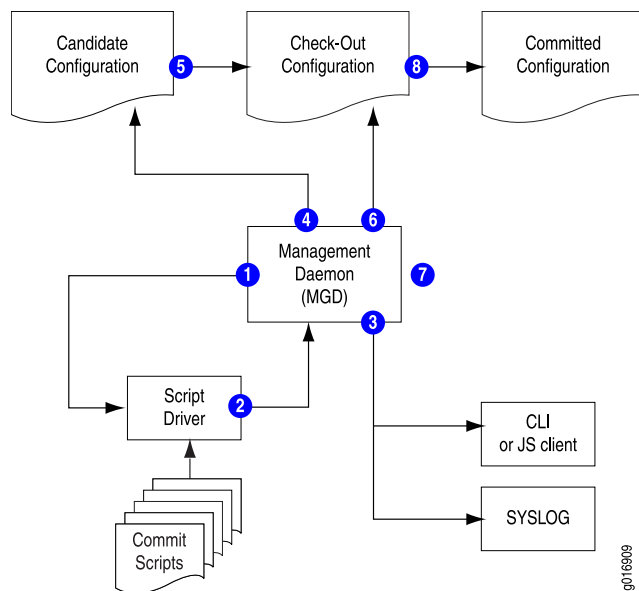
1. When the candidate configuration is committed, it is copied to become the checkout configuration.

2. The mgd process validates the checkout configuration.
3. If no error occurs, the checkout configuration is copied as the current active configuration.

### Commit Model with Commit Scripts

When commit scripts are added to the standard commit model, the process becomes more complex. The mgd process first passes an XML-formatted checkout configuration to a script driver, which handles the verification of the checkout configuration by the commit scripts. When verification is complete, the script driver returns an XML *action file* to the mgd process. The mgd process follows the instructions in the action file to update the candidate and checkout configurations, issue messages to the CLI, and write information to the system log as required. After processing the action file, the mgd process performs the standard Junos OS validation. [Figure 220 on page 6088](#) and the accompanying discussion explain this process.

**Figure 220: Commit Model with Commit Scripts Added**



In the commit script model, Junos OS performs the following steps:

1. When the candidate configuration is committed, the mgd process sends the XML-formatted candidate configuration to the script driver.
2. Each enabled commit script is invoked against the candidate configuration, and each script can generate a set of actions for the mgd process to perform. The actions are collected in an XML action file.
3. The mgd process performs the following actions in response to **<error>**, **<warning>**, and **<syslog>** tag elements in the action file:
  - **<error>**—The mgd process halts the commit process (that is, the commit operation fails), returns an error message to the CLI or Junos XML protocol client, and takes no further action.

- **<warning>**—The mgd process forwards the message to the CLI or the Junos XML protocol client.
  - **<syslog>**—The mgd process forwards the message to the system log process.
4. If the action file includes any **<change>** tag elements, the mgd process loads the requested changes into the candidate configuration.
  5. The candidate configuration is copied to become the checkout configuration.
  6. If the action file includes any **<transient-change>** tag elements, the mgd process loads the requested changes into the checkout configuration.
  7. The mgd process validates the checkout configuration.
  8. If there are no validation errors, the checkout configuration is copied to become the current active configuration.



**NOTE:** Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

Changes that are made to the candidate configuration during the commit operation are not evaluated by the custom rules during that commit operation. However, persistent changes are maintained in the candidate configuration and are evaluated by the custom rules during subsequent commit operations. For more information about how commit scripts change the candidate configuration, see [“Avoiding Potential Conflicts When Using Multiple Commit Scripts” on page 6089](#).

Transient changes are never evaluated by the custom rules in commit scripts, because they are made to the checkout configuration only after the commit scripts have evaluated the candidate configuration and the candidate is copied to become the checkout configuration. To remove a transient change from the configuration, remove, disable, or deactivate the commit script (as discussed in *Controlling Execution of Commit Scripts During Commit Operations*), or comment out the code that generates the transient change.

For more information about differences between persistent and transient changes, see [“Overview of Generating Persistent or Transient Configuration Changes” on page 6090](#).

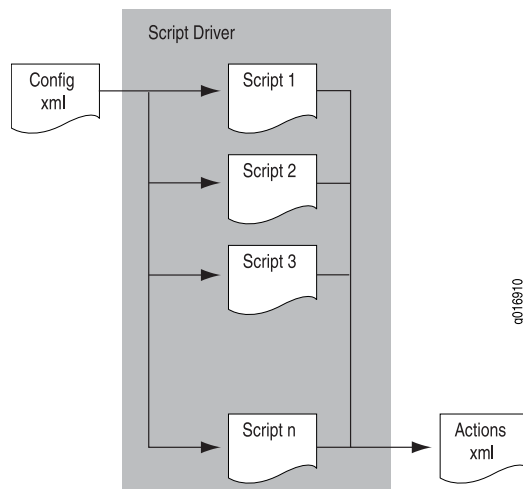
#### Related Documentation

- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 6089](#)

## Avoiding Potential Conflicts When Using Multiple Commit Scripts

When you use multiple commit scripts, each script evaluates the original candidate configuration file. Changes made by one script are not evaluated by the other scripts. This means that conflicts between scripts might not be resolved when the scripts are first applied to the configuration. The commit scripts are executed in the order they are listed at the **[edit system scripts commit]** hierarchy level, as illustrated in [Figure 221 on page 6090](#).

Figure 221: Configuration Evaluation by Multiple Commit Scripts



As an example of a conflict between commit scripts, suppose that commit script **A.xsl** is created to ensure that the device uses the domain name server with IP address 192.168.0.255. Later, the DNS server's address is changed to 192.168.255.255 and a second script, **B.xsl**, is added to check that the device uses the DNS server with that address. However, script **A.xsl** is not removed or disabled.

Because each commit script evaluates the original candidate configuration, the final result of executing both scripts **A.xsl** and **B.xsl** depends on which DNS server address is configured in the original candidate configuration. If the now outdated address of 192.168.0.255 is configured, script **B.xsl** changes it to 192.168.255.255. However, if the correct address of 192.168.255.255 is configured, script **A.xsl** changes it to the incorrect value 192.168.0.255.

As another example of a potential conflict between commit scripts, suppose that a commit script protects a hierarchy using the **protect** attribute. If a second commit script attempts to modify or delete the hierarchy or the statements within the hierarchy, Junos OS issues a warning during the commit process and prevents the configuration change.

Exercise care to ensure that you do not introduce conflicts between scripts like those described in the examples. As a method of checking for conflicts with persistent changes, you can issue two separate **commit** commands.

**Related Documentation**

- [How Commit Scripts Work on page 6084](#)

## Overview of Generating Persistent or Transient Configuration Changes

Junos OS commit scripts enforce custom configuration rules. When a candidate configuration includes statements that you have decided must not be included in your configuration, or when the candidate configuration omits statements that you have

decided are required, commit scripts can automatically change the configuration and thereby correct the problem.

- [Differences Between Persistent and Transient Changes on page 6091](#)
- [Interaction of Configuration Changes and Configuration Groups on page 6094](#)
- [Tag Elements and Templates for Generating Changes on page 6094](#)

### Differences Between Persistent and Transient Changes

---

Configuration changes made by commit scripts can be *persistent* or *transient*.

A persistent change remains in the candidate configuration and affects routing operations until you explicitly delete it, even if you subsequently remove or disable the commit script that generated the change and reissue the **commit** command. In other words, removing the commit script does not cause a persistent change to be removed from the configuration.

A transient change, in contrast, is made in the *checkout configuration* but not in the candidate configuration. The checkout configuration is the configuration database that is inspected for standard Junos OS syntax just before it is copied to become the active configuration on the device. If you subsequently remove or disable the commit script that made the change and reissue the **commit** command, the change is no longer made to the checkout configuration and so does not affect the active configuration. In other words, removing the commit script effectively removes a transient change from the configuration.

A common use for transient changes is to eliminate the need to repeatedly configure and display well-known policies, thus allowing these policies to be enforced implicitly. For example, if MPLS must be enabled on every interface with an International Organization for Standardization (ISO) protocol enabled, the change can be transient, so that the repetitive or redundant configuration data need not be carried or displayed in the candidate configuration. Furthermore, transient changes allow you to write script instructions that apply the change only if a set of conditions is met.

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute to a configuration element produces the same behavior as a **replace:** tag in a **load replace** operation.

By default, Junos OS merges the incoming configuration and the candidate configuration. New statements and hierarchies are added, and conflicting statements are overridden. When generating a persistent or transient change, if you add the **replace="replace"** attribute to a configuration element, Junos OS replaces the existing configuration element with the incoming configuration element. If the **replace="replace"** attribute is added to a configuration element, but there is no existing element of the same name in the current configuration, the incoming configuration element is added into the configuration. Elements that do not have the **replace** attribute are merged into the configuration.

Persistent and transient changes are loaded before the standard Junos validation checks are performed. This means any configuration changes introduced by a commit script are

validated for correct syntax. If the syntax is correct, the new configuration becomes the active, operational device configuration.

Protected elements in the configuration hierarchy cannot be modified or deleted by either a persistent or a transient change. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made, and proceeds with the commit.

Persistent and transient changes have several important differences, as described in [Table 583 on page 6092](#).

**Table 583: Differences Between Persistent and Transient Changes**

| Persistent Changes                                                                                                                                                                                                                                                                                                                                           | Transient Changes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>A persistent change is represented in a commit script by the <b>&lt;change&gt;</b> tag.</p> <p>Another way to represent a persistent change is with the <b>content</b> parameter inside a call to the <b>jcs:emit-change</b> template.</p> <p>The <b>jcs:emit-change</b> template is a helper template contained in the <b>junos.xsl</b> import file.</p> | <p>A transient change is represented in a commit script by the <b>&lt;transient-change&gt;</b> tag.</p> <p>Another way to represent a transient change is to use the <b>content</b> parameter and the <b>tag transient</b> parameter inside a call to the <b>jcs:emit-change</b> template.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <p>You can use persistent changes to perform any Junos XML protocol operation, such as activate, deactivate, delete, insert (reorder), comment (annotate), and replace sections of the configuration.</p>                                                                                                                                                    | <p>Like persistent changes, you can use transient changes to perform any Junos XML protocol operation. However, some Junos XML protocol operations do not make sense to use with transient changes, such as generating comments and inactive settings.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <p>Persistent changes are always loaded during the commit process if no errors are generated by any commit scripts or by the standard Junos OS validity check.</p>                                                                                                                                                                                           | <p>For transient changes to be loaded, you must include the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level. If you enable a commit script that generates transient changes and you do not include the <b>allow-transients</b> statement in the configuration, the CLI generates an error message and the commit operation fails.</p> <p>Like persistent changes, transient changes must pass the standard Junos OS validity check.</p> <p>You cannot use a commit script to generate the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level. Rather, you must include this statement directly by using the CLI.</p> |



Table 583: Differences Between Persistent and Transient Changes (*continued*)

| Persistent Changes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Transient Changes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Persistent changes work like the <b>load replace</b> configuration mode command, and the change is added to the candidate configuration.</p> <p>When generating a persistent change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the candidate configuration with the incoming configuration element. If there is no existing element of the same name in the candidate configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p> | <p>Transient changes work like the <b>load replace</b> configuration mode command, and the change is added to the checkout configuration.</p> <p>When generating a transient change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the checkout configuration with the incoming configuration element. If there is no existing element of the same name in the checkout configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p> <p>Transient changes are not copied to the candidate configuration. For this reason, transient changes are not saved in the configuration if the associated commit script is deleted or deactivated.</p> |
| <p>After a persistent change is committed, the software treats it like a change you make by directly editing and committing the candidate configuration.</p> <p>After the persistent changes are copied to the candidate configuration, they are copied to the checkout configuration. If the changes pass the standard Junos OS validity checks, the changes are propagated to the switch, router, or security device components.</p>                                                                                                                                                                                                  | <p>Each time a transient change is committed, the software updates the checkout configuration database. After the transient changes pass the standard Junos OS validity checks, the changes are propagated to the device components.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <p>After committing a script that causes a persistent change to be generated, you can view the persistent change by issuing the <b>show</b> configuration mode command:</p> <pre>user@host# show</pre> <p>This command displays persistent changes only, not transient changes.</p>                                                                                                                                                                                                                                                                                                                                                     | <p>After committing a script that causes a transient change to be generated, you can view the transient change by issuing the <b>show   display commit-scripts</b> configuration mode command:</p> <pre>user@host# show   display commit-scripts</pre> <p>This command displays both persistent and transient changes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <p>Persistent changes must conform to your custom configuration design rules as dictated by commit scripts.</p> <p>This does not become apparent until after a second commit operation because persistent changes are not evaluated by commit script rules on the current commit operation. The subsequent commit operation fails if the persistent changes do not conform to the rules imposed by the commit scripts configured during the first commit operation.</p>                                                                                                                                                                 | <p>Transient changes are never tested by and do not need to conform to your custom rules. This is caused by the order of operations in the Junos OS commit model, which is explained in detail in <a href="#">“Commit Scripts and the Junos OS Commit Model” on page 6087</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <p>A persistent change remains in the configuration even if you delete, disable, or deactivate the commit script instructions that generated the change.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>If you delete, disable, or deactivate the commit script instructions that generate a transient change, the change is removed from the configuration after the next commit operation. In short, if the associated instructions or the entire commit script is removed, the transient change is also removed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 583: Differences Between Persistent and Transient Changes (*continued*)

| Persistent Changes                                                                                                                                                                                                                                                                                                                                                                                                                       | Transient Changes                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| As with direct CLI configuration, you can remove a persistent change by rolling back to a previous configuration that did not include the change and issuing the <b>commit</b> command. However, if you do not disable or deactivate the associated commit script, and the problem that originally caused the change to be generated still exists, the change is automatically regenerated when you issue another <b>commit</b> command. | You cannot remove a transient change by rolling back to a previous configuration.                                                                                                                                                                                                              |
| You can alter persistent changes directly by editing the configuration using the CLI.                                                                                                                                                                                                                                                                                                                                                    | <p>You cannot directly alter or delete a transient change by using the Junos OS CLI, because the change is not in the candidate configuration.</p> <p>To alter the contents of a transient change, you must alter the statements in the commit script that generates the transient change.</p> |

### Interaction of Configuration Changes and Configuration Groups

Any configuration change you can make by directly editing the configuration using the Junos OS command-line interface (CLI) can also be generated by a commit script as a persistent or transient change. This includes values specified at a specific hierarchy level or in configuration groups. As with direct CLI configuration, values specified in the *target* override values inherited from a configuration group. The target is the statement to which you apply a configuration group by including the **apply-groups** statement.

If you define persistent or transient changes as belonging to a configuration group, the configuration groups are applied in the order you specify in the **apply-groups** statements, which you can include at any hierarchy level except the top level. You can also disable inheritance of a configuration group by including the **apply-groups-except** statement at any hierarchy level except the top level.



**CAUTION:** Each commit script inspects the postinheritance view of the configuration. If a candidate configuration contains a configuration group, be careful when using a commit script to change the related target configuration, because doing so might alter the intended inheritance from the configuration group.

Also be careful when using a commit script to change a configuration group, because the configuration group might be generated by an application that performs a load replace operation on the group during each commit operation.

For more information about configuration groups, see the *CLI User Guide*.

### Tag Elements and Templates for Generating Changes

To generate changes, you can use the **jcs:emit-change** template, which implicitly includes **<change>** and **<transient-change>** XML elements; or you can explicitly include **<change>**

and `<transient-change>` XML elements. Using the `jcs:emit-change` template allows you to set the hierarchical context of the change once rather than multiple times.

The `<change>` and `<transient-change>` elements are similar to the `<load-configuration>` operation defined by the Junos XML management protocol. The possible contents of the `<change>` and `<transient-change>` elements are the same as the contents of the `<configuration>` tag element used in the Junos XML protocol operation `<load-configuration>`. For complete details about the `<load-configuration>` element, see the *Junos XML Management Protocol Developer Guide*.

## Required Boilerplate for Commit Scripts

When you write commit scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all commit scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make commit scripts easier to read and write, which you import from a file called `junos.xsl`. For more information about the extension functions and templates, see *Junos Script Automation: Understanding Extension Functions in the jcs and slax Namespaces* and *Junos Script Automation: Named Templates in the jcs Namespace Overview*.

Commit scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all commit scripts that you create. The XSLT boilerplate follows:

### XSLT Boilerplate for Commit Scripts

```

1 <?xml version="1.0" standalone="yes"?>
2 <xsl:stylesheet version="1.0"
3 xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4 xmlns:junos="http://xml.juniper.net/junos/*/junos"
5 xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6 xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7 <xsl:import href="../../../import/junos.xsl"/>

8 <xsl:template match="configuration">
9 <!-- ... Insert your code here ... -->
10 </xsl:template>
11 </xsl:stylesheet>

```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI). This PI specifies that the code is written in XML using version 1.0. The XML PI, if present, must be the first noncomment token in the script file.

```

1 <?xml version="1.0"?>

```

Lines 2 through 6 set the style sheet element and the associated namespaces. Line 2 sets the style sheet version as 1.0. Lines 3 through 6 list all the namespace mappings commonly used in commit scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```
2 <xsl:stylesheet version="1.0"
3 xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4 xmlns:junos="http://xml.juniper.net/junos/*/junos"
5 xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6 xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of the Junos OS. The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Junos Script Automation: Named Templates in the jcs Namespace Overview* and *Junos Named Templates in the jcs Namespace Summary*.

```
7 <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `<configuration>` element, which is the node selected by the `<xsl:template match="/">` template, contained in the `junos.xsl` import file. The `<xsl:template match="configuration">` element allows you to exclude the `/configuration/` root element from all XML Path Language (XPath) expressions in the script and begin XPath expressions with the top Junos OS hierarchy level. For more information, see *XPath Overview*.

```
8 <xsl:template match="configuration">
```

Add your code between Lines 8 and 9.

Line 9 closes the template.

```
9 </xsl:template>
```

Line 10 closes the style sheet and the commit script.

```
10 </xsl:stylesheet>
```

### SLAX Boilerplate for Commit Scripts

The corresponding SLAX boilerplate is as follows:

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
/*
* Insert your code here
*/
}
```

## How Op Scripts Work


Op scripts execute Junos OS operational commands and inspect the resulting output. After inspection, op scripts can automatically correct errors within the device running Junos OS based on this output.

You add op scripts to device operations by listing the filenames of one or more op script files within the **[edit system scripts op]** hierarchy level. These files must be added to the appropriate op script file directory. For more information about op script file directories, see *Storing Scripts in Flash Memory*. Once added to the device, op scripts are invoked from the command line, using the **op filename** command.

You can use op scripts to generate changes to the device configuration by including the **<load-configuration>** tag element. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 222 on page 6097 shows a high-level view of the flow of op script input and output.

**Figure 222: Op Script Input and Output**



## Required Boilerplate for Op Scripts

When you write operation (op) scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all op scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make scripts easier to read and write, which you import from a file called **junos.xml**. For more information about the extension functions and templates, see *Junos Script Automation: Understanding Extension Functions in the jcs and slax Namespaces* and *Junos Script Automation: Named Templates in the jcs Namespace Overview*.

Op scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all op scripts that you create. The XSLT boilerplate follows:

### XSLT Boilerplate for Op Scripts

```

1 <?xml version="1.0" standalone="yes"?>
2 <xsl:stylesheet version="1.0"
3 xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4 xmlns:junos="http://xml.juniper.net/junos/*/junos"
5 xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6 xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7 <xsl:import href="../import/junos.xml"/>
8 <xsl:template match="/">

```

```
9 <op-script-results>
 <!-- ... insert your code here ... -->
10 </op-script-results>
11 </xsl:template>
 <!-- ... insert additional template definitions here ... -->
12 </xsl:stylesheet>
```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI), which marks this file as XML and specifies the version of XML as 1.0. The XML PI, if present, must be the first non-comment token in the script file.

```
1 <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2 <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in operation scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```
3 xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4 xmlns:junos="http://xml.juniper.net/junos/*/junos"
5 xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6 xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of Junos OS (in the file `/usr/libdata/cscript/import/junos.xsl`). The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Junos Script Automation: Named Templates in the jcs Namespace Overview* and *Junos Named Templates in the jcs Namespace Summary*.

```
7 <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `</>` element. The `<xsl:template match="/">` element is the root element and represents the top level of the XML hierarchy. All XML Path Language (XPath) expressions in the script must start at the top level. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see *XPath Overview*.

```
8 <xsl:template match="/">
```

After the `<xsl:template match="/">` tag element, the `<op-script-results>` and `</op-script-results>` container tags must be the top-level child tags, as shown in Lines 9 and 10.

```
9 <op-script-results>
 <!-- ... insert your code here ... -->
10 </op-script-results>
```

Line 11 closes the template.

```
11 </xsl:template>
```

Between Line 11 and Line 12, you can define additional XSLT templates that are called from within the `<xsl:template match="/">` template.

Line 12 closes the style sheet and the op script.

```
12 </xsl:stylesheet>
```

#### SLAX Boilerplate for Op Scripts

The corresponding SLAX boilerplate is as follows:

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match / {
 <op-script-results> {
 /*
 * Insert your code here
 */
 }
}
```

## Fabric OAM

- [Overview of Internal Fabric Monitoring on page 6099](#)

### Overview of Internal Fabric Monitoring

Internal fabric monitoring is a feature of the Operation, Administration, and Maintenance (OAM) of the QFabric system. This feature enables you to validate the flow path of protocol data units (PDUs) across a given VLAN on the QFabric system using the unicast ping, multicast ping, and traceroute operations.

Internal fabric monitoring is useful for fault detection on the QFabric system. For example, if a PDU reaches a destination that is not part of the VLAN configuration, the operation (unicast ping, multicast ping, or traceroute) displays the exception on the console at runtime.

The unicast and multicast ping operations send PDUs from a source interface (called the source fabric maintenance endpoint [FMEP]) to a destination FMEP. The destination FMEP sends a response to the source FMEP when the PDUs are received.

The traceroute operation, also called a flow linktrace, traces the path taken by a specific, learned unicast flow from a source FMEP to a destination FMEP in a VLAN. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices. The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the PDU travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.

The following internal fabric monitoring commands are supported:

- **show oam fabric flow specification**
- **show oam fabric interfaces**
- **ping fabric unicast-flow**
- **ping fabric multicast-flow**
- **traceroute fabric unicast-flow**

The following are the main components of the internal fabric monitoring feature:

- **FMEP**—Represents the source or destination point (endpoint) in the monitoring operations. An FMEP is an interface through which PDUs are sent (source FMEP) or received (destination FMEP). Upon receipt of the PDUs, the destination FMEP sends a response to the source FMEP to validate the monitoring flow. FMEPs are associated with a VLAN in the fabric maintenance association (FMA) configuration, and the source and destination FMEP addresses are configured in the flow specification.
- **FMA**—Associates a set of FMEPs with a VLAN. The FMA defines the VLAN and FMEP parameters, including the VLAN name, FMEP identifiers, FMEP names, and the interface names of the FMEPs. The FMEPs defined in the FMA are the source and destination FMEPs used in the monitoring commands.



**NOTE:** A default FMA is automatically created for each Node group in the QFabric system. The default FMA is used to send error response PDUs (for example, in the case of a VLAN leak) and responses to PDUs that are not mapped to a specific interface in the QFabric system.

---

- **Flow specification**—Configures the flow type and FMEP addressing parameters. Unicast flow types include an Ethernet type (other than IPv4) and Ethernet IPv4. Multicast flow types include the Ethernet IPv4 and VLAN flood type. The flow specification also defines parameters within each flow type, such as MAC or IPv4 addresses of the source and destination FMEPs. The names and identifiers of these FMEPs are configured in the FMA parameters.

To enable internal fabric monitoring, configure the **fabric-maintenance-associations** and **flow-specs** statements at the **[edit protocols oam fabric]** hierarchy level.

The **ping fabric unicast-flow**, **ping fabric multicast-flow**, and **traceroute fabric unicast-flow** commands require that you specify the flow specification and FMA names, as well as the source and destination FMEP names.

#### Related Documentation

- [Configuring a Fabric Maintenance Association on page 6184](#)
- [Configuring Flow Specifications on page 6185](#)
- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)



## Junos Space

- [Understanding Junos Space Support on page 6101](#)
- [Preparing the Device for Junos Space Management on page 6104](#)

### Understanding Junos Space Support

The Juniper Networks Junos Space application, running on a JA1500 appliance or a Junos Space Virtual Appliance, is a comprehensive platform for building and deploying applications for collaboration, productivity, and network infrastructure and operations management. Junos Space provides a runtime environment implemented as a fabric of virtual and physical appliances.

The Junos Space software comprises various applications for network management and configuration. Starting with Junos Space Release 11.1, the following Junos Space applications provide support for the QFX Series devices:

- Network Application Platform (Platform) automates network management operations, including device discovery and management, device templates for configuration, scripts management, image management, and SSH secure console.
- Ethernet Design provides tools to simultaneously configure and manage multiple Junos OS devices within a network. You can manage and customize port profiles, and apply port profiles to your devices.
- Service Now enables you to display and manage information about problem events. When problems are detected on the device by Advanced Insight Scripts (AI-Scripts) that are installed on the device, the data is collected and sent to Service Now for your review and action. This application is available only if you have a Juniper Networks service contract.



**NOTE:** Do not install Junos Space and AIS on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

- Service Insight provides proactive support in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their end of life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now. This application is available only if you have a Juniper Networks service contract.
- Virtual Control is a web-based solution that enables you to manage virtual networks deployed in virtualized environments in data centers. It provides a single management interface for you to monitor and control the virtual and physical elements of the virtual environment.
- IBM Systems Director and Junos Space Launch in Context (LiC) integrate the IBM Systems Director and Junos Space software to manage Juniper Networks devices.

For more information about Junos Space, go to:

[http://www.juniper.net/techpubs/en\\_US/release-independent/junos-space/](http://www.juniper.net/techpubs/en_US/release-independent/junos-space/)



**NOTE:** Before you can perform the tasks listed in [Table 584 on page 6102](#), you must first ensure the configuration on the device meets the requirements for device discovery in Junos Space. For more information, see [“Preparing the Device for Junos Space Management” on page 6104](#).

[Table 584 on page 6102](#) lists some device and network management tasks you can perform in Junos Space.

**Table 584: Device and Network Management Tasks in Junos Space**

| Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Application Workspace Task                                                      | Documentation                                                                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set up and manage user access                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>Platform &gt; Administration &gt; Users &gt; Create User</b>                 | <i>Junos Space Network Application Platform User Guide</i>                                                                                                                                             |
| Bring the device under Junos Space management (device discovery)<br><br><b>NOTE:</b> Before the device can be discovered in Junos Space, the following requirements must be met: <ul style="list-style-type: none"> <li>• The device configuration has a static management IP address that is reachable from the Junos Space server.</li> <li>• There is a user with full administrative privileges for Junos Space administration.</li> <li>• If you plan on using SNMP to probe devices as part of device discovery, SNMP is enabled with appropriate read-only v1/v2c/v3 credentials.</li> </ul> | <b>Platform &gt; Devices &gt; Discover Devices</b>                              | <ul style="list-style-type: none"> <li>• <a href="#">Preparing the Device for Junos Space Management on page 6104</a></li> <li>• <i>Junos Space Network Application Platform User Guide</i></li> </ul> |
| View the hardware inventory or status of physical interfaces on the device                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>Platform &gt; Devices &gt; Manage Devices</b>                                | <i>Junos Space Network Application Platform User Guide</i>                                                                                                                                             |
| Upload Junos OS software image, deploy the image on the device, and manage device images                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>Platform &gt; Device Images &gt; Manage Images</b>                           |                                                                                                                                                                                                        |
| Edit the configuration of a managed device                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>Platform &gt; Devices &gt; Manage Devices &gt; Edit Device Configuration</b> | <i>Junos Space Network Application Platform User Guide</i>                                                                                                                                             |

Table 584: Device and Network Management Tasks in Junos Space (*continued*)

| Task                                                                                          | Application Workspace Task                                                                                               | Documentation                                                       |
|-----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| Design a device template definition for configuring the device                                | <b>Platform &gt; Device Templates &gt; Manage Definitions</b>                                                            | <i>Junos Space Network Application Platform User Guide</i>          |
| Centrally configure one or more devices by deploying a device template                        | <b>Platform &gt; Device Templates &gt; Manage Templates</b>                                                              | <i>Junos Space Network Application Platform User Guide</i>          |
| Manage device management interface (DMI) schemas                                              | <b>Platform &gt; Administration &gt; Manage DMI Schemas</b>                                                              | <i>Junos Space Network Application Platform User Guide</i>          |
| Import, modify, and deploy Junos OS scripts                                                   | <b>Platform &gt; Scripts &gt; Manage Scripts</b>                                                                         | <i>Junos Space Network Application Platform User Guide</i>          |
| Access the device remotely using a secure console                                             | <b>Platform &gt; Devices &gt; Secure Console</b>                                                                         | <i>Junos Space Network Application Platform User Guide</i>          |
| Create and manage QFabric system Node groups                                                  | <b>Ethernet Design &gt; EZ Design &gt; Manage QFabric Node Groups</b>                                                    | <i>Ethernet Design User Guide</i>                                   |
| Provision a port or port group using a port profile                                           | <b>Ethernet Design &gt; EZ Design &gt; Port Profile or Ethernet Design &gt; EZ Design &gt; Manage Fabric Port Groups</b> | <i>Ethernet Design User Guide</i>                                   |
| View incident reports and submit a case to Juniper Networks support                           | <b>Service Now &gt; Service Central &gt; Incidents</b>                                                                   | <i>Service Automation User Guide</i>                                |
| Get notification from Juniper support about known issues (Proactive Bug Notifications [PBNs]) | <b>Service Insight &gt; Insight Central &gt; Targeted PBNs</b>                                                           | <i>Service Automation User Guide</i>                                |
| Use the Virtual Control Getting Started Assistant to set up a virtual network                 | Navigate to the Virtual Control application, click the Help icon, and expand the Getting Started assistant               | <i>Virtual Control User Guide</i>                                   |
| Use IBM Systems Director to launch Junos Space and manage devices                             | <b>IBM Systems Director &gt; Navigate Resources</b>                                                                      | <i>IBM Systems Director and Junos Space Launch in Context (LIC)</i> |

**Related Documentation**

- [Installing and Connecting a QFX3500 Device](#)
- [Preparing the Device for Junos Space Management on page 6104](#)

## Preparing the Device for Junos Space Management

Before you can use the Juniper Networks Junos Space application to manage the QFX Series device, you must ensure that the configuration on the device meets the following requirements for device discovery in Junos Space:

- The device configuration has a static management IP address that is reachable from the Junos Space server.
- There is a user with full administrative privileges for Junos Space administration.
- SNMP is enabled (only if you plan on using SNMP as part of the device discovery).
- In Junos Space, set up a default device management interface (DMI) schema for the QFX Series device.



**NOTE:** Do not install Junos Space and AI-Scripts (AIS) on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

To prepare the device before using Junos Space:

1. Perform the initial configuration of the device through the console port using the Junos OS CLI. This task includes the configuration of a static management IP address and a user with root administrative privileges.

For the QFX3500 switch, see [“Configuring a QFX3500 Device as a Standalone Switch” on page 163](#).

For the QFabric system, see [“QFabric System Initial and Default Configuration Information” on page 1255](#) and [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 1335](#).

2. (Optional) Configure SNMP if you plan on using SNMP to probe devices during device discovery.

See [“Configuring SNMP” on page 1703](#).

3. (Optional) Enable SSH if you wish to use the Secure Console feature in Junos Space.

See [“Configuring SSH Service for Remote Access to the Router or Switch” on page 1709](#).

4. In Junos Space, set up a default DMI schema by navigating to **Network Application Platform > Administration > Manage DMI Schemas**.

For more information, see “Managing DMI Schemas” in the *Junos Space Network Application Platform User Guide* at:

[http://www.juniper.net/techpubs/en\\_US/release-independent/junos-space/](http://www.juniper.net/techpubs/en_US/release-independent/junos-space/) .

---

## sFlow Technology

- [Overview of sFlow Technology on page 6105](#)

## Overview of sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a QFX Series device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

sFlow technology implements the following two sampling mechanisms:

- **Packet-based sampling**—Samples one packet out of a specified number of packets from an interface enabled for sFlow technology. Only the first 128 bytes of each packet are sent to the collector. Data collected include the Ethernet, IP, and TCP headers, along with other application-level headers (if present). Although this type of sampling might not capture infrequent packet flows, the majority of flows are reported over time, allowing the collector to generate a reasonably accurate representation of network activity. You configure packet-based sampling when you specify a sample rate.
- **Time-based sampling**—Samples interface statistics (counters) at a specified interval from an interface enabled for sFlow technology. Statistics such as Ethernet interface errors are captured. You configure time-based sampling when you specify a polling interval.

An sFlow monitoring system consists of an sFlow agent embedded in the device and up to four external collectors. On a QFX3500 or QFX3600 standalone switch, the sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to the sFlow collectors.

On a QFabric system, the sFlow technology architecture is distributed. The global sFlow technology configuration defined on the QFabric system Director device is distributed to Node groups that have sFlow sampling configured on their interfaces. The sFlow agent has a separate sampling entity, known as a *subagent*, running on each Node device. Each subagent has its own independent state and forwards its own sample information (datagrams) directly to the sFlow collectors.

On the QFabric system, an sFlow collector must be reachable through the data network. Because each Node device has all routes stored in the default routing instance, the collector IP address should be included in the default routing instance to ensure the collector's reachability from the Node device.

Regardless of the rate of traffic or the configured sampling interval, a datagram is sent whenever its size reaches the maximum Ethernet transmission unit (MTU) of 1500 bytes, or whenever a 250-ms timer expires, whichever occurs first. The timer ensures that a collector receives regularly sampled data.

To ensure sampling accuracy and efficiency, QFX Series devices use adaptive sFlow sampling. Adaptive sampling monitors the overall incoming traffic rate on the device and

provides feedback to the interfaces to dynamically adapt their sampling rate to traffic conditions. The sFlow agent reads the statistics on the interfaces every 5 seconds and identifies five interfaces with the highest number of samples. On a standalone switch, when the CPU processing limit is reached, a binary backoff algorithm is implemented to reduce the sampling load of the top five interfaces by half. The adapted sampling rate is then applied to other interfaces.

On a QFabric system, sFlow technology monitors the interfaces on each Node device as a group, and implements the binary backoff algorithm based on the traffic on that group of interfaces.

Using adaptive sampling prevents overloading of the CPU and keeps the device operating at its optimum level even when there is a change in traffic patterns on the interfaces. The reduced sampling rate is used until the device is rebooted or when a new sampling rate is configured.

The sFlow collector uses the IP address of the sFlow agent to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID for the sFlow agent remains constant. If you do not assign an IP address to the agent, an IP address will be assigned to the agent using the IP address of a configured interface.

On the QFX3500 and QFX3600 standalone switches, the following priority is used to determine which interface will be used:

1. Management Ethernet interface me0 IP address
2. Any Layer 3 interface if the me0 IP address is not available

If a particular interface is not configured, the IP address of the next interface in the priority list is used as the IP address for the agent. Once an IP address is assigned to the agent, the agent ID is not modified until the sFlow service is restarted. At least one interface has to be configured for an IP address to be assigned to the agent.

In addition, you can explicitly configure the IP address for the source data (sFlow datagrams). On the QFX3500 and QFX3600 standalone switches, if you do not configure that address, the following priority is used:

- Any Layer 3 interface IP address
- The me0 IP address if no Layer 3 interface IP address is available

On the QFabric system, the following default values are used if the optional parameters are not configured:

- Agent ID is the management IP address of the default partition.
- Source IP is the management IP address of the default partition.

In addition, the QFabric system subagent ID (which is included in the sFlow datagrams) is the ID of the Node group from which the datagram is sent to the collector.

On the QFX Series, limitations of sFlow traffic sampling include:

- sFlow sampling on ingress interfaces does not capture CPU-bound traffic.
- sFlow sampling on egress interfaces does not support broadcast and multicast packets.
- Egress samples do not contain modifications made to the packet in the egress pipeline.
- If a packet is discarded because of a firewall filter, the reason code for discarding the packet is not sent to the collector.
- The **out-priority** field for a VLAN is always set to 0 (zero) on ingress and egress samples.
- You cannot configure sFlow monitoring on a link aggregation group (LAG), but you can configure it individually on a LAG member interface.

#### Related Documentation

- [Example: Monitoring Network Traffic Using sFlow Technology on page 6165](#)
- [Configuring sFlow Technology on page 6189](#)

## SNMP

- [Understanding the Implementation of SNMP on page 6107](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 6110](#)
- [Fabric Chassis MIB on page 6112](#)
- [Utility MIB on page 6116](#)
- [SNMPv3 Overview on page 6117](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6118](#)
- [Understanding RMON on page 6119](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6121](#)
- [Understanding Health Monitoring on page 6123](#)
- [SNMP MIBs Support on page 6124](#)
- [SNMP Traps Support on page 6139](#)
- [MIB Objects for the QFX Series on page 6152](#)

### Understanding the Implementation of SNMP

The QFX Series products support the Simple Network Management Protocol (SNMP) that is implemented in the Junos OS software.



**NOTE:** By default, SNMP is not enabled on devices running Junos OS. For information on enabling SNMP on a device running Junos OS, see [“Configuring SNMP” on page 1703](#).

A typical SNMP implementation includes the following components:

- Network management system (NMS)—The NMS is a combination of hardware and software that is used to monitor and administer a network. Software running on the

NMS includes the SNMP manager, which collects information about network connectivity, activity, and events by polling the managed devices.

- **Managed device**—A managed device (also called a network element) is any device managed by the NMS. Routers and switches are common examples of managed devices. The SNMP agent is the SNMP process that resides on the managed device and communicates with the NMS.
- **SNMP agent**—The SNMP agent exchanges network management information with SNMP manager software running on an NMS, or host. The agent responds to requests for information and actions from the manager. The agent also controls access to the agent's MIB, the collection of objects that can be viewed or changed by the SNMP manager.

SNMP data is stored in a highly structured, hierarchical format known as a management information base (MIB). The MIB structure is based on a tree structure, which defines a grouping of objects into related sets. Each object in the MIB is associated with an object identifier (OID), which names the object. The “leaf” in the tree structure is the actual managed object instance, which represents a resource, event, or activity that occurs in your network device. The SNMP implementation in Junos OS uses both standard (developed by IETF and documented in RFCs) and Juniper Networks enterprise-specific MIBs.

Communication between the agent and the manager occurs in one of the following forms:

- **Get, GetBulk, and GetNext requests**—The manager requests information from the agent; the agent returns the information in a **Get** response message.
- **Set requests**—The manager changes the value of a MIB object controlled by the agent; the agent indicates status in a **Set** response message.
- **Traps notification**—The agent sends traps to notify the manager of significant events that occur on the network device.

The processes maintaining the SNMP management data include:

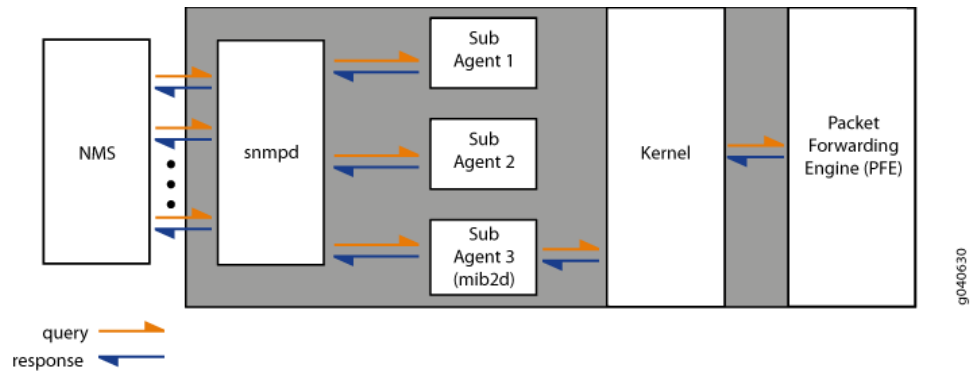
- A master SNMP agent (known as SNMP process, or `snmpd`) that resides on the managed device and is managed by the NMS or host.
- Various subagents that reside on different modules of Junos OS, such as the Routing Engine, and are managed by the master SNMP agent.
- Junos OS processes that share data with the subagents when polled for SNMP data (for example, interface-related MIBs).

When an NMS polls the master agent for data, the master agent immediately shares the data with the NMS if the requested data is available from the master agent or one of the subagents. However, if the requested data is not maintained by the master agent or subagents, the subagent polls the Junos OS kernel or the process that maintains that data. The Junos OS kernel may need to get the data from the Packet Forwarding Engine. On receiving the required data, the subagent passes the response back on to the master agent, which in turn passes it on to the NMS.



Figure 223 on page 6109 shows the communication flow among the NMS, SNMP master agent (snmpd), SNMP subagents, Junos OS kernel, and Packet Forwarding Engine.

Figure 223: SNMP Communication Flow



When a significant event, most often an error or a failure, occurs on a network device, the SNMP agent sends notifications to the SNMP manager. SNMP notifications can be sent as traps (unconfirmed notifications) or inform requests (confirmed notifications).

Junos OS supports trap queuing to ensure that traps are not lost because of temporary unavailability of routes. Two types of queues, destination queues and a throttle queue, are formed to ensure delivery of traps and control the trap traffic. On QFX Series products, the maximum size of trap queues (throttle queue plus destination queue) is 40,960 traps. The maximum size of any one queue is 20,480 traps.

Junos OS forms a destination queue when a trap to a particular destination is returned because the host is not reachable, and it adds the subsequent traps to the same destination to the queue. Junos OS checks for availability of routes every 30 seconds, and sends the traps from the destination queue in a round-robin fashion.

If the trap delivery fails, the trap is added back to the queue, and the delivery attempt counter and the next delivery attempt timer for the queue are reset. Subsequent attempts occur at progressive intervals of 1 minute, 2 minutes, 4 minutes, and 8 minutes. The maximum delay between the attempts is 8 minutes, and the maximum number of attempts is ten. After ten unsuccessful attempts, the destination queue and all the traps in the queue are deleted.

Junos OS also has a throttle mechanism to control the number of traps (throttle threshold) sent during a particular time period (throttle interval). The throttle mechanism ensures consistency in trap traffic, especially when large numbers of traps are generated because of interface status changes. The throttle interval period begins when the first trap arrives at the throttle. All traps within the trap threshold are processed, and the traps beyond the threshold limit are queued. The default throttle threshold is 500 traps, and the throttle interval default is 5 seconds.



**NOTE:** You cannot configure trap queueing in Junos OS. You cannot view information about trap queues except for what is provided in the system logs.

- Related Documentation**
- [Configuring SNMP on page 1703](#)
  - [SNMP MIBs Support on page 6124](#)
  - [SNMP Traps Support on page 6139](#)

## Understanding the Implementation of SNMP on the QFabric System

SNMP monitors network devices from a central location. The QFabric system supports the basic SNMP architecture of Junos OS, but its implementation of SNMP differs from that of other devices running Junos OS. This topic provides an overview of the SNMP implementation on the QFabric system.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent resides in the QFabric Director software and is responsible for receiving and distributing all traps as well as responding to all the queries of the SNMP manager. For example, traps that are generated by a Node device are sent to the SNMP agent in the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.



**NOTE:** In its SNMP implementation, the QFabric system acts as an SNMP proxy server, and requires more time to process SNMP requests than a typical Junos OS device does. The default timeout setting on most SNMP client applications is 3 seconds, which is not enough time for the QFabric system to respond to SNMP requests, so the results of your `mibwalk` command may be incomplete. For this reason, we recommend that you change the SNMP timeout setting to 5 seconds or longer for the QFabric system to complete the responses to your requests.

Support for SNMP on the QFabric system includes:

- Support for the SNMP Version 1 (v1) and v2.



**NOTE:** Only SNMPv2 traps are supported on the QFabric system.

- Support for the following standard MIBs:
  - RFC 1155, *Structure and Identification of Management Information for TCP/IP-based Internets*
  - RFC 1157, *A Simple Network Management Protocol (SNMP)*
  - RFC 1212, *Concise MIB Definitions*
  - RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II* (partial support, including the system group and interfaces group)
  - RFC 1215, *A Convention for Defining Traps for use with the SNMP*
  - RFC 1901, *Introduction to Community-based SNMPv2*

- RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol Using SMIv2*
- RFC 2012, *SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2*
- RFC 2013, *SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2*
- RFC 2233, *The Interfaces Group MIB Using SMIv2*
- RFC 2571, *An Architecture for Describing SNMP Management Frameworks* (read-only access) (excluding SNMPv3)
- RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (read-only access) (excluding SNMPv3)
- RFC 2576, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
- RFC 2579, *Textual Conventions for SMIv2*
- RFC 2580, *Conformance Statements for SMIv2*
- RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*
- RFC 2863, *The Interfaces Group MIB*
- RFC 3410, *Introduction and Applicability Statements for Internet Standard Management Framework* (excluding SNMPv3)
- RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Framework* (excluding SNMPv3)
- RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (excluding SNMPv3)
- RFC 3413, *Simple Network Management Protocol (SNMP) Applications* (excluding SNMPv3)
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)*
- RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 4188, *Definitions of Managed Objects for Bridges*

- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4363b, *Q-Bridge VLAN MIB*
- Support for the following Juniper Networks enterprise-specific MIBs:
  - Chassis MIB (mib-jnx-chassis.txt)
  - Class-of-Service MIB (mib-jnx-cos.txt)
  - Configuration Management MIB (mib-jnx-cfgmngmt.txt)
  - Fabric Chassis MIB (mib-jnx-fabric-chassis.txt)
  - Interface MIB Extensions (mib-jnx-if-extensions.txt)
  - Power Supply Unit MIB (mib-jnx-power-supply-unit.txt)
  - QFabric MIB (mib-jnx-qf-smi.txt)
  - Utility MIB (mib-jnx-util.txt)
- Support for operational mode commands—Limited to the **show snmp statistics** command. You may issue other SNMP requests, including **get**, **get next**, and **walk** requests, by using external SNMP client applications.

**Related  
Documentation**

- [SNMP MIBs Support on page 6124](#)
- [SNMP Traps Support on page 6139](#)

## Fabric Chassis MIB

The Juniper Networks enterprise-specific SNMP Fabric Chassis MIB (mib-jnx-fabric-chassis) provides hardware information about the QFabric system and its component devices in a single MIB. The Fabric Chassis MIB is based on the Juniper Networks enterprise-specific Chassis MIB that provides information for individual devices. Unlike the Chassis MIB, the Fabric Chassis MIB represents the QFabric system component devices as part of the QFabric system. Only the information from the Fabric Chassis MIB (and not from individual Chassis MIBs) is available to SNMP management clients of the QFabric system.

The Fabric Chassis MIB uses the basic information structure of the Chassis MIB, but adds another level of indexing that provides detailed information about QFabric system devices. Each physical device in a QFabric system (such as a Node device or an Interconnect device) is represented with its hardware components, including the power supply, fans, and front and rear cards.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent (snmpd) resides in the QFabric system Director software and is responsible for receiving and distributing all traps as well as responding to all queries from the SNMP manager. In addition, there is an SNMP subagent running in the Routing Engine of each Node group and Interconnect device. The SNMP subagent manages the information about the component device, and that information is communicated to the SNMP agent in the Director software as needed. Traps that are generated by a Node device are sent to the

SNMP agent in the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.

Table 585 on page 6113 describes the tables and objects in the Fabric Chassis MIB.

**Table 585: Fabric Chassis MIB Tables and Objects**

| Table or Object Name                               | Root OID                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Tables with Counterparts in the Chassis MIB</b> |                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| jnxFabricContainersTable                           | 1.3.6.1.4.1.2636.3.42.2.2.2 | <p>Provides information about different types of containers in QFabric system devices.</p> <ul style="list-style-type: none"> <li>Containers for Interconnect devices include fan trays, power supply units, control boards, and so on.</li> <li>Containers for Node devices include fan trays, power supply units, Flexible PIC Concentrator (FPC), PICs, and so on.</li> <li>Containers for the Director devices include CPU, memory, fan trays, power supply units, and hard disks. The containers have a non-hierarchical or flat structure, and components in them are organized as siblings to each other.</li> </ul>                                                                             |
| jnxFabricContentsTable                             | 1.3.6.1.4.1.2636.3.42.2.2.3 | <p>Contains contents that are present across all devices represented in the jnxFabricDeviceTable object. This table includes all field replaceable units (FRUs) and non-FRUs for QFabric system devices.</p> <ul style="list-style-type: none"> <li>Contents in the Interconnect devices include fan trays and control boards.</li> <li>Contents in the Node devices include fan trays and power supply units.</li> <li>Contents in the Director devices include CPUs, memory, fan trays, power supply units, and hard disks, but do not include network interface cards (NICs).</li> </ul>                                                                                                             |
| jnxFabricFilledTable                               | 1.3.6.1.4.1.2636.3.42.2.2.4 | <p>Shows the status of containers in QFabric devices. The jnxFabricFilledState object represents the state of the component: (1) unknown, (2) empty, or (3) filled.</p> <p><b>NOTE:</b> The jnxFabricFilledTable object does not contain information about the Director group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| jnxFabricOperatingTable                            | 1.3.6.1.4.1.2636.3.42.2.2.5 | <p>Represents different operating parameters for the contents that are populated in the jnxFabricContentsTable object.</p> <ul style="list-style-type: none"> <li>Contents in each Node device and Interconnect device include fan trays, power supply units, FPC, PIC, and Routing Engine.</li> <li>Contents in the Director device include CPUs, memory, fan trays, power supply units, and hard disks, but do not include network interface cards (NICs).</li> </ul> <p>The jnxFabricOperatingState object provides the state of the device: (1) unknown, (2) running, (3) ready, (4) reset, (5) runningAtFullSpeed (for fans only), (6) down, (6) off (for power supply units), or (7) standby.</p> |

Table 585: Fabric Chassis MIB Tables and Objects (*continued*)

| Table or Object Name     | Root OID                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jnxFabricRedundancyTable | 1.3.6.1.4.1.2636.3.42.2.2.6 | <p>Represents the redundancy information that is available at different subsystem levels across the QFabric system. Information about the Routing Engines in Node devices is included, but there are no corresponding entries for Interconnect devices in this table. The jnxFabricRedundancyState object indicates the state of the subsystem: (1) unknown, (2) master, (3) backup, or (4) disabled.</p> <p><b>NOTE:</b> Information about redundant Director devices, virtual machines (VMs) within Director groups, and Virtual Chassis devices is not available at this time.</p> |
| jnxFabricFruTable        | 1.3.6.1.4.1.2636.3.42.2.2.7 | <p>Contains all FRUs for the QFabric system in the jnxFabricDeviceTable table. The FRUs are listed regardless of whether or not they are installed or online. The jnxFabricFruState object represents the state of the FRU, including online, offline, or empty, and so on. This table also contains information about each FRU, such as name, type, temperature, time last powered on, and time last powered off.</p> <p><b>NOTE:</b> The jnxFabricFruTable table does not include network interface cards (NICs) on Director devices.</p>                                           |

#### Table Specific to the Fabric Chassis MIB

|                      |                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jnxFabricDeviceTable | 1.3.6.1.4.1.2636.3.42.2.2.1 | <p>Contains information about all devices in the QFabric system. This table organizes scalar variables represented in the Chassis MIB into a table format for the QFabric system component devices. Columns in this table include device information such as model, device alias, and serial number. The jnxFabricDeviceIndex identifies each QFabric system device (Node device, Interconnect device, and Director device).</p> <p><b>NOTE:</b> At this time, information about the Virtual Chassis is not available.</p> <p><b>NOTE:</b> The following objects are not supported:</p> <ul style="list-style-type: none"> <li>jnxFabricDeviceEntryRevision</li> <li>jnxFabricDeviceEntryFirmwareRevision</li> <li>jnxFabricDeviceEntryKernelMemoryUsedPercent</li> </ul> |
|----------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

#### Scalar Variables

Table 585: Fabric Chassis MIB Tables and Objects (*continued*)

| Table or Object Name                                                                                                                                                                                                                                                                                                      | Root OID                  | Description                                                                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>The following scalar variables are supported:</p> <ul style="list-style-type: none"> <li>• jnxFabricClass</li> <li>• jnxFabricDescr</li> <li>• jnxFabricSerialNo</li> <li>• jnxFabricRevision</li> <li>• jnxFabricLastInstalled</li> <li>• jnxFabricContentsLastChange</li> <li>• jnxFabricFilledLastChange</li> </ul> | 1.3.6.1.4.1.2636.3.42.2.1 | <p>Describe the QFabric system as a whole.</p> <p><b>NOTE:</b> The jnxFabricFirmwareRevision scalar variable is not supported at this time.</p> |

Table 586 on page 6115 describes the SNMPv2 traps that are defined in the Fabric Chassis MIB.



**NOTE:** Only SNMPv2 traps are supported on the QFabric system.

Table 586: Fabric Chassis MIB SNMPv2 Traps

| Trap Group and Name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Root OID              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>jnxFabricChassisTraps group—Includes the following traps:</p> <ul style="list-style-type: none"> <li>• jnxFabricPowerSupplyFailure</li> <li>• jnxFabricFanFailure</li> <li>• jnxFabricOverTemperature</li> <li>• jnxFabricRedundancySwitchover</li> <li>• jnxFabricFruRemoval</li> <li>• jnxFabricFruInsertion</li> <li>• jnxFabricFruPowerOff</li> <li>• jnxFabricFruPowerOn</li> <li>• jnxFabricFruFailed</li> <li>• jnxFabricFruOffline</li> <li>• jnxFabricFruOnline</li> <li>• jnxFabricFruCheck</li> <li>• jnxFabricFEBSwitchover</li> <li>• jnxFabricHardDiskFailed</li> <li>• jnxFabricHardDiskMissing</li> <li>• jnxFabricBootFromBackup</li> </ul> | 1.3.6.1.4.1.2636.4.19 | <p>Indicates an alarm condition.</p> <p><b>NOTE:</b> Hardware events on the Director group are detected by scanning. As a result, a trap may not be generated until up to 30 seconds after the event has occurred.</p> <p><b>NOTE:</b> The software does not distinguish between the fan removal and fan failure events on the Director group. In each case, both the jnxFabricFanFailure and jnxFabricFruFailed traps are generated.</p> <p><b>NOTE:</b> The software does not distinguish between the fan insertion and fan OK events on the Director group. In each case, both the jnxFabricFanOK and jnxFabricFruOK traps are generated.</p> |

Table 586: Fabric Chassis MIB SNMPv2 Traps (*continued*)

| Trap Group and Name                                                                                                                                                                                                         | Root OID              | Description                           |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|---------------------------------------|
| <p>jnxFabricChassisOKTraps group—Includes the following traps:</p> <ul style="list-style-type: none"> <li>jnxFabricPowerSupplyOK</li> <li>jnxFabricFanOK</li> <li>jnxFabricTemperatureOK</li> <li>jnxFabricFruOK</li> </ul> | 1.3.6.1.4.1.2636.4.20 | Indicates an alarm cleared condition. |

For more information, see the Fabric Chassis MIB at:

[http://www.juniper.net/techpubs/en\\_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt](http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt)

**Related  
Documentation**

- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
- *Chassis MIBs*

## Utility MIB

The Juniper Networks enterprise-specific Utility MIB, whose object ID is {jnxUtilMibRoot 1}, defines objects for counters, integers, and strings. The Utility MIB contains one table for each of the following five data types:

- 32-bit counters
- 64-bit counters
- Signed integers
- Unsigned integers
- Octet strings

Each data type has an arbitrary ASCII name, which is defined when the data is populated, and a timestamp that shows the last time when the data instance was modified. For a downloadable version of this MIB, see

[http://www.juniper.net/techpubs/en\\_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt](http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt).

For information about the enterprise-specific Utility MIB objects, see the following topics:

- *jnxUtilCounter32Table*
- *jnxUtilCounter64Table*
- *jnxUtilIntegerTable*
- *jnxUtilUintTable*
- *jnxUtilStringTable*

**Related  
Documentation**

- *Juniper Networks Enterprise-Specific MIBs*
- *Standard SNMP MIBs Supported by Junos OS*



- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)

## SNMPv3 Overview

The QFX3500 switch supports SNMP version 3 (SNMPv3). SNMPv3 enhances the functionality of SNMPv1 and SNMPv2c by supporting user authentication and data encryption. SNMPv3 uses the user-based security model (USM) to provide security for SNMP messages, and the view-based access control model (VACM) for user access control.

SNMPv3 features include:

- With USM, the SNMP messages between the SNMP manager and the agent can have the message source authenticated and the data integrity checked. USM reduces messaging delays and message replays by enforcing timeout limits and by checking for duplicate message request IDs.
- VACM complements USM by providing user access control for SNMP queries to the agent. You define access privileges that you wish to extend to a group of one or more users. Access privileges are determined by the security model parameters (**usm**, **v1**, or **v2**) and security level parameters (**authentication**, **privacy**, or **none**). For each security level, you must associate one MIB view for the group. Associating a MIB view with a group grants the read, write, or notify permission to a set of MIB objects for the group.
- You configure security parameters for each user, including the username, authentication type and authentication password, and privacy type and privacy password. The username given to each user is in a format that is dependent on the security model configured for that user.
- To ensure messaging security, another type of username, called the security name, is included in the messaging data that is sent between the local SNMP server and the destination SNMP server. Each user name is mapped to a security name, but the security name is in a format that is independent of the security model.
- Trap entries in SNMPv3 are created by configuring the notify, notify filter, target address, and target parameters. The **notify** statement specifies the type of notification (trap) and contains a single tag that defines a set of target addresses to receive a trap. The notify filter defines access to a collection of trap object identifiers (OIDs). The target address defines the address of an SNMP management application and other attributes used in sending notifications. Target parameters define the message processing and security parameters used in sending notifications to a particular target.

### Related Documentation

- [Assigning a Security Name to a Group on page 6204](#)
- [Configuring Access Privileges for a Group on page 6202](#)
- [Configuring SNMP Informs on page 6206](#)
- [Creating SNMPv3 Users on page 6201](#)

## Minimum SNMPv3 Configuration on a Device Running Junos OS

To configure the minimum requirements for SNMPv3, include the following statements at the `[edit snmp v3]` and `[edit snmp]` hierarchy levels:



**NOTE:** You must configure at least one view (notify, read, or write) at the `[edit snmp view-name]` hierarchy level.

```
[edit snmp]
view view-name {
 oid object-identifier (include | exclude);
}
[edit snmp v3]
notify name {
 tag tag-name;
}
notify-filter profile-name {
 oid object-identifier (include | exclude);
}
snmp-community community-index {
 security-name security-name;
}
target-address target-address-name {
 address address;
 target-parameters target-parameters-name;
}
target-parameters target-parameters-name {
 notify-filter profile-name;
 parameters {
 message-processing-model (v1 | v2c | v3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 security-name security-name;
 }
}
usm {
 local-engine {
 user username {
 }
 }
}
vacm {
 access {
 group group-name {
 (default-context-prefix | context-prefix context-prefix){
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }
 }
}
```

```

 }
 }
 security-to-group {
 security-model (usm | v1 | v2c) {
 security-name security-name {
 group group-name;
 }
 }
 }
}

```

#### Related Documentation

- [Creating SNMPv3 Users on page 6201](#)
- [Configuring MIB Views on page 6197](#)
- [Defining Access Privileges for an SNMP Group](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6204](#)
- [Configuring SNMP Informs on page 6206](#)
- [Complete SNMPv3 Configuration Statements](#)
- [Example: SNMPv3 Configuration](#)

## Understanding RMON

- [RMON Overview on page 6119](#)
- [Alarm Thresholds and Events on page 6120](#)

### RMON Overview

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

An operational support system (OSS) or a fault-monitoring system can be used to automatically monitor events that track many different metrics, including performance, availability, faults, and environmental data. For example, an administrator might want to know when the internal temperature of a chassis has risen above a configured threshold, which might indicate that a chassis fan tray is faulty, the chassis air flow is impeded, or the facility cooling system in the vicinity of the chassis is not operating normally.

The RMON MIB also defines tables that store various statistics for Ethernet interfaces, including the **etherStatsTable** and the **etherHistoryTable**. The **etherStatsTable** contains cumulative real-time statistics for Ethernet interfaces, such as the number of unicast, multicast, and broadcast packets received on an interface. The **etherHistoryTable** maintains a historical sample of statistics for Ethernet interfaces. The control of the **etherHistoryTable**, including the interfaces to track and the sampling interval, is defined by the RMON **historyControlTable**.

To enable RMON alarms, you perform the following steps:


1. Configure SNMP, including trap groups. You configure SNMP at the `[edit snmp]` hierarchy level.
2. Configure rising and falling events in the `eventTable`, including the event types and trap groups. You can also configure events using the CLI at the `[edit snmp rmon event]` hierarchy level.
3. Configure alarms in the `alarmTable`, including the variables to monitor, rising and falling thresholds, the sampling types and intervals, and the corresponding events to generate when alarms occur. You can also configure alarms using the CLI at the `[edit snmp rmon alarm]` hierarchy level.

Extensions to the `alarmTable` are defined in the Juniper Networks enterprise-specific MIB `jnxRmon` (`mib-jnx-rmon.txt`).

### Alarm Thresholds and Events

By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range (see [Figure 224 on page 6120](#)).

Figure 224: Setting Thresholds



Events are only generated when the alarm threshold is first crossed in any one direction rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs. This considerably reduces the quantity of events that are produced by the system, making it easier for operations staff to react when events do occur.

Before you configure remote monitoring, you should identify what variables need to be monitored and their allowable operational range. This requires some period of baselining to determine the allowable operational ranges. An initial baseline period of at least

3 months is not unusual when you first identify the operational ranges and define thresholds, but baseline monitoring should continue over the life span of each monitored variable.

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6198](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6121](#)

## RMON MIB Event, Alarm, Log, and History Control Tables

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

[Table 587 on page 6121](#) provides each field in the RMON **eventTable**, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the **[edit snmp rmon]** hierarchy level.

**Table 587: RMON Event Table**

| Field                   | Description                                                                                                                     | Statement [edit snmp rmon] |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| <b>eventDescription</b> | Text description of this event.                                                                                                 | <b>description</b>         |
| <b>eventType</b>        | Type of event (for example, <b>log</b> , <b>trap</b> , or <b>log and trap</b> ).                                                | <b>type</b>                |
| <b>eventCommunity</b>   | Trap group to which to send this event, as defined in the Junos OS configuration. (This is not the same as the SNMP community.) | <b>community</b>           |
| <b>eventOwner</b>       | Entity (for example, <b>manager</b> ) that created this event.                                                                  | —                          |
| <b>eventStatus</b>      | Status of this row (for example, <b>valid</b> , <b>invalid</b> , or <b>createRequest</b> ).                                     | —                          |

[Table 588 on page 6121](#) provides each field in the RMON **alarmTable**, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the **[edit snmp rmon]** hierarchy level.

**Table 588: RMON Alarm Table**

| Field                | Description                                                                                | Statement [edit snmp rmon] |
|----------------------|--------------------------------------------------------------------------------------------|----------------------------|
| <b>alarmStatus</b>   | Status of this row (for example, <b>valid</b> , <b>invalid</b> , or <b>createRequest</b> ) | —                          |
| <b>alarmInterval</b> | Sampling period (in seconds) of the monitored variable                                     | <b>interval</b>            |
| <b>alarmVariable</b> | Object identifier (OID) and instance of the variable to be monitored                       | —                          |

Table 588: RMON Alarm Table (*continued*)

| Field                         | Description                                                         | Statement [edit snmp rmon] |
|-------------------------------|---------------------------------------------------------------------|----------------------------|
| <b>alarmValue</b>             | Actual value of the sampled variable                                | —                          |
| <b>alarmSampleType</b>        | Sample type ( <b>absolute</b> or <b>delta</b> changes)              | <b>sample-type</b>         |
| <b>alarmStartupAlarm</b>      | Initial alarm ( <b>rising</b> , <b>falling</b> , or <b>either</b> ) | <b>startup-alarm</b>       |
| <b>alarmRisingThreshold</b>   | Rising threshold against which to compare the value                 | <b>rising-threshold</b>    |
| <b>alarmFallingThreshold</b>  | Falling threshold against which to compare the value                | <b>falling-threshold</b>   |
| <b>alarmRisingEventIndex</b>  | Index (row) of the rising event in the event table                  | <b>rising-event-index</b>  |
| <b>alarmFallingEventIndex</b> | Index (row) of the falling event in the event table                 | <b>falling-event-index</b> |

Table 589 on page 6122 provides each field in the **jnxRmon jnxRmonAlarmTable**, which is an extension to the RMON **alarmTable**. You can troubleshoot the RMON agent, **rmopd**, that runs on a switch by inspecting the contents of the **jnxRmonAlarmTable** object.

Table 589: jnxRmon Alarm Table

| Field                            | Description                                                                       |
|----------------------------------|-----------------------------------------------------------------------------------|
| <b>jnxRmonAlarmGetFailCnt</b>    | Number of times the internal <b>Get</b> request for the variable failed           |
| <b>jnxRmonAlarmGetFailTime</b>   | Value of the <b>sysUpTime</b> object when the last failure occurred               |
| <b>jnxRmonAlarmGetFailReason</b> | Reason why the <b>Get</b> request failed                                          |
| <b>jnxRmonAlarmGetOkTime</b>     | Value of the <b>sysUpTime</b> object when the variable moved out of failure state |
| <b>jnxRmonAlarmState</b>         | Status of this alarm entry                                                        |

Table 590 on page 6122 provides each field in the RMON **historyControlTable**, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the **[edit snmp rmon history]** hierarchy level. The **historyControlTable** controls the RMON **etherHistoryTable**.

Table 590: RMON History Control Table

| Field                           | Description                                                                | Statement [edit snmp rmon history] |
|---------------------------------|----------------------------------------------------------------------------|------------------------------------|
| <b>historyControlDataSource</b> | Identifies the source of the data for which historical data was collected. | <b>interface</b>                   |

Table 590: RMON History Control Table (*continued*)

| Field                                       | Description                                                                 | Statement [edit snmp rmon history] |
|---------------------------------------------|-----------------------------------------------------------------------------|------------------------------------|
| <code>historyControlBucketsRequested</code> | Requested number of discrete time intervals over which data is to be saved. | <b>bucket-size</b>                 |
| <code>historyControlBucketsGranted</code>   | Number of discrete sampling intervals over which data is to be saved.       | —                                  |
| <code>historyControlInterval</code>         | Interval, in seconds, over which the data is sampled for each bucket.       | <b>interval</b>                    |
| <code>historyControlOwner</code>            | Entity that configured this entry.                                          | <b>owner</b>                       |
| <code>historyControlStatus</code>           | Status of this entry.                                                       | —                                  |

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6198](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [Understanding RMON on page 6119](#)

## Understanding Health Monitoring

Health monitoring is an SNMP feature that extends the RMON alarm infrastructure to provide monitoring for a predefined set of objects (such as file system usage, CPU usage, and memory usage), and for Junos OS processes.

You enable the health monitor feature using the **health-monitor** statement at the **[edit snmp]** hierarchy level. You can also configure health monitor parameters such as a falling threshold, rising threshold, and interval. If the value of a monitored object exceeds the rising or falling threshold, an alarm is triggered and an event may be logged.

The falling threshold is the lower threshold for the monitored object instance. The rising threshold is the upper threshold for the monitored object instance. Each threshold is expressed as a percentage of the maximum possible value. The interval represents the period of time, in seconds, over which the object instance is sampled and compared with the rising and falling thresholds.

Events are only generated when a threshold is first crossed in any one direction, rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs.

System log entries for health monitor events have a corresponding **HEALTHMONITOR** tag and not a generic **SNMPD\_RMON\_EVENTLOG** tag. However, the health monitor sends generic RMON **risingThreshold** and **fallingThreshold** traps. You can use the **show snmp health-monitor** operational command to view information about health monitor alarms and logs.

When you configure the health monitor, monitoring information for certain object instances is available, as shown in [Table 591 on page 6124](#).

**Table 591: Monitored Object Instances**

| Object                                 | Description                                                                                                                                         |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>jnxHrStoragePercentUsed.1</code> | Monitors the <code>/dev/ad0s1a</code> : file system on the switch. This is the root file system mounted on <code>/</code> .                         |
| <code>jnxHrStoragePercentUsed.2</code> | Monitors the <code>/dev/ad0s1e</code> : file system on the switch. This is the configuration file system mounted on <code>/config</code> .          |
| <code>jnxOperatingCPU (RE0)</code>     | Monitors CPU usage by the Routing Engine (RE0).                                                                                                     |
| <code>jnxOperatingBuffer (RE0)</code>  | Monitors the amount of memory available on the Routing Engine (RE0).                                                                                |
| <code>sysAppElmtRunCPU</code>          | Monitors the CPU usage for each Junos OS process (also called daemon). Multiple instances of the same process are monitored and indexed separately. |
| <code>sysAppElmtRunMemory</code>       | Monitors the memory usage for each Junos OS process. Multiple instances of the same process are monitored and indexed separately.                   |

- Related Documentation**
- [Configuring Health Monitoring on page 6200](#)
  - [falling-threshold \(Health Monitor\) on page 1772](#)
  - [interval \(Health Monitor\) on page 1777](#)
  - [rising-threshold \(Health Monitor\) on page 1807](#)
  - [show snmp health-monitor on page 6445](#)

## SNMP MIBs Support

The QFX3500 and QFX3600 switches and QFabric systems support standard MIBs and Juniper Networks enterprise-specific MIBs.

For more information, see:

- [MIBs Supported on QFX3500 and QFX3600 Switches on page 6124](#)
- [MIBs Supported on QFabric Systems on page 6133](#)

### MIBs Supported on QFX3500 and QFX3600 Switches

The QFX3500 and QFX3600 switches support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 592 on page 6125](#) for standard MIBs
- [Table 593 on page 6130](#) for Juniper Networks enterprise-specific MIBs



Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches

| RFC                                                                                                   | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IEEE 802.1ab section 12.1, <i>Link Layer Discovery Protocol (LLDP) MIB</i>                            | Supported tables and objects: <ul style="list-style-type: none"> <li>• lldpRemManAddrOID</li> <li>• lldpLocManAddrOID</li> <li>• lldpReinitDelay</li> <li>• lldpNotificationInterval</li> <li>• lldpStatsRxPortFramesDiscardedTotal</li> <li>• lldpStatsRxPortFramesError</li> <li>• lldpStatsRxPortTLVsDiscardedTotal</li> <li>• lldpStatsRxPortTLVsUnrecognizedTotal</li> <li>• lldpStatsRxPortAgeoutsTotal</li> </ul>                                                                                                                                                                                                                                          |
| IEEE 802.3ad, <i>Aggregation of Multiple Link Segments</i>                                            | The following tables and objects are supported: <ul style="list-style-type: none"> <li>• dot3adAggPortTable, dot3adAggPortListTable, dot3adAggTable, and dot3adAggPortStatsTable</li> <li>• dot3adAggPortDebugTable (only dot3adAggPortDebugRxState, dot3adAggPortDebugMuxState, dot3adAggPortDebugActorSyncTransitionCount, dot3adAggPortDebugPartnerSyncTransitionCount, dot3adAggPortDebugActorChangeCount, and dot3adAggPortDebugPartnerChangeCount)</li> <li>• dot3adTablesLastChanged</li> </ul>                                                                                                                                                            |
| RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>    | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>                                          | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| RFC 1212, <i>Concise MIB Definitions</i>                                                              | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i> | The following areas are supported: <ul style="list-style-type: none"> <li>• MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> <li>• Statistics counters</li> <li>• IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>)</li> <li>• ipAddrTable</li> <li>• SNMP management</li> <li>• Interface management</li> </ul> </li> <li>• SNMPv1 <b>Get</b>, <b>GetNext</b> requests, and SNMPv2 <b>GetBulk</b> request</li> <li>• Junos OS-specific secured access list</li> <li>• Master configuration keywords</li> <li>• Reconfigurations upon SIGHUP</li> </ul> |
| RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>                                | Support is limited to MIB II SNMP version 1 traps and version 2 notifications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)

| RFC                                                                                                                       | Additional Information                                                                                                                                                                                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 1286, <i>Definitions of Managed Objects for Bridges</i>                                                               | —                                                                                                                                                                                                                                                                                      |
| RFC 1657, <i>Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2</i> | —                                                                                                                                                                                                                                                                                      |
| RFC 1850, <i>OSPF Version 2 Management Information Base</i>                                                               | <p>The following table, objects, and traps are <i>not</i> supported:</p> <ul style="list-style-type: none"> <li>• Host Table</li> <li>• ospfOriginateNewLsas and ospfRxNewLsas objects</li> <li>• ospfOriginateLSA, ospfLsdbOverflow, and ospfLsdbApproachingOverflow traps</li> </ul> |
| RFC 1901, <i>Introduction to Community-based SNMPv2</i>                                                                   | —                                                                                                                                                                                                                                                                                      |
| RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>                     | —                                                                                                                                                                                                                                                                                      |
| RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>             | —                                                                                                                                                                                                                                                                                      |
| RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>                                 | —                                                                                                                                                                                                                                                                                      |
| RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>                     | —                                                                                                                                                                                                                                                                                      |
| RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>                            | —                                                                                                                                                                                                                                                                                      |
| RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>                                                                     | <b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.                                                                                                                                                                            |
| RFC 2287, <i>Definitions of System-Level Managed Objects for Applications</i>                                             | <p>The following objects are supported:</p> <ul style="list-style-type: none"> <li>• sysApplInstallPkgTable</li> <li>• sysApplInstallElmtTable</li> <li>• sysApplElmtRunTable</li> <li>• sysApplMapTable</li> </ul>                                                                    |
| RFC 2570, <i>Introduction to Version 3 of the Internet-standard Network Management Framework</i>                          | —                                                                                                                                                                                                                                                                                      |

Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)

| RFC                                                                                                                            | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)                                  | <b>NOTE:</b> RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.                                                                                                                                                                                                                                                                                                                                            |
| RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)       | <b>NOTE:</b> RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.                                                                                                                                                                                                                                                                                                                                            |
| RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> | <b>NOTE:</b> RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.                                                                                                                                                                                                                                                                                                                                            |
| RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>                                                         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2579, <i>Textual Conventions for SMIv2</i>                                                                                 | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2580, <i>Conformance Statements for SMIv2</i>                                                                              | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>                                          | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>                                     | Support does not include row creation, the Set operation, and the <code>vrpStatsPacketLengthErrors</code> object.                                                                                                                                                                                                                                                                                                                                      |
| RFC 2790, <i>Host Resources MIB</i>                                                                                            | Support is limited to the following objects: <ul style="list-style-type: none"> <li>Only <code>hrStorageTable</code>. The file systems <code>/</code>, <code>/config</code>, <code>/var</code>, and <code>/tmp</code> always return the same index number. When SNMP restarts, the index numbers for the remaining file systems might change.</li> <li>Only the objects of the <code>hrSystem</code> and <code>hrSWInstalled</code> groups.</li> </ul> |
| RFC 2819, <i>Remote Network Monitoring Management Information Base</i>                                                         | The following objects are supported: <ul style="list-style-type: none"> <li><code>etherStatsTable</code> (for Ethernet interfaces only), <code>alarmTable</code>, <code>eventTable</code>, and <code>logTable</code>.</li> <li><code>historyControlTable</code> and <code>etherHistoryTable</code> (except the <code>etherHistoryUtilization</code> object).</li> </ul>                                                                                |
| RFC 2863, <i>The Interfaces Group MIB</i>                                                                                      | <b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.                                                                                                                                                                                                                                                                                                                                            |
| RFC 2932, <i>IPv4 Multicast Routing MIB</i>                                                                                    | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>                                                                 | —                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| RFC 2934, <i>Protocol Independent Multicast MIB for IPv4</i>                                                                   | In Junos OS, RFC 2934 is implemented based on a draft version, <code>pimmib.mib</code> , of the now standard RFC.                                                                                                                                                                                                                                                                                                                                      |

Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)

| RFC                                                                                                                            | Additional Information                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>                          | —                                                                                               |
| RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>                | <b>NOTE:</b> RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571. |
| RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>                          | <b>NOTE:</b> RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572. |
| RFC 3413, <i>Simple Network Management Protocol (SNMP) Applications</i>                                                        | All MIBs are supported except for the Proxy MIB.                                                |
| RFC 3414, <i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i>              | —                                                                                               |
| RFC 3415, <i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i>                      | —                                                                                               |
| RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>                        | <b>NOTE:</b> RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.   |
| RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>                                          | —                                                                                               |
| RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>                           | <b>NOTE:</b> RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.   |
| RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> | —                                                                                               |
| RFC 3826, <i>The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model</i>                 | —                                                                                               |

Table 592: Standard MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)

| RFC                                                                                                                                                               | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 4188, <i>Definitions of Managed Objects for Bridges</i>                                                                                                       | <p>The QFX3500 and QFX3600 switches support 802.1D STP (1998) and the following subtrees and objects only:</p> <ul style="list-style-type: none"> <li>• dot1dTp subtree—dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus objects from the dot1dTpFdbTable table.</li> <li>• dot1dBase subtree—dot1dBasePort and dot1dBasePortIfIndex objects from the dot1dBasePortTable table.</li> </ul> <p><b>NOTE:</b> On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (RFC 4363b, <i>Q-Bridge VLAN MIB</i>) when you issue the <b>show snmp mib walk</b> command.</p> |
| RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>                                                                                       | Supports the ipAddrTable table only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| RFC 4318, <i>Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol</i>                                                                     | Supports 802.1w and 802.1t extensions for RSTP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| RFC 4363b, <i>Q-Bridge VLAN MIB</i>                                                                                                                               | <p><b>NOTE:</b> On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table (RFC 4188, <i>Definitions of Managed Objects for Bridges</i>) is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (in this MIB) when you issue the <b>show snmp mib walk</b> command.</p>                                                                                                                                                                                                                                                                                                                                                             |
| RFC 4444, <i>IS-IS MIB</i>                                                                                                                                        | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Internet Assigned Numbers Authority, <i>IANAiftype Textual Convention MIB</i> (referenced by RFC 2233)                                                            | See <a href="http://www.iana.org/assignments/ianaiftype-mib">http://www.iana.org/assignments/ianaiftype-mib</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Internet draft<br>draft-reeder-snmpv3-usm-3desede-00.txt, <i>Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in 'Outside' CBC Mode</i> | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Internet draft<br>draft-ietf-idmr-igmp-mib-13.txt, <i>Internet Group Management Protocol (IGMP) MIB</i>                                                           | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| ESO Consortium MIB                                                                                                                                                | <p><b>NOTE:</b> The ESO Consortium MIB has been replaced by RFC 3826. See <a href="http://www.snmp.com/eso/">http://www.snmp.com/eso/</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 593: Juniper Networks Enterprise-Specific MIBs Supported on QFX3500 and QFX3600 Switches**

| MIB                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm MIB (mib-jnx-chassis-alarm)                               | <p>Provides support for alarms from the switch.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chassis-alarm.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chassis-alarm.txt</a>.</p> <p>For more information, see <i>Alarm MIB</i>.</p>                                                                                                                                                                                                                                                             |
| Analyzer MIB (mib-jnx-analyzer)                                 | <p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-analyzer.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-analyzer.txt</a>.</p> <p>For more information, see <i>Analyzer MIB</i>.</p>                                                                                                                                                                                                                                           |
| Chassis MIB (mib-jnx-chassis)                                   | <p>Provides support for environmental monitoring (power supply state, board voltages, fans, temperatures, and airflow) and inventory support for the chassis, Flexible PIC Concentrators (FPCs), and PICs.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chassis.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chassis.txt</a>.</p> <p>For more information, see <i>Chassis MIBs</i>.</p>                                                                                                           |
| Chassis Definitions for Router Model MIB (mib-jnx-chas-defines) | <p>Contains the object identifiers (OIDs) that are used by the Chassis MIB to identify routing and switching platforms and chassis components. The Chassis MIB provides information that changes often, whereas the Chassis Definitions for Router Model MIB provides information that changes less often.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chas-defines.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-chas-defines.txt</a>.</p> <p>For more information, see <i>Chassis MIBs</i>.</p> |
| Class-of-Service MIB (mib-jnx-cos)                              | <p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cos.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cos.txt</a>.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p>                                                                                                                                                                                                         |

**Table 593: Juniper Networks Enterprise-Specific MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)**

| MIB                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Management MIB (mib-jnx-cfgmgmt) | <p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cfgmgmt.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cfgmgmt.txt</a>.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p> |
| Ethernet MAC MIB (mib-jnx-mac)                 | <p>Monitors media access control (MAC) statistics on Gigabit Ethernet intelligent queuing (IQ) interfaces. It collects MAC statistics; for example, inoctets, inframes, outoctets, and outframes on each source MAC address and virtual LAN (VLAN) ID for each Ethernet port.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-mac.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-mac.txt</a>.</p> <p>For more information, see <i>Ethernet MAC MIB</i>.</p>                                                                                                               |
| Event MIB (mib-jnx-event)                      | <p>Defines a generic trap that can be generated using an operations script or event policy. This MIB provides the ability to specify a system log string and raise a trap if that system log string is found.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-event.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-event.txt</a>.</p> <p>For more information, see <i>Event MIB</i>.</p>                                                                                                                                                                                  |
| Firewall MIB (mib-jnx-firewall)                | <p>Provides support for monitoring firewall filter counters.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-firewall.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-firewall.txt</a>.</p> <p>For more information, see <i>Firewall MIB</i>.</p>                                                                                                                                                                                                                                                                                                                          |
| Host Resources MIB (mib-jnx-hostresources)     | <p>Extends the hrStorageTable object, providing a measure of the usage of each file system on the switch as a percentage. Previously, the objects in the hrStorageTable measured the usage in allocation units—hrStorageUsed and hrStorageAllocationUnits—only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-hostresources.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-hostresources.txt</a>.</p> <p>For more information, see <i>Host Resources MIB</i>.</p>          |

**Table 593: Juniper Networks Enterprise-Specific MIBs Supported on QFX3500 and QFX3600 Switches (continued)**

| MIB                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface MIB (Extensions)<br>(mib-jnx-if-extensions) | <p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-if-extensions.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-if-extensions.txt</a>.</p> <p>For more information, see <i>Interface MIB</i>.</p>                                                                                                                 |
| Ping MIB (mib-jnx-ping)                               | <p>Extends the standard Ping MIB control table (RFC 2925). Items in this MIB are created when entries are created in pingCtlTable of the Ping MIB. Each item is indexed exactly as it is in the Ping MIB.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-ping.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-ping.txt</a>.</p> <p>For more information, see <i>PING MIB</i>.</p>                                                                                                              |
| RMON Events and Alarms MIB<br>(mib-jnx-rmon)          | <p>Supports Junos OS extensions to the standard Remote Monitoring (RMON) Events and Alarms MIB (RFC 2819). The extension augments the alarmTable object with additional information about each alarm. Two additional traps are also defined to indicate when problems are encountered with an alarm.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-rmon.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-rmon.txt</a>.</p> <p>For more information, see <i>RMON Events and Alarms MIB</i>.</p> |
| Structure of Management Information MIB (mib-jnx-smi) | <p>Explains how the Juniper Networks enterprise-specific MIBs are structured.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-smi.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-smi.txt</a>.</p> <p>For more information, see <i>Structure of Management Information MIB</i>.</p>                                                                                                                                                                                                             |
| System Log MIB (mib-jnx-syslog)                       | <p>Enables notification of an SNMP trap-based application when an important system log message occurs.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-syslog.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-syslog.txt</a>.</p> <p>For more information, see <i>System Log MIB</i>.</p>                                                                                                                                                                                                       |



**Table 593: Juniper Networks Enterprise-Specific MIBs Supported on QFX3500 and QFX3600 Switches (*continued*)**

| MIB                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Utility MIB (mib-jnx-util) | <p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt</a>.</p> <p>For more information, see “Utility MIB” on page 6116 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 6396.</p> |
| VLAN MIB (mib-jnx-vlan)    | <p>Contains information about prestandard IEEE 802.10 VLANs and their association with LAN emulation clients.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-vlan.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-vlan.txt</a>.</p> <p>For more information, see <i>VLAN MIB</i>.</p>                                                                                                                                                                                                                                                |

### MIBs Supported on QFabric Systems

The QFabric systems support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 594 on page 6133](#) for standard MIBs
- [Table 595 on page 6137](#) for Juniper Networks enterprise-specific MIBs

**Table 594: Standard MIBs Supported on QFabric Systems**

| RFC                                                                                                | Additional Information |
|----------------------------------------------------------------------------------------------------|------------------------|
| RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i> | —                      |
| RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>                                       | —                      |
| RFC 1212, <i>Concise MIB Definitions</i>                                                           | —                      |

Table 594: Standard MIBs Supported on QFabric Systems (*continued*)

| RFC                                                                                                           | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>         | <p>The following areas are supported:</p> <ul style="list-style-type: none"> <li>MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> <li>Statistics counters</li> <li>IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>)</li> <li>ipAddrTable</li> <li>SNMP management</li> <li>Interface management</li> </ul> </li> <li>SNMPv1 <b>Get</b>, <b>GetNext</b> requests, and version 2 <b>GetBulk</b> request</li> <li>Junos OS-specific secured access list</li> <li>Master configuration keywords</li> <li>Reconfigurations upon SIGHUP</li> </ul> |
| RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>                                        | Support is limited to MIB II SNMP version 1 traps and version 2 notifications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| RFC 1286, <i>Definitions of Managed Objects for Bridges</i>                                                   | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 1901, <i>Introduction to Community-based SNMPv2</i>                                                       | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>                     | <b>NOTE:</b> On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>                | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>                                                         | <p><b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p><b>NOTE:</b> The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 594: Standard MIBs Supported on QFabric Systems (*continued*)

| RFC                                                                                                                            | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)                                  | <b>NOTE:</b> RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)       | <b>NOTE:</b> RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> | <b>NOTE:</b> RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>                                                         | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RFC 2579, <i>Textual Conventions for SMIv2</i>                                                                                 | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RFC 2580, <i>Conformance Statements for SMIv2</i>                                                                              | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>                                          | <p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> <li>dot3StatsTable—There is one row with statistics for each Ethernet-like interface in the QFabric system. The dot3StatsIndex is an interface index that is unique across the system.</li> <li>dot3ControlTable—There is one row in this table for each Ethernet-like interface in the QFabric system that implements the MAC control sublayer. OIDs supported are dot3ControlFunctionsSupported and dot3ControlInUnknownOpcode.</li> <li>dot3PauseTable—There is one row in this table for each Ethernet-like interface in the QFabric system that supports the MAC control PAUSE function. OIDs supported are dot3PauseAdminMode, dot3PauseOperMode, dot3InPauseFrames, and dot3OutPauseFrames.</li> </ul> <p><b>NOTE:</b> Scalar variables are not supported on the QFabric system.</p> |
| RFC 2863, <i>The Interfaces Group MIB</i>                                                                                      | <p><b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p><b>NOTE:</b> The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>                                                                 | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>                          | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 594: Standard MIBs Supported on QFabric Systems (*continued*)

| RFC                                                                                                                            | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>                | <b>NOTE:</b> RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>                          | <b>NOTE:</b> RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>                        | <b>NOTE:</b> RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>                                          | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>                           | <b>NOTE:</b> RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| RFC 4188, <i>Definitions of Managed Objects for Bridges</i>                                                                    | <p>The QFabric system support is limited to the following objects:</p> <ul style="list-style-type: none"> <li>• Under the dot1dBase OID, the dot1dBasePortTable table supports only the first two columns in the table: dot1dBasePort and dot1dBasePortIfIndex.</li> <li>• The system does not implement the optional traps supporting dot1dNotifications (dot1dBridge 0).</li> <li>• Under the dot1dStp OID, supports only the dot1dStpPortTable table. Does not support the scalar variables under dot1dStp.</li> <li>• The system does not support scalar variables under dot1dTp, but under that, the dot1dTpFdbTable table is supported (dot1dBridge 4).</li> <li>• For OIDs with tables support only, scalar values that are returned by the SNMP agent may not be meaningful and are therefore not recommended for use.</li> </ul> |
| RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>                                                    | <p>Supports the ipAddrTable table only.</p> <p>On the QFabric system, supported objects in the ipAddrTable table include ipAdEntAddr, ipAdEntIfIndex, ipAdEntNetMask, ipAdEntBcastAddr, and ipAdEntReasmMaxSize.</p> <p><b>NOTE:</b> On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.</p>                                                                                                                                                                                                                                                                                                                                                                                         |

Table 594: Standard MIBs Supported on QFabric Systems (*continued*)

| RFC                                 | Additional Information                                                                                                                                                                                                   |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RFC 4363b, <i>Q-Bridge VLAN MIB</i> | <p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> <li>• dot1qTpFdbTable</li> <li>• dot1qVlanStaticTable</li> <li>• dot1qPortVlanTable</li> <li>• dot1qFdbTable</li> </ul> |

Table 595: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems

| MIB                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analyzer MIB (mib-jnx-analyzer)    | <p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>The QFabric system supports:</p> <ul style="list-style-type: none"> <li>• Analyzer table—jnxAnalyzerName, jnxMirroringRatio, jnxLossPriority.</li> <li>• Analyzer input table—jnxAnalyzerInputValue, jnxAnalyzerInputOption, jnxAnalyzerInputType.</li> <li>• Analyzer output table—jnxAnalyzerOutputValue, jnxAnalyzerOutputType.</li> </ul> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-analyzer.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-analyzer.txt</a>.</p> <p>For more information, see <i>Analyzer MIB</i>.</p>                                                                                                                                                              |
| Chassis MIB (mib-jnx-chassis)      | <p><b>NOTE:</b> The Chassis MIB has been deprecated for the QFabric system. We recommend that you use the Fabric Chassis MIB (mib-jnx-fabric-chassis) for information about the QFabric system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Class-of-Service MIB (mib-jnx-cos) | <p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>The QFabric system supports the following tables and objects:</p> <ul style="list-style-type: none"> <li>• Jnxcosifstatflagtable—jnxCosIfstatFlags and jnxCosIfIndex.</li> <li>• Jnxcosqstattable—jnxCosQstatTxedPkts, jnxCosQstatTxedPktRate, jnxCosQstatTxedBytes, and jnxCosQstatTxedByteRate.</li> <li>• Jnxcosfcidtable—jnxCosFcIdToFcName.</li> <li>• Jnxcosfctable—jnxCosFcQueueNr.</li> </ul> <p>The QFabric system does not support any traps for this MIB.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cos.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cos.txt</a>.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p> |

**Table 595: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems (continued)**

| MIB                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Management MIB<br>(mib-jnx-cfgmgmt)     | <p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p><b>NOTE:</b> On the QFabric system, these conditions apply:</p> <ul style="list-style-type: none"> <li>• All scalar variables under the jnxCmCfgChg table are supported.</li> <li>• Supported scalar OIDs are jnxCmCfgChgLatestIndex, jnxCmCfgChgLatestTime, jnxCmCfgChgLatestDate, jnxCmCfgChgLatestSource, jnxCmCfgChgLatestUser, and jnxCmCfgChgMaxEventEntries.</li> <li>• Scalar variables under the jnxCmRescueChg table are not supported.</li> </ul> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cfgmgmt.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-cfgmgmt.txt</a>.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p> |
| Fabric Chassis MIB<br>(mib-jnx-fabric-chassis)        | <p>Provides hardware information about the QFabric system and its component devices. This MIB is based on the Juniper Networks enterprise-specific Chassis MIB but adds another level of indexing that provides information for QFabric system component devices.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt</a>.</p> <p>For more information, see “Fabric Chassis MIB” on page 6112.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Host Resources MIB<br>(mib-jnx-hostresources)         | <p>Extends the hrStorageTable object, providing a measure of the usage of each file system on the switch as a percentage. Previously, the objects in the hrStorageTable measured the usage in allocation units—hrStorageUsed and hrStorageAllocationUnits—only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-hostresources.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-hostresources.txt</a>.</p> <p>For more information, see <i>Host Resources MIB</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Interface MIB (Extensions)<br>(mib-jnx-if-extensions) | <p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p><b>NOTE:</b> On the QFabric system, scalar variables are not supported.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-if-extensions.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-if-extensions.txt</a>.</p> <p>For more information, see <i>Interface MIB</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 595: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems (*continued*)**

| MIB                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Power Supply Unit MIB<br>(mib-jnx-power-supply-unit) | <p>Provides support for environmental monitoring of the power supply unit for the Interconnect device of the QFabric system.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-power-supply-unit.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-power-supply-unit.txt</a>.</p> <p>For more information, see <i>Power Supply Unit MIB</i>.</p> <p><b>NOTE:</b> On the QFabric system, scalar variables for the jnxPsuObjects 1 object ID in the jnxPsuScalars table are not supported.</p>                                              |
| QFabric MIB (jnx-qf-smi)                             | <p>Explains how the Juniper Networks enterprise-specific QFabric MIBs are structured. Defines the MIB objects that are reported by the QFabric system and the contents of the traps that can be issued by the QFabric system.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-qf-smi.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-qf-smi.txt</a>.</p>                                                                                                                                                                              |
| Utility MIB (mib-jnx-util)                           | <p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt">http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-util.txt</a>.</p> <p>For more information, see “Utility MIB” on page 6116 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 6396.</p> |

- Related Documentation**
- *SNMP MIBs and Traps Reference*
  - [Understanding the Implementation of SNMP on page 6107](#)
  - [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
  - [SNMP Traps Support on page 6139](#)

## SNMP Traps Support

The QFX3500 and QFX3600 switches and QFabric systems support standard SNMP traps and Juniper Networks enterprise-specific traps.

For more information, see:

- [SNMP Traps Supported on QFX3500 and QFX3600 Switches on page 6140](#)
- [SNMP Traps Supported on QFabric Systems on page 6148](#)

## SNMP Traps Supported on QFX3500 and QFX3600 Switches

The QFX3500 and QFX3600 standalone switches support SNMPv1 and v2 traps. For more information, see:

- [SNMPv1 Traps on page 6140](#)
- [SNMPv2 Traps on page 6144](#)

### SNMPv1 Traps

QFX3500 and QFX3600 switches support both standard SNMPv1 traps and Juniper Networks enterprise-specific SNMPv1 traps. See:

- [Table 596 on page 6140](#) for standard SNMPv1 traps.
- [Table 597 on page 6142](#) for enterprise-specific SNMPv1 traps.

The traps are organized first by trap category and then by trap name. The system logging severity levels are listed for those traps that have them. Traps that do not have corresponding system logging severity levels are marked with an en dash (–).

**Table 596: Standard SNMP Version 1 Traps Supported on QFX3500 and QFX3600 Switches**

| Defined in                                                                                         | Trap Name            | Enterprise ID    | Generic Trap Number | Specific Trap Number | System Logging Severity Level | Syslog Tag                        |
|----------------------------------------------------------------------------------------------------|----------------------|------------------|---------------------|----------------------|-------------------------------|-----------------------------------|
| <b>Link Notifications</b>                                                                          |                      |                  |                     |                      |                               |                                   |
| RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>                              | linkDown             | 1.3.6.1.4.1.2636 | 2                   | 0                    | Warning                       | SNMP_TRAP_LINK_DOWN               |
|                                                                                                    | linkUp               | 1.3.6.1.4.1.2636 | 3                   | 0                    | Info                          | SNMP_TRAP_LINK_UP                 |
| <b>Remote Operations Notifications</b>                                                             |                      |                  |                     |                      |                               |                                   |
| RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i> | pingProbeFailed      | 1.3.6.1.2.1.80.0 | 6                   | 1                    | Info                          | SNMP_TRAP_PING_PROBE_FAILED       |
|                                                                                                    | pingTestFailed       | 1.3.6.1.2.1.80.0 | 6                   | 2                    | Info                          | SNMP_TRAP_PING_TEST_FAILED        |
|                                                                                                    | pingTestCompleted    | 1.3.6.1.2.1.80.0 | 6                   | 3                    | Info                          | SNMP_TRAP_PING_TEST_COMPLETED     |
|                                                                                                    | traceRoutePathChange | 1.3.6.1.2.1.81.0 | 6                   | 1                    | Info                          | SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE |
|                                                                                                    | traceRouteTestFailed | 1.3.6.1.2.1.81.0 | 6                   | 2                    | Info                          | SNMP_TRAP_TRACE_ROUTE_TEST_FAILED |



**Table 596: Standard SNMP Version 1 Traps Supported on QFX3500 and QFX3600 Switches (continued)**

| Defined in                       | Trap Name               | Enterprise ID       | Generic Trap Number | Specific Trap Number | System Logging Severity Level | Syslog Tag                           |
|----------------------------------|-------------------------|---------------------|---------------------|----------------------|-------------------------------|--------------------------------------|
|                                  | traceRouteTestCompleted | 1.3.6.1.2.1.81.0    | 6                   | 3                    | Info                          | SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED |
| <b>RMON Alarms</b>               |                         |                     |                     |                      |                               |                                      |
| RFC 2819a, <i>RMON MIB</i>       | fallingAlarm            | 1.3.6.1.2.1.16      | 6                   | 2                    | —                             | —                                    |
|                                  | risingAlarm             | 1.3.6.1.2.1.16      | 6                   | 1                    | —                             | —                                    |
| <b>Routing Notifications</b>     |                         |                     |                     |                      |                               |                                      |
| <i>BGP 4 MIB</i>                 | bgpEstablished          | 1.3.6.1.2.1.15.7    | 6                   | 1                    | —                             | —                                    |
|                                  | bgpBackwardTransition   | 1.3.6.1.2.1.15.7    | 6                   | 2                    | —                             | —                                    |
| <i>OSPF TRAP MIB</i>             | ospfVirtIfStateChange   | 1.3.6.1.2.1.14.16.2 | 6                   | 1                    | —                             | —                                    |
|                                  | ospfNbrStateChange      | 1.3.6.1.2.1.14.16.2 | 6                   | 2                    | —                             | —                                    |
|                                  | ospfVirtNbrStateChange  | 1.3.6.1.2.1.14.16.2 | 6                   | 3                    | —                             | —                                    |
|                                  | ospfIfConfigError       | 1.3.6.1.2.1.14.16.2 | 6                   | 4                    | —                             | —                                    |
|                                  | ospfVirtIfConfigError   | 1.3.6.1.2.1.14.16.2 | 6                   | 5                    | —                             | —                                    |
|                                  | ospfIfAuthFailure       | 1.3.6.1.2.1.14.16.2 | 6                   | 6                    | —                             | —                                    |
|                                  | ospfVirtIfAuthFailure   | 1.3.6.1.2.1.14.16.2 | 6                   | 7                    | —                             | —                                    |
|                                  | ospfIfRxBadPacket       | 1.3.6.1.2.1.14.16.2 | 6                   | 8                    | —                             | —                                    |
|                                  | ospfVirtIfRxBadPacket   | 1.3.6.1.2.1.14.16.2 | 6                   | 9                    | —                             | —                                    |
|                                  | ospfTxRetransmit        | 1.3.6.1.2.1.14.16.2 | 6                   | 10                   | —                             | —                                    |
|                                  | ospfVirtIfTxRetransmit  | 1.3.6.1.2.1.14.16.2 | 6                   | 11                   | —                             | —                                    |
|                                  | ospfMaxAgeLsa           | 1.3.6.1.2.1.14.16.2 | 6                   | 13                   | —                             | —                                    |
|                                  | ospfIfStateChange       | 1.3.6.1.2.1.14.16.2 | 6                   | 16                   | —                             | —                                    |
| <b>Startup Notifications</b>     |                         |                     |                     |                      |                               |                                      |
| RFC 1215, <i>Conventions for</i> | authenticationFailure   | 1.3.6.1.4.1.2636    | 4                   | 0                    | Notice                        | SNMPD_TRAP_GEN_FAILURE               |

**Table 596: Standard SNMP Version 1 Traps Supported on QFX3500 and QFX3600 Switches (continued)**

| Defined in                                                                                 | Trap Name           | Enterprise ID    | Generic Trap Number | Specific Trap Number | System Logging Severity Level | Syslog Tag              |
|--------------------------------------------------------------------------------------------|---------------------|------------------|---------------------|----------------------|-------------------------------|-------------------------|
| <i>Defining Traps for Use with the SNMP</i>                                                | coldStart           | 1.3.6.1.4.1.2636 | 0                   | 0                    | Critical                      | SNMPD_TRAP_COLD_START   |
|                                                                                            | warmStart           | 1.3.6.1.4.1.2636 | 1                   | 0                    | Error                         | SNMPD_TRAP_WARM_START   |
| <b>VRRP Notifications</b>                                                                  |                     |                  |                     |                      |                               |                         |
| <i>RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i> | vrrpTrapNewMaster   | 1.3.6.1.2.1.68   | 6                   | 1                    | Warning                       | VRRPD_NEW_MASTER_TRAP   |
|                                                                                            | vrrpTrapAuthFailure | 1.3.6.1.2.1.68   | 6                   | 2                    | Warning                       | VRRPD_AUTH_FAILURE_TRAP |

**Table 597: Enterprise-Specific SNMPv1 Traps Supported on QFX3500 and QFX3600 Switches**

| Defined in                                      | Trap Name             | Enterprise ID        | Generic Trap Number | Specific Trap Number | System Logging Severity Level | System Log Tag     |
|-------------------------------------------------|-----------------------|----------------------|---------------------|----------------------|-------------------------------|--------------------|
| <b>Chassis Notifications (Alarm Conditions)</b> |                       |                      |                     |                      |                               |                    |
| <i>Chassis MIB (jnx-chassis. mib)</i>           | jnxPowerSupplyFailure | 1.3.6.1.4.1.2636.4.1 | 6                   | 1                    | Warning                       | CHASSISD_SNMP_TRAP |
|                                                 | jnxFanFailure         | 1.3.6.1.4.1.2636.4.1 | 6                   | 2                    | Critical                      | CHASSISD_SNMP_TRAP |
|                                                 | jnxOverTemperature    | 1.3.6.1.4.1.2636.4.1 | 6                   | 3                    | Alert                         | CHASSISD_SNMP_TRAP |
|                                                 | jnxFruRemoval         | 1.3.6.1.4.1.2636.4.1 | 6                   | 5                    | Notice                        | CHASSISD_SNMP_TRAP |
|                                                 | jnxFruInsertion       | 1.3.6.1.4.1.2636.4.1 | 6                   | 6                    | Notice                        | CHASSISD_SNMP_TRAP |
|                                                 | jnxFruPowerOff        | 1.3.6.1.4.1.2636.4.1 | 6                   | 7                    | Notice                        | CHASSISD_SNMP_TRAP |
|                                                 | jnxFruPowerOn         | 1.3.6.1.4.1.2636.4.1 | 6                   | 8                    | Notice                        | CHASSISD_SNMP_TRAP |

**Table 597: Enterprise-Specific SNMPv1 Traps Supported on QFX3500 and QFX3600 Switches (continued)**

| Defined in                                                  | Trap Name                             | Enterprise ID        | Generic Trap Number | Specific Trap Number | System Logging Severity Level | System Log Tag     |
|-------------------------------------------------------------|---------------------------------------|----------------------|---------------------|----------------------|-------------------------------|--------------------|
|                                                             | jnxFruFailed                          | 1.3.6.1.4.1.2636.4.1 | 6                   | 9                    | Warning                       | CHASSISD_SNMP_TRAP |
|                                                             | jnxFruOffline                         | 1.3.6.1.4.1.2636.4.1 | 6                   | 10                   | Notice                        | CHASSISD_SNMP_TRAP |
|                                                             | jnxFruOnline                          | 1.3.6.1.4.1.2636.4.1 | 6                   | 11                   | Notice                        | CHASSISD_SNMP_TRAP |
|                                                             | jnxFruCheck                           | 1.3.6.1.4.1.2636.4.1 | 6                   | 12                   | Warning                       | CHASSISD_SNMP_TRAP |
|                                                             | jnxPowerSupplyOk                      | 1.3.6.1.4.1.2636.4.2 | 6                   | 1                    | Critical                      | CHASSISD_SNMP_TRAP |
|                                                             | jnxFanOK                              | 1.3.6.1.4.1.2636.4.2 | 6                   | 2                    | Critical                      | CHASSISD_SNMP_TRAP |
|                                                             | jnxTemperatureOK                      | 1.3.6.1.4.1.2636.4.2 | 6                   | 3                    | Alert                         | CHASSISD_SNMP_TRAP |
| <b>Configuration Notifications</b>                          |                                       |                      |                     |                      |                               |                    |
| <i>Configuration Management MIB</i><br>(jnx-configmgmt.mib) | jnxCmCfgChange                        | 1.3.6.1.4.1.2636.4.5 | 6                   | 1                    | –                             | –                  |
|                                                             | jnxCmRescueChange                     | 1.3.6.1.4.1.2636.4.5 | 6                   | 2                    | –                             | –                  |
| <b>Remote Operations</b>                                    |                                       |                      |                     |                      |                               |                    |
| <i>Ping MIB</i><br>(jnx-ping.mib)                           | jnxPingRttThresholdExceeded           | 1.3.6.1.4.1.2636.4.9 | 6                   | 1                    | –                             | –                  |
|                                                             | jnxPingRttStdDevThreshold Exceeded    | 1.3.6.1.4.1.2636.4.9 | 6                   | 2                    | –                             | –                  |
|                                                             | jnxPingRttJitterThreshold Exceeded    | 1.3.6.1.4.1.2636.4.9 | 6                   | 3                    | –                             | –                  |
|                                                             | jnxPingEgressThreshold Exceeded       | 1.3.6.1.4.1.2636.4.9 | 6                   | 4                    | –                             | –                  |
|                                                             | jnxPingEgressStdDev ThresholdExceeded | 1.3.6.1.4.1.2636.4.9 | 6                   | 5                    | –                             | –                  |
|                                                             | jnxPingEgressJitterThreshold Exceeded | 1.3.6.1.4.1.2636.4.9 | 6                   | 6                    | –                             | –                  |

**Table 597: Enterprise-Specific SNMPv1 Traps Supported on QFX3500 and QFX3600 Switches (continued)**

| Defined in                        | Trap Name                             | Enterprise ID        | Generic Trap Number | Specific Trap Number | System Logging Severity Level | System Log Tag |
|-----------------------------------|---------------------------------------|----------------------|---------------------|----------------------|-------------------------------|----------------|
|                                   | jnxPingIngressThresholdExceeded       | 1.3.6.1.4.1.2636.4.9 | 6                   | 7                    | –                             | –              |
|                                   | jnxPingIngressStddevThresholdExceeded | 1.3.6.1.4.1.2636.4.9 | 6                   | 8                    | –                             | –              |
|                                   | jnxPingIngressJitterThresholdExceeded | 1.3.6.1.4.1.2636.4.9 | 6                   | 9                    | –                             | –              |
| <b>RMON Alarms</b>                |                                       |                      |                     |                      |                               |                |
| <i>RMON MIB</i><br>(jnx-rmon.mib) | jnxRmonAlarmGetFailure                | 1.3.6.1.4.1.2636.4.3 | 6                   | 1                    | –                             | –              |
|                                   | jnxRmonGetOk                          | 1.3.6.1.4.1.2636.4.3 | 6                   | 2                    | –                             | –              |

**SNMPv2 Traps**

- [Table 598 on page 6144](#) lists the standard SNMP traps
- [Table 599 on page 6147](#) lists the Juniper Networks enterprise-specific traps

**Table 598: Standard SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches**

| Defined in                                                                                         | Trap Name         | SNMP Trap OID       | System Logging Severity Level | Syslog Tag                    |
|----------------------------------------------------------------------------------------------------|-------------------|---------------------|-------------------------------|-------------------------------|
| <b>Link Notifications</b>                                                                          |                   |                     |                               |                               |
| RFC 2863, <i>The Interfaces Group MIB</i>                                                          | linkDown          | 1.3.6.1.6.3.1.1.5.3 | Warning                       | SNMP_TRAP_LINK_DOWN           |
|                                                                                                    | linkUp            | 1.3.6.1.6.3.1.1.5.4 | Info                          | SNMP_TRAP_LINK_UP             |
| <b>Remote Operations Notifications</b>                                                             |                   |                     |                               |                               |
| RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i> | pingProbeFailed   | 1.3.6.1.2.1.80.0.1  | Info                          | SNMP_TRAP_PING_PROBE_FAILED   |
|                                                                                                    | pingTestFailed    | 1.3.6.1.2.1.80.0.2  | Info                          | SNMP_TRAP_PING_TEST_FAILED    |
|                                                                                                    | pingTestCompleted | 1.3.6.1.2.1.80.0.3  | Info                          | SNMP_TRAP_PING_TEST_COMPLETED |

**Table 598: Standard SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches** (*continued*)

| Defined in                   | Trap Name               | SNMP Trap OID      | System Logging Severity Level | Syslog Tag                           |
|------------------------------|-------------------------|--------------------|-------------------------------|--------------------------------------|
|                              | traceRoutePathChange    | 1.3.6.1.2.1.81.0.1 | Info                          | SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE    |
|                              | traceRouteTestFailed    | 1.3.6.1.2.1.81.0.2 | Info                          | SNMP_TRAP_TRACE_ROUTE_TEST_FAILED    |
|                              | traceRouteTestCompleted | 1.3.6.1.2.1.81.0.3 | Info                          | SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED |
| <b>RMON Alarms</b>           |                         |                    |                               |                                      |
| RFC 2819a, <i>RMON MIB</i>   | fallingAlarm            | 1.3.6.1.2.1.16.0.1 | —                             | —                                    |
|                              | risingAlarm             | 1.3.6.1.2.1.16.0.2 | —                             | —                                    |
| <b>Routing Notifications</b> |                         |                    |                               |                                      |
| <i>BGP 4 MIB</i>             | bgpEstablished          | 1.3.6.1.2.1.15.7.1 | —                             | —                                    |
|                              | bgpBackwardTransition   | 1.3.6.1.2.1.15.7.2 | —                             | —                                    |

**Table 598: Standard SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches (continued)**

| Defined in                                                                                                    | Trap Name              | SNMP Trap OID          | System Logging Severity Level | Syslog Tag              |
|---------------------------------------------------------------------------------------------------------------|------------------------|------------------------|-------------------------------|-------------------------|
| <i>OSPF Trap MIB</i>                                                                                          | ospfVirtIfStateChange  | 1.3.6.1.2.1.14.16.2.1  | –                             | –                       |
|                                                                                                               | ospfNbrStateChange     | 1.3.6.1.2.1.14.16.2.2  | –                             | –                       |
|                                                                                                               | ospfVirtNbrStateChange | 1.3.6.1.2.1.14.16.2.3  | –                             | –                       |
|                                                                                                               | ospfIfConfigError      | 1.3.6.1.2.1.14.16.2.4  | –                             | –                       |
|                                                                                                               | ospfVirtIfConfigError  | 1.3.6.1.2.1.14.16.2.5  | –                             | –                       |
|                                                                                                               | ospfIfAuthFailure      | 1.3.6.1.2.1.14.16.2.6  | –                             | –                       |
|                                                                                                               | ospfVirtIfAuthFailure  | 1.3.6.1.2.1.14.16.2.7  | –                             | –                       |
|                                                                                                               | ospfIfRxBadPacket      | 1.3.6.1.2.1.14.16.2.8  | –                             | –                       |
|                                                                                                               | ospfVirtIfRxBadPacket  | 1.3.6.1.2.1.14.16.2.9  | –                             | –                       |
|                                                                                                               | ospfTxRetransmit       | 1.3.6.1.2.1.14.16.2.10 | –                             | –                       |
|                                                                                                               | ospfVirtIfTxRetransmit | 1.3.6.1.2.1.14.16.2.11 | –                             | –                       |
|                                                                                                               | ospfMaxAgeLsa          | 1.3.6.1.2.1.14.16.2.13 | –                             | –                       |
|                                                                                                               | ospfIfStateChange      | 1.3.6.1.2.1.14.16.2.16 | –                             | –                       |
| <b>Startup Notifications</b>                                                                                  |                        |                        |                               |                         |
| RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> | coldStart              | 1.3.6.1.6.3.1.1.5.1    | Critical                      | SNMPD_TRAP_COLD_START   |
|                                                                                                               | warmStart              | 1.3.6.1.6.3.1.1.5.2    | Error                         | SNMPD_TRAP_WARM_START   |
|                                                                                                               | authenticationFailure  | 1.3.6.1.6.3.1.1.5.5    | Notice                        | SNMPD_TRAP_GEN_FAILURE  |
| <b>VRRP Notifications</b>                                                                                     |                        |                        |                               |                         |
| RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>                    | vrrpTrapNewMaster      | 1.3.6.1.2.1.68.0.1     | Warning                       | VRRPD_NEWMASTER_TRAP    |
|                                                                                                               | vrrpTrapAuthFailure    | 1.3.6.1.2.1.68.0.2     | Warning                       | VRRPD_AUTH_FAILURE_TRAP |

Table 599: Enterprise-Specific SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches

| Source MIB                                      | Trap Name             | SNMP Trap OID           | System Logging Severity Level | System Log Tag       |
|-------------------------------------------------|-----------------------|-------------------------|-------------------------------|----------------------|
| <b>Chassis (Alarm Conditions) Notifications</b> |                       |                         |                               |                      |
| <i>Chassis MIB</i><br>(mib-jnx-chassis)         | jnxPowerSupplyFailure | 1.3.6.1.4.1.2636.4.1.1  | Alert                         | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFanFailure         | 1.3.6.1.4.1.2636.4.1.2  | Critical                      | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxOverTemperature    | 1.3.6.1.4.1.2636.4.1.3  | Critical                      | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruRemoval         | 1.3.6.1.4.1.2636.4.1.5  | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruInsertion       | 1.3.6.1.4.1.2636.4.1.6  | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruPowerOff        | 1.3.6.1.4.1.2636.4.1.7  | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruPowerOn         | 1.3.6.1.4.1.2636.4.1.8  | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruFailed          | 1.3.6.1.4.1.2636.4.1.9  | Warning                       | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruOffline         | 1.3.6.1.4.1.2636.4.1.10 | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruOnline          | 1.3.6.1.4.1.2636.4.1.11 | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFruCheck           | 1.3.6.1.4.1.2636.4.1.12 | Notice                        | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxPowerSupplyOK      | 1.3.6.1.4.1.2636.4.2.1  | Critical                      | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxFanOK              | 1.3.6.1.4.1.2636.4.2.2  | Critical                      | CHASSISD_ SNMP_ TRAP |
|                                                 | jnxTemperatureOK      | 1.3.6.1.4.1.2636.4.2.3  | Alert                         | CHASSISD_ SNMP_ TRAP |
| <b>Configuration Notifications</b>              |                       |                         |                               |                      |

Table 599: Enterprise-Specific SNMPv2 Traps Supported on QFX3500 and QFX3600 Switches (*continued*)

| Source MIB                                               | Trap Name                              | SNMP Trap OID            | System Logging Severity Level | System Log Tag |
|----------------------------------------------------------|----------------------------------------|--------------------------|-------------------------------|----------------|
| <i>Configuration Management MIB</i><br>(mib-jnx-cfgmgmt) | jnxCmCfgChange                         | 1.3.6.1.4.1.2636.4.5.0.1 | –                             | –              |
|                                                          | jnxCmRescueChange                      | 1.3.6.1.4.1.2636.4.5.0.2 | –                             | –              |
| <b>Remote Operations Notifications</b>                   |                                        |                          |                               |                |
| <i>Ping MIB</i><br>(mib-jnx-ping)                        | jnxPingRttThreshold Exceeded           | 1.3.6.1.4.1.2636.4.9.0.1 | –                             | –              |
|                                                          | jnxPingRttStdDevThreshold Exceeded     | 1.3.6.1.4.1.2636.4.9.0.2 | –                             | –              |
|                                                          | jnxPingRttJitterThreshold Exceeded     | 1.3.6.1.4.1.2636.4.9.0.3 | –                             | –              |
|                                                          | jnxPingEgressThreshold Exceeded        | 1.3.6.1.4.1.2636.4.9.0.4 | –                             | –              |
|                                                          | jnxPingEgressStdDevThreshold Exceeded  | 1.3.6.1.4.1.2636.4.9.0.5 | –                             | –              |
|                                                          | jnxPingEgressJitterThreshold Exceeded  | 1.3.6.1.4.1.2636.4.9.0.6 | –                             | –              |
|                                                          | jnxPingIngressThreshold Exceeded       | 1.3.6.1.4.1.2636.4.9.0.7 | –                             | –              |
|                                                          | jnxPingIngressStddevThreshold Exceeded | 1.3.6.1.4.1.2636.4.9.0.8 | –                             | –              |
|                                                          | jnxPingIngressJitterThreshold Exceeded | 1.3.6.1.4.1.2636.4.9.0.9 | –                             | –              |
| <b>RMON Alarms</b>                                       |                                        |                          |                               |                |
| <i>RMON MIB</i><br>(mib-jnx-rmon)                        | jnxRmonAlarmGetFailure                 | 1.3.6.1.4.1.2636.4.3.0.1 | –                             | –              |
|                                                          | jnxRmonGetOk                           | 1.3.6.1.4.1.2636.4.3.0.2 | –                             | –              |

### SNMP Traps Supported on QFabric Systems

QFabric systems support standard SNMPv2 traps and Juniper Networks enterprise-specific SNMPv2 traps.





**NOTE:** QFabric systems do not support SNMPv1 traps.

For more information, see:

- [Table 600 on page 6149](#) for standard SNMPv2 traps
- [Table 601 on page 6150](#) for Juniper Networks enterprise-specific SNMPv2 traps

**Table 600: Standard SNMPv2 Traps Supported on QFabric Systems**

| Defined in                                                                                                    | Trap Name             | SNMP Trap OID       | System Logging Severity Level | Syslog Tag             |
|---------------------------------------------------------------------------------------------------------------|-----------------------|---------------------|-------------------------------|------------------------|
| <b>Link Notifications</b>                                                                                     |                       |                     |                               |                        |
| RFC 2863, <i>The Interfaces Group MIB</i>                                                                     | linkDown              | 1.3.6.1.6.3.1.1.5.3 | Warning                       | SNMP_TRAP_LINK_DOWN    |
|                                                                                                               | linkUp                | 1.3.6.1.6.3.1.1.5.4 | Info                          | SNMP_TRAP_LINK_UP      |
| <b>Startup Notifications</b>                                                                                  |                       |                     |                               |                        |
| RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> | coldStart             | 1.3.6.1.6.3.1.1.5.1 | Critical                      | SNMPD_TRAP_COLD_START  |
|                                                                                                               | warmStart             | 1.3.6.1.6.3.1.1.5.2 | Error                         | SNMPD_TRAP_WARM_START  |
|                                                                                                               | authenticationFailure | 1.3.6.1.6.3.1.1.5.5 | Notice                        | SNMPD_TRAP_GEN_FAILURE |

Table 601: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems

| Source MIB                                            | Trap Name                                                      | SNMP Trap OID            | System Logging Severity Level | System Log Tag |
|-------------------------------------------------------|----------------------------------------------------------------|--------------------------|-------------------------------|----------------|
| <i>Fabric Chassis MIB</i><br>(mib-jnx-fabric-chassis) | <b>Fabric Chassis (Alarm Conditions) Notifications</b>         |                          |                               |                |
|                                                       | jnxFabricPowerSupplyFailure                                    | 1.3.6.1.4.1.2636.4.19.1  | Warning                       | –              |
|                                                       | jnxFabricFanFailure                                            | 1.3.6.1.4.1.2636.4.19.2  | Critical                      | –              |
|                                                       | jnxFabricOverTemperature                                       | 1.3.6.1.4.1.2636.4.19.3  | Alert                         | –              |
|                                                       | jnxFabricRedundancySwitchover                                  | 1.3.6.1.4.1.2636.4.19.4  | Notice                        | –              |
|                                                       | jnxFabricFruRemoval                                            | 1.3.6.1.4.1.2636.4.19.5  | Notice                        | –              |
|                                                       | jnxFabricFruInsertion                                          | 1.3.6.1.4.1.2636.4.19.6  | Notice                        | –              |
|                                                       | jnxFabricFruPowerOff                                           | 1.3.6.1.4.1.2636.4.19.7  | Notice                        | –              |
|                                                       | jnxFabricFruPowerOn                                            | 1.3.6.1.4.1.2636.4.19.8  | Notice                        | –              |
|                                                       | jnxFabricFruFailed                                             | 1.3.6.1.4.1.2636.4.19.9  | Warning                       | –              |
|                                                       | jnxFabricFruOffline                                            | 1.3.6.1.4.1.2636.4.19.10 | Notice                        | –              |
|                                                       | jnxFabricFruOnline                                             | 1.3.6.1.4.1.2636.4.19.11 | Notice                        | –              |
|                                                       | jnxFabricFruCheck                                              | 1.3.6.1.4.1.2636.4.19.12 | Warning                       | –              |
|                                                       | jnxFabricFEBSwitchover                                         | 1.3.6.1.4.1.2636.4.19.13 | Warning                       | –              |
|                                                       | jnxFabricHardDiskFailed                                        | 1.3.6.1.4.1.2636.4.19.14 | Warning                       | –              |
|                                                       | jnxFabricHardDiskMissing                                       | 1.3.6.1.4.1.2636.4.19.15 | Warning                       | –              |
|                                                       | jnxFabricBootFromBackup                                        | 1.3.6.1.4.1.2636.4.19.16 | Warning                       | –              |
|                                                       | <b>Fabric Chassis (Alarm Cleared Conditions) Notifications</b> |                          |                               |                |
|                                                       | jnxFabricPowerSupplyOK                                         | 1.3.6.1.4.1.2636.4.20.1  | Critical                      | –              |
|                                                       | jnxFabricFanOK                                                 | 1.3.6.1.4.1.2636.4.20.2  | Critical                      | –              |
|                                                       | jnxFabricTemperatureOK                                         | 1.3.6.1.4.1.2636.4.20.3  | Alert                         | –              |
|                                                       | jnxFabricFruOK                                                 | 1.3.6.1.4.1.2636.4.20.4  | –                             | –              |

Table 601: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems (*continued*)

| Source MIB                                               | Trap Name                              | SNMP Trap OID               | System Logging Severity Level | System Log Tag |
|----------------------------------------------------------|----------------------------------------|-----------------------------|-------------------------------|----------------|
| <i>QFabric MIB</i><br>(mib-jnx-qf-smi)                   | <b>QFabric MIB Notifications</b>       |                             |                               |                |
|                                                          | jnxQFabricDownloadIssued               | 1.3.6.1.4.1.2636.3.42.1.0.1 | –                             | –              |
|                                                          | jnxQFabricDownloadFailed               | 1.3.6.1.4.1.2636.3.42.1.0.2 | –                             | –              |
|                                                          | jnxQFabricDownloadSucceeded            | 1.3.6.1.4.1.2636.3.42.1.0.3 | –                             | –              |
|                                                          | jnxQFabricUpgradeIssued                | 1.3.6.1.4.1.2636.3.42.1.0.4 | –                             | –              |
|                                                          | jnxQFabricUpgradeFailed                | 1.3.6.1.4.1.2636.3.42.1.0.5 | –                             | –              |
|                                                          | jnxQFabricUpgradeSucceeded             | 1.3.6.1.4.1.2636.3.42.1.0.6 | –                             | –              |
| <b>Configuration Notifications</b>                       |                                        |                             |                               |                |
| <i>Configuration Management MIB</i><br>(mib-jnx-cfgmgmt) | jnxCmCfgChange                         | 1.3.6.1.4.1.2636.4.5.0.1    | –                             | –              |
|                                                          | jnxCmRescueChange                      | 1.3.6.1.4.1.2636.4.5.0.2    | –                             | –              |
| <b>Remote Operations Notifications</b>                   |                                        |                             |                               |                |
| <i>Ping MIB</i><br>(mib-jnx-ping)                        | jnxPingRttThreshold Exceeded           | 1.3.6.1.4.1.2636.4.9.0.1    | –                             | –              |
|                                                          | jnxPingRttStdDevThreshold Exceeded     | 1.3.6.1.4.1.2636.4.9.0.2    | –                             | –              |
|                                                          | jnxPingRttJitterThreshold Exceeded     | 1.3.6.1.4.1.2636.4.9.0.3    | –                             | –              |
|                                                          | jnxPingEgressThreshold Exceeded        | 1.3.6.1.4.1.2636.4.9.0.4    | –                             | –              |
|                                                          | jnxPingEgressStdDevThreshold Exceeded  | 1.3.6.1.4.1.2636.4.9.0.5    | –                             | –              |
|                                                          | jnxPingEgressJitterThreshold Exceeded  | 1.3.6.1.4.1.2636.4.9.0.6    | –                             | –              |
|                                                          | jnxPingIngressThreshold Exceeded       | 1.3.6.1.4.1.2636.4.9.0.7    | –                             | –              |
|                                                          | jnxPingIngressStddevThreshold Exceeded | 1.3.6.1.4.1.2636.4.9.0.8    | –                             | –              |
|                                                          | jnxPingIngressJitterThreshold Exceeded | 1.3.6.1.4.1.2636.4.9.0.9    | –                             | –              |

- Related Documentation**
- [SNMP MIBs and Traps Reference](#)
  - [Understanding the Implementation of SNMP on page 6107](#)
  - [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
  - [SNMP MIBs Support on page 6124](#)

## MIB Objects for the QFX Series

This topic lists the Juniper Networks enterprise-specific SNMP Chassis MIB definition objects for the QFX Series:

- [QFX3500 and QFX3600 Switches on page 6152](#)
- [QFabric Systems on page 6152](#)
- [QFabric System QFX3100 Director Device on page 6153](#)
- [QFabric System QFX3008-I Interconnect Device on page 6153](#)
- [QFabric System QFX3600-I Interconnect Device on page 6153](#)
- [QFabric System Node Devices on page 6154](#)

### QFX3500 and QFX3600 Switches

|                               |                                                          |
|-------------------------------|----------------------------------------------------------|
| jnxProductLineQFXSwitch       | OBJECT IDENTIFIER ::= { jnxProductLine 82 }              |
| jnxProductNameQFXSwitch       | OBJECT IDENTIFIER ::= { jnxProductName 82 }              |
| jnxProductModelQFXSwitch      | OBJECT IDENTIFIER ::= { jnxProductModel 82 }             |
| jnxProductVariationQFXSwitch  | OBJECT IDENTIFIER ::= { jnxProductVariation 82 }         |
| jnxProductQFX3500s            | OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 1 } |
| jnxProductQFX360016QS         | OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 2 } |
| jnxProductQFX350048T4QS       | OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 3 } |
| jnxChassisQFXSwitch           | OBJECT IDENTIFIER ::= { jnxChassis 82 }                  |
| jnxSlotQFXSwitch              | OBJECT IDENTIFIER ::= { jnxSlot 82 }                     |
| jnxQFXSwitchSlotFPC           | OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 1 }             |
| jnxQFXSwitchSlotHM            | OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 2 }             |
| jnxQFXSwitchSlotPower         | OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 3 }             |
| jnxQFXSwitchSlotFan           | OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 4 }             |
| jnxQFXSwitchSlotFPB           | OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 5 }             |
| jnxMediaCardSpaceQFXSwitch    | OBJECT IDENTIFIER ::= { jnxMediaCardSpace 82 }           |
| jnxQFXSwitchMediaCardSpacePIC | OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXSwitch 1 }   |

### QFabric Systems

|                            |                                                        |
|----------------------------|--------------------------------------------------------|
| jnxProductLineQFX3000      | OBJECT IDENTIFIER ::= { jnxProductLine 84 }            |
| jnxProductNameQFX3000      | OBJECT IDENTIFIER ::= { jnxProductName 84 }            |
| jnxProductModelQFX3000     | OBJECT IDENTIFIER ::= { jnxProductModel 84 }           |
| jnxProductVariationQFX3000 | OBJECT IDENTIFIER ::= { jnxProductVariation 84 }       |
| jnxProductQFX3000-G        | OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 1 } |
| jnxProductQFX3000-M        | OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 2 } |
| jnxChassisQFX3000          | OBJECT IDENTIFIER ::= { jnxChassis 84 }                |

### QFabric System QFX3100 Director Device

```
jnxProductLineQFX3100 OBJECT IDENTIFIER ::= { jnxProductLine 100 }
jnxProductNameQFX3100 OBJECT IDENTIFIER ::= { jnxProductName 100 }
jnxProductModelQFX3100 OBJECT IDENTIFIER ::= { jnxProductModel 100 }
jnxProductVariationQFX3100 OBJECT IDENTIFIER ::= { jnxProductVariation 100 }
jnxChassisQFX3100 OBJECT IDENTIFIER ::= { jnxChassis 100 }

jnxSlotQFX3100 OBJECT IDENTIFIER ::= { jnxSlot 100 }
jnxQFX3100SlotCPU OBJECT IDENTIFIER ::= { jnxSlotQFX3100 1 }
jnxQFX3100SlotMemory OBJECT IDENTIFIER ::= { jnxSlotQFX3100 2 }
jnxQFX3100SlotPower OBJECT IDENTIFIER ::= { jnxSlotQFX3100 3 }
jnxQFX3100SlotFan OBJECT IDENTIFIER ::= { jnxSlotQFX3100 4 }
jnxQFX3100SlotHardDisk OBJECT IDENTIFIER ::= { jnxSlotQFX3100 5 }
jnxQFX3100SlotNIC OBJECT IDENTIFIER ::= { jnxSlotQFX3100 6 }
```

### QFabric System QFX3008-I Interconnect Device

```
jnxProductLineQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductLine 60 }
jnxProductNameQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductName 60 }
jnxProductModelQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductModel 60 }
jnxProductVariationQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 60 }
jnxProductQFX3008 OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 1 }
jnxProductQFXC083008 OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 2 }
jnxProductQFX3008I OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 3 }

jnxChassisQFXInterconnect OBJECT IDENTIFIER ::= { jnxChassis 60 }

jnxSlotQFXInterconnect OBJECT IDENTIFIER ::= { jnxSlot 60 }
jnxQFXInterconnectSlotFPC OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 1 }
jnxQFXInterconnectSlotHM OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 2 }
jnxQFXInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 3 }
jnxQFXInterconnectSlotFan OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 4 }
jnxQFXInterconnectSlotCBD OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 5 }
jnxQFXInterconnectSlotFPB OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect 6 }

jnxMediaCardSpaceQFXInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 60 }
jnxQFXInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXInterconnect 1 }

jnxMidplaneQFXInterconnect OBJECT IDENTIFIER ::= { jnxBackplane 60 }
```

### QFabric System QFX3600-I Interconnect Device

```
jnxProductLineQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductLine 91 }
jnxProductNameQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductName 91 }
jnxProductModelQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductModel 91 }
jnxProductVariationQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 91 }
jnxProductQFX3600I OBJECT IDENTIFIER ::= { jnxProductVariationQFXMInterconnect 1 }

jnxChassisQFXMInterconnect OBJECT IDENTIFIER ::= { jnxChassis 91 }

jnxSlotQFXMInterconnect OBJECT IDENTIFIER ::= { jnxSlot 91 }
jnxQFXMInterconnectSlotFPC OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect 1 }
jnxQFXMInterconnectSlotHM OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect 2 }
jnxQFXMInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect 3 }
jnxQFXMInterconnectSlotFan OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect 4 }
jnxQFXMInterconnectSlotFPB OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect 5 }
```

```
jnxMediaCardSpaceQFXMInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 91 }
jnxQFXMInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXMInterconnect 1 }
```

### QFabric System Node Devices

```
jnxProductLineQFXNode OBJECT IDENTIFIER ::= { jnxProductLine 61 }
jnxProductNameQFXNode OBJECT IDENTIFIER ::= { jnxProductName 61 }
jnxProductModelQFXNode OBJECT IDENTIFIER ::= { jnxProductModel 61 }
jnxProductVariationQFXNode OBJECT IDENTIFIER ::= { jnxProductVariation 61 }
jnxProductQFX3500 OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 1 }
jnxProductQFX360016Q OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 3 }

jnxChassisQFXNode OBJECT IDENTIFIER ::= { jnxChassis 61 }

jnxSlotQFXNode OBJECT IDENTIFIER ::= { jnxSlot 61 }
jnxQFXNodeSlotFPC OBJECT IDENTIFIER ::= { jnxSlotQFXNode 1 }
jnxQFXNodeSlotHM OBJECT IDENTIFIER ::= { jnxSlotQFXNode 2 }
jnxQFXNodeSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXNode 3 }
jnxQFXNodeSlotFan OBJECT IDENTIFIER ::= { jnxSlotQFXNode 4 }
jnxQFXNodeSlotFPB OBJECT IDENTIFIER ::= { jnxSlotQFXNode 5 }

jnxMediaCardSpaceQFXNode OBJECT IDENTIFIER ::= { jnxMediaCardSpace 61 }
jnxQFXNodeMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXNode 1 }
```

- Related Documentation**
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
  - [Fabric Chassis MIB on page 6112](#)

## System Logging

- [Overview of Junos OS System Log Messages on page 6154](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Understanding the Implementation of System Log Messages on the QFabric System on page 6156](#)

### Overview of Junos OS System Log Messages

The Junos OS, running on the QFX Series, generates system log messages (also called *syslog messages*) to record events that occur on the switch, including the following:

- Routine operations, such as a user login into the configuration database.
- Failure and error conditions, such as failure to access a configuration file.
- Emergency or critical conditions, such as power-down of the switch due to excessive temperature.

Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred. For detailed information about specific system log messages, see the *Junos OS System Log Messages Reference*.

- Related Documentation**
- [Junos OS System Log Configuration Statements on page 6208](#)
  - [Junos OS Minimum System Logging Configuration on page 6207](#)

## Overview of Single-Chassis System Logging Configuration

The Junos OS system logging utility on the QFX Series is similar to the UNIX **syslogd** utility. This topic describes how to configure system logging for a single-chassis system that runs the Junos OS.

Each system log message belongs to a *facility*, which groups together related messages. Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects router functions. You always specify the facility and severity of the messages to include in the log. For more information, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6222](#).

You direct messages to one or more destinations by including the appropriate statement at the **[edit system syslog]** hierarchy level:

- To a named file in a local file system, by including the **file** statement. See [“Directing System Log Messages to a Log File” on page 6209](#).
- To the terminal session of one or more specific users (or all users) when they are logged in to the switch, by including the **user** statement. See [“Directing System Log Messages to a User Terminal” on page 6211](#).
- To the switch console, by including the **console** statement. See [“Directing System Log Messages to the Console” on page 6211](#).
- To a remote machine that is running the **syslogd** utility, by including the **host** statement. See [“Directing System Log Messages to a Remote Machine” on page 6210](#).

By default, messages are logged in a standard format, which is based on a UNIX system log format; for detailed information about message formatting, see the *Junos OS System Log Messages Reference*. You can alter the content and format of logged messages in the following ways:

- You can log messages to a file in structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from the message. For more information, see [“Logging Messages in Structured-Data Format” on page 6215](#).
- A message’s facility and severity level are together referred to as its *priority*. By default, the standard Junos OS format for messages does not include priority information (structured-data format includes a priority code by default). To include priority information in standard-format messages directed to a file or a remote destination, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 6213](#).
- By default, the standard Junos OS format for messages specifies the month, date, hour, minute, and second when the message was logged. You can modify the timestamp on standard-format system log messages to include the year, the millisecond, or both. (Structured-data format specifies the year and millisecond by default.) For more information, see [“Including the Year or Millisecond in Timestamps” on page 6215](#).

- When directing messages to a remote machine, you can specify the IP address that is reported in messages as their source. You can also configure features that make it easier to separate messages generated by Junos OS or messages generated on particular switches. For more information, see [“Directing System Log Messages to a Remote Machine” on page 6210](#).
- The predefined facilities group together related messages, but you can also use regular expressions to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination. For more information, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6227](#).



**NOTE:** During a commit check, warnings about the `traceoptions` configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

#### Related Documentation

- [Examples: Configuring System Logging on page 6159](#)
- [Specifying the Facility and Severity of Messages to Include in the Log on page 6222](#)
- [Junos OS System Logging Facilities and Message Severity Levels on page 6222](#)
- [Directing System Log Messages to a Log File on page 6209](#)
- [Directing System Log Messages to a Remote Machine on page 6210](#)
- [Directing System Log Messages to a User Terminal on page 6211](#)
- [Directing System Log Messages to the Console on page 6211](#)

## Understanding the Implementation of System Log Messages on the QFabric System

This topic provides an overview of system log (syslog) messages as implemented on the QFabric system.

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

You configure system log messages by using the **host** and **file** statements at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view the messages.





**NOTE:** On the QFabric system, a syslog file named **messages** with a size of 100 MB is configured by default. If you do not configure a filename, you can use the default filename **messages** with the **show log filename** command.

All messages with a severity level of **notice** or higher are logged. Messages with a facility level of **interactive-commands** on Node devices are not logged.

The QFabric system supports the following system log message features:

- The **file filename** and **host hostname** statements at the **[edit system syslog]** hierarchy level are supported. Other statements at that hierarchy level are not supported.
- You can specify the maximum amount of data that is displayed when you issue the **show log filename** command by configuring the **file filename archive maximum-file-size** statement.
- You can specify that one or more system log message servers receive messages, which are sent to each server that is configured.
- If you configured an alias for a device or interface, the alias is displayed in the message for the device or interface.
- The level of detail that is included in a message depends on the facility and severity levels that are configured. Messages include the highest level of detail available for the configured facility and severity levels.
- The unit of time is measured and displayed in seconds, and not milliseconds. If you attempt to configure the **time-format** option in milliseconds, the log output displays **000**.

Starting in Junos OS Release 13.1, the QFabric system supports these additional syslog features:

- You can filter the output of the **show log filename** operational mode command by device type and device ID or device alias when you specify the **device-type (device-id | device-alias)** optional parameters. Device types include **director-device**, **infrastructure-device**, **interconnect-device**, and **node-device**.
- You can specify the syslog structured data output format when you configure the **structured-data** statement at the **[edit system syslog file filename]** and **[edit system syslog host hostname]** hierarchy levels.



**NOTE:** Information displayed in the structured data output for system logs originating from the Director software may not be complete.

- You can filter the types of logs that the Director group collects from a component device when you configure the **filter all facility severity** or **filter all match "regular-expression"** statements at the **[edit system syslog]** hierarchy level.

Unsupported syslog features include:

- File access to syslog messages
- Monitoring of syslog messages

**Related  
Documentation**

- [Example: Configuring System Log Messages on page 1372](#)
- [syslog \(QFabric System\) on page 1403](#)

## CHAPTER 68

# Configuration

- [Configuration Examples on page 6159](#)
- [Configuration Tasks for Network Management on page 6177](#)
- [Configuration Tasks for Automation Scripts on page 6181](#)
- [Configuration Tasks for Fabric OAM on page 6183](#)
- [Configuration Tasks for sFlow Technology on page 6188](#)
- [Configuration Tasks for SNMP on page 6190](#)
- [Configuration Tasks for System Log Messages on page 6207](#)
- [Configuration Statements for Network Management on page 6229](#)
- [Configuration Statements for Automation Scripts on page 6236](#)
- [Configuration Statements for Fabric OAM on page 6254](#)
- [Configuration Statements for sFlow Technology on page 6267](#)
- [Configuration Statements for SNMP on page 6274](#)
- [Configuration Statements for System Log Messages on page 6364](#)

### Configuration Examples

---

- [Examples: Configuring System Logging on page 6159](#)
- [Examples: Assigning an Alternative Facility on page 6161](#)
- [Example: Configuring System Log Messages on page 6162](#)
- [Example: Monitoring Network Traffic Using sFlow Technology on page 6165](#)
- [Example: Configuring SNMP on page 6169](#)
- [Example: Configuring Internal Fabric OAM Monitoring on page 6171](#)

### Examples: Configuring System Logging

The system log provides an excellent way of tracking all management activity on the switch by recording events such as user authentication, access authorization, and command execution. Logged command executions include commands entered by users at the CLI prompt or by client applications such as the Junos XML protocol or NETCONF XML client. Because system log files contain information about commands executed on the switch and the user who executed the commands, checking system log files for failed authentication events can help identify attempts to hack in to the switch. You can also

analyze network activity by correlating executed commands with events and changes that occurred on the network at a particular time.

System log files are stored locally on the switch in the default **/var/log** directory.

The following example shows how to configure system log messages to record all commands entered by users and all authentication or authorization attempts. Logged commands include those entered by users at the CLI prompt and by client applications. Authentication and authorization attempts include events that are saved in the file named **cli-commands** and those that are sent to the terminal of a user who is logged in.

```
[edit system]
syslog {
 file cli-commands {
 interactive-commands info;
 authorization info;
 }
 user * {
 interactive-commands info;
 authorization info;
 }
}
```

The following example shows how to log all alarms state changes to the file **/var/log/alarms**:

```
[edit system]
syslog {
 file alarms {
 kernel warning;
 }
}
```

The following example shows how to configure the handling of messages of various types, as described in the comments. Information is logged to two files, to the terminal of user **alex**, to a remote machine, and to the console:

```
[edit system]
syslog {
 /* write all security-related messages to file /var/log/security */
 file security {
 authorization info;
 interactive-commands info;
 }
 /* write messages about potential problems to file /var/log/messages: */
 /* messages from "authorization" facility at level "notice" and above, */
 /* messages from all other facilities at level "warning" and above */
 file messages {
 authorization notice;
 any warning;
 }
 /* write all messages at level "critical" and above to terminal of user
 "alex" if */
 /* that user is logged in */
 user alex {
```

```

 any critical;
 }
 /* write all messages from the “daemon” facility at level “info”
 and above, and */
 /* messages from all other facilities at level “warning” and above, to the
 */
 /* machine monitor.mycompany.com */
 host monitor.mycompany.com {
 daemon info;
 any warning;
 }
 /* write all messages at level “error” and above to the system console */
 console {
 any error;
 }
}

```

The following example shows how to configure the handling of messages generated when users issue Junos OS CLI commands, by specifying the **interactive-commands** facility at the **info**, **notice**, and **warning** severity levels:

```

[edit system]
file user-actions {
 interactive-commands info;
}
user philip {
 interactive-commands notice;
}
console {
 interactive-commands warning;
}
}

```

The following list describes the security levels used in the example:

- **info**—Logs a message when users issue any command at the CLI operational or configuration mode prompt. The example writes the messages to the file `/var/log/user-actions`.
- **notice**—Logs a message when users issue the configuration mode commands **rollback** and **commit**. The example writes the messages to the terminal of user **philip**.
- **warning**—Logs a message when users issue a command that restarts a software process. The example writes the messages to the console.

#### Related Documentation

- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

### Examples: Assigning an Alternative Facility

This topic contains examples of configuring system log messages to use an alternative facility for logging.

The following example shows how to log all messages generated on the switch at the **error** level or higher to the **local0** facility on the remote host called **monitor.mycompany.com**:

```
[edit system syslog]
host monitor.mycompany.com {
 any error;
 facility-override local0;
}
```

The following example contains two sets of statements that show how to configure switches located in California and in New York to send messages to a single remote host called **central-logger.mycompany.com**. The messages from California are assigned to alternative facility **local0** and the messages from New York are assigned to alternative facility **local2**.

- The following statements configure the California switch to aggregate messages in the **local0** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
 change-log info;
 facility-override local0;
}
```

- The following statements configure the New York switch to aggregate messages in the **local2** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
 change-log info;
 facility-override local2;
}
```

On the remote host named **central-logger** you can subsequently configure the system logging utility to write messages from the **local0** facility to one file (for example, **california-config**) and the messages from the **local2** facility to another file (for example, **new-york-config**).

#### Related Documentation

- [Junos OS System Log Alternate Facilities for Remote Logging on page 6224](#)

### Example: Configuring System Log Messages

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

This example describes how to configure system log messages on the QFabric system.

- [Requirements on page 6163](#)
- [Overview on page 6163](#)
- [Configuration on page 6163](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.1
- QFabric system
- External servers that can be configured as system log message hosts

### Overview

Component devices that generate system log message events may include Node devices, Interconnect devices, Director devices, and the control plane switches. The following configuration example includes these components in the QFabric system:

- Director software running on the Director group
- Control plane switches
- Interconnect device
- Multiple Node devices

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog host 10.1.1.12 any error
set system syslog file qflogs structured-data brief
set system syslog file qflogs archive size 1g
```

#### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure system messages from the QFabric Director device:

1. Specify a host, any facility, and the **error** severity level.
 

```
[edit system syslog]
user@switch# set host 10.1.1.12 any error
```



**NOTE:** You can configure more than one system log message server (host). The QFabric system sends the messages to each server configured.

2. (Optional) Specify a filename to capture log messages.



**NOTE:** On the QFabric system, a syslog file named `messages` is configured implicitly with facility and severity levels of `any any` and a file size of 100 MBs.

[edit system syslog]

user@switch# **set file qflogs structured-data brief**

3. (Optional) Configure the maximum size of your system log message archive file. This example specifies an archive size of 1 GB.

[edit system syslog]

user@switch# **set file qflogs archive size 1g**

**Results** From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

[edit]

user@switch# **show system**

```
syslog {
 file qflogs {
 archive {
 size 1g;
 }
 structured-data {
 brief;
 }
 }
 host 10.1.1.12 {
 any error;
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Overview of QFabric System Logs](#)
  - [syslog \(QFabric System\) on page 1403](#)



## Example: Monitoring Network Traffic Using sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a QFX Series device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

This example describes how to configure and use sFlow monitoring on a QFX3500 switch in standalone mode.

- [Requirements on page 6165](#)
- [Overview on page 6165](#)
- [Configuration on page 6166](#)
- [Verification on page 6167](#)

---

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.3 or later
- One QFX3500 switch


---

### Overview

An sFlow monitoring system consists of an sFlow agent embedded in the device and a centralized collector on the network. The two main activities of the sFlow agent are random sampling and statistics gathering. The sFlow agent combines interface counters and flow samples and sends them to the IP address and UDP destination port of the sFlow collector in UDP datagrams.

[Figure 225 on page 6166](#) depicts the basic elements of an sFlow system.

Figure 225: sFlow Technology Monitoring System



### Configuration

#### CLI Quick Configuration

To quickly configure sFlow technology, copy the following commands and paste them into the terminal window of the switch:

```
[edit protocols sflow]
set collector 10.204.32.46 udp-port 5600
set interfaces xe-0/0/1.0
set polling-interval 20
set sample-rate 1000
```

#### Step-by-Step Procedure

To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@switch# set collector 10.204.32.46 udp-port 5600
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces xe-0/0/1.0
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example, ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@switch# set polling-interval 20
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets must be sampled at the global level. The following example sets a sample rate of 1 in 1000 packets:

```
[edit protocols sflow]
user@switch# set sample-rate 1000
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show protocols
sflow {
 collector 10.204.32.46 {
 udp-port 5600;
 }
 interfaces xe-0/0/1.0 {
 polling-interval 20;
 sample-rate 1000;
 }
}
```

### Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That sFlow Technology Has Been Configured Properly on page 6167](#)
- [Verifying That sFlow Technology Is Enabled on an Interface on page 6168](#)
- [Verifying the sFlow Collector Configuration on page 6168](#)

#### *Verifying That sFlow Technology Has Been Configured Properly*

**Purpose** Verify that sFlow technology has been configured properly.

**Action** Enter the **show sflow** operational mode command:

```
user@switch> show sflow
```

```
sFlow : Enabled
Sample limit : 300 packets/second
Polling interval : 20 second
Sample rate : 1:1000
Agent ID : 10.1.1.2
```



**NOTE:** The sample limit cannot be configured and is set to 300 packets per second.

**Meaning** The output shows that sFlow technology is enabled and specifies the values for the sampling limit, polling interval, and sampling rate.

#### *Verifying That sFlow Technology Is Enabled on an Interface*

**Purpose** Verify that sFlow technology is enabled on interfaces and display the sampling parameters.

**Action** Enter the **show sflow interface** operational mode command:

```
user@switch> show sflow interface
Interface Status Sample Polling
 rate interval
xe-0/0/1.0 Enabled 1000 20
```

**Meaning** The output indicates that sFlow technology is enabled on the **Node1:xe-0/0/1.0** interface on the Node device with a sampling rate of 1000 and a polling interval of 20 seconds.

#### *Verifying the sFlow Collector Configuration*

**Purpose** Verify the sFlow collector configuration.

**Action** Enter the **show sflow collector** operational mode command:

```
user@switch> show sflow collector
Collector Udp-port No. of samples
address
10.204.32.46 5600 7516
```

**Meaning** The output displays the IP address of the collector, the UDP port, and the number of samples collected.

**Related Documentation**

- [Configuring sFlow Technology on page 6189](#)
- [Overview of sFlow Technology on page 6105](#)

## Example: Configuring SNMP

By default, SNMP is disabled on devices running Junos OS. This example describes the steps for configuring SNMP on the QFabric system.

- [Requirements on page 6169](#)
- [Overview on page 6169](#)
- [Configuration on page 6169](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- Network management system (NMS) (running the SNMP manager)
- QFabric system (running the SNMP agent) with multiple Node devices

### Overview

Because SNMP is disabled by default on devices running Junos OS, you must enable SNMP on your device by including configuration statements at the **[edit snmp]** hierarchy level. At a minimum, you must configure the **community public** statement. The community defined as public grants read-only access to MIB data to any client.

If no **clients** statement is configured, all clients are allowed. We recommend that you always include the **restrict** option to limit SNMP client access to the switch.

The network topology in this example includes an NMS, a QFabric system with four Node devices, and external SNMP servers that are configured for receiving traps.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set snmp name "snmp qfabric" description "qfabric0 switch"
set snmp location "Lab 4 Row 11" contact "qfabric-admin@qfabric0"
set snmp community public authorization read-only
set snmp client-list list0 192.168.0.0/24
set snmp community public client-list-name list0
set snmp community public clients 192.170.0.0/24 restrict
set snmp trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure SNMP on the QFabric system:



**NOTE:** If the name, description, location, contact, or community name contains spaces, enclose the text in quotation marks (" ").

1. Configure the SNMP system name:

```
[edit snmp]
user@switch# set name "snmp qfabric"
```

2. Specify a description.

```
[edit snmp]
user@switch# set description "qfabric0 system"
```

This string is placed into the MIB II sysDescription object.

3. Specify the physical location of the QFabric system.

```
[edit snmp]
user@switch# set location "Lab 4 Row 11"
```

This string is placed into the MIB II sysLocation object.

4. Specify an administrative contact for the SNMP system.

```
[edit snmp]
user@switch# set contact "qfabric-admin@qfabric0"
```

This name is placed into the MIB II sysContact object.

5. Specify a unique SNMP community name and the read-only authorization level.



**NOTE:** The read-write option is not supported on the QFabric system.

```
[edit snmp]
user@switch# set community public authorization read-only
```

6. Create a client list with a set of IP addresses that can use the SNMP community.

```
[edit snmp]
user@switch# set client-list list0 192.168.0.0/24
user@switch# set community public client-list-name list0
```

7. Specify IP addresses of clients that are restricted from using the community.

```
[edit snmp]
user@switch# set community public clients 192.170.0.0/24 restrict
```

8. Configure a trap group, destination port, and a target to receive the SNMP traps in the trap group.

```
[edit snmp]
user@switch# set trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```



**NOTE:** You do not need to include the `destination-port` statement if you use the default port 162.

The trap group `qf-traps` is configured to send traps to 192.168.0.100.

**Results** From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show
snmp {
 name "snmp qfabric";
 description "qfabric0 system";
 location "Lab 4 Row 11";
 contact "qfabric-admin@qfabric0";
 client-list list0 {
 192.168.0.0/24;
 }
 community public {
 authorization read-only;
 clients {
 197.170.0.0/24 restrict;
 }
 }
 trap-group qf-traps {
 destination-port 155;
 targets {
 192.168.0.100;
 }
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)
  - [snmp on page 1810](#)

## Example: Configuring Internal Fabric OAM Monitoring

This example shows how to configure the internal fabric Operation, Administration, Maintenance (OAM) monitoring feature on the QFabric system, including the fabric maintenance association (FMA) and flow specifications for unicast Ethernet traffic.

Internal fabric monitoring enables you to validate the flow path of protocol data units (PDUs) across a given VLAN on the QFabric system using the ping operation.

Internal fabric monitoring is useful for fault detection on the QFabric system. For example, if a PDU reaches a destination that is not part of the VLAN configuration, the ping operation displays the exception on the console at runtime.

- [Requirements on page 6172](#)
- [Overview on page 6172](#)
- [Configuration on page 6174](#)
- [Verification on page 6176](#)

---

## Requirements

This example uses the following hardware and software components:

- One QFabric system with the following components:
  - A Director group (two Director devices)
  - Two QFabric system Interconnect devices
  - Node devices, including:
    - One Node device (Node-81) configured in a default network Node group (NNW-NG-0)
    - One Node device (Node-80) in an autogenerated server Node group (SNG-80)
    - Two Node devices (Node-83 and Node-84) configured in a redundant server Node group (RSNG-8384)
- Junos OS Release 12.2 or later

Before you configure the FMA, first configure a VLAN in which the internal fabric OAM monitoring occurs. You can usually configure a VLAN by specifying a VLAN ID or VLAN name. However, in the case of the VLAN used for internal fabric OAM monitoring, you must specify a VLAN name (not a VLAN ID). The FMA configuration requires a VLAN name and not a VLAN ID.

If you have already configured a VLAN ID for the VLAN, you must delete the VLAN ID from the associated interfaces and then add a VLAN name.

---

## Overview

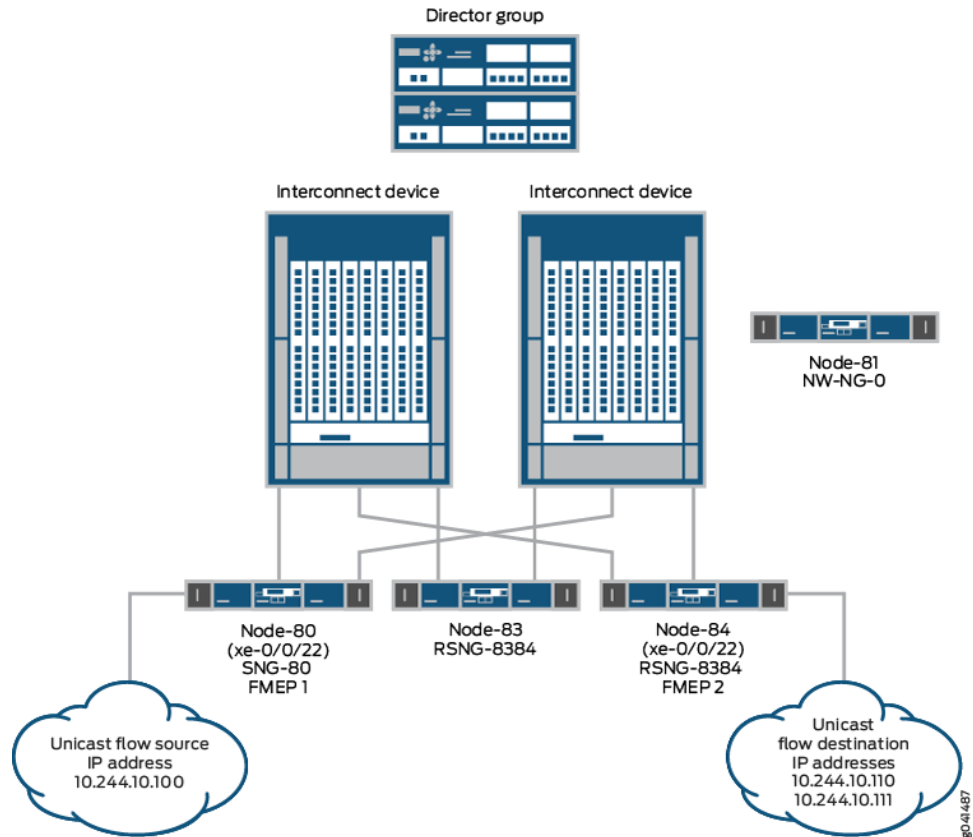
On a QFabric system, enabling internal fabric OAM monitoring requires that you configure an FMA and a flow specification.

- The *FMA* associates a particular VLAN with a set of source and destination interfaces called fabric maintenance endpoints (FMEPs) that are used for tracing the data path between the endpoints. Configuration of the VLAN name is mandatory when you configure an FMA.
- A *flow specification* is used to determine the contents of a packet that is used for a particular operation. For example, an Ethernet type flow specification describes the contents of a pure Layer 2 packet, whereas an Ethernet IPv4 type flow specification describes an Ethernet IPv4 packet.



### Topology

This example uses the following topology:



- The QFabric system Node devices are connected to the Interconnect devices.
- Node device Node-80 is configured in the server Node group SNG-80.
- Node devices Node-83 and Node-84 are configured in a redundant server Node group RSNG-8384.
- The device with the unicast flow source IP address is connected to Node-80.
- The device with the unicast flow destination IP addresses is connected to Node-84.

Table 602 on page 6173 maps the association of the fabric OAM elements configured in this example.

**Table 602: Fabric OAM Configuration Elements**

| Interface         | FMA  | VLAN Name | FMEP ID | Flow Specifications                                                            |
|-------------------|------|-----------|---------|--------------------------------------------------------------------------------|
| Node-80:xe-0/0/22 | fma1 | v10       | 1       | <ul style="list-style-type: none"> <li>flowspec1</li> <li>flowspec2</li> </ul> |
| Node-84:xe-0/0/22 | fma1 | v10       | 2       | <ul style="list-style-type: none"> <li>flowspec1</li> <li>flowspec2</li> </ul> |

## Configuration

---

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set vlans v10 vlan-id 10
set interfaces Node-80:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
set interfaces Node-84:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
set protocols oam fabric fabric-maintenance-associations fma1 vlan-name v10
set protocols oam fabric fabric-maintenance-associations fma1
 fabric-maintenance-end-points 1 fmep-interface Node-80:xe-0/0/22.0
set protocols oam fabric fabric-maintenance-associations fma1
 fabric-maintenance-end-points 2 fmep-interface Node-84:xe-0/0/22.0
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 source-ip
 10.244.10.100
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-ip
 10.244.10.110
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-mac
 00:00:80:8d:2d:15
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 source-l4-port 60
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-l4-port
 80
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 source-ip
 10.244.10.100
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-ip
 10.244.10.111
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-mac
 00:00:80:8d:2d:17
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 source-l4-port 60
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-l4-port
 80
```

### Step-by-Step Procedure

To configure internal fabric OAM monitoring:

1. Configure a name and ID for the VLAN.  

```
[edit]
set vlans v10 vlan-id 10
```
2. Assign a set of interfaces to the VLAN.  

```
[edit]
set interfaces Node-80:xe-0/0/22 unit 0 family ethernet-switching vlan members
v10
set interfaces Node-84:xe-0/0/22 unit 0 family ethernet-switching vlan members
v10
```
3. Define the FMA and associate it with a VLAN name.  

```
[edit protocols oam fabric]
set fabric-maintenance-associations fma1 vlan-name v10
```
4. Configure FMEPs for the FMA.  

```
[edit protocols oam fabric]
```

```

set fabric-maintenance-associations fma1 fabric-maintenance-end-points 1
fmem-interface Node-80:xe-0/0/22.0
set fabric-maintenance-associations fma1 fabric-maintenance-end-points 2
fmem-interface Node-84:xe-0/0/22.0

```

5. Configure two IPv4 flow specifications for unicast traffic from FMEP 1 to FMEP 2.

```

[edit protocols oam fabric]
set flow-specs flowspec1 unicast-ethernet-ipv4 source-ip 10.244.10.100
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-ip 10.244.10.110
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:15
set flow-specs flowspec1 unicast-ethernet-ipv4 source-l4-port 60
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-l4-port 80
set flow-specs flowspec2 unicast-ethernet-ipv4 source-ip 10.244.10.100
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-ip 10.244.10.111
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:17
set flow-specs flowspec2 unicast-ethernet-ipv4 source-l4-port 60
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-l4-port 80

```

**Results** Display the results of the configuration.

```

root@qfabric> show protocols oam fabric
fabric-maintenance-associations {
 fma1 {
 vlan-name v10;
 fabric-maintenance-end-points {
 1 {
 fmem-interface Node-80:xe-0/0/22.0;
 }
 2 {
 fmem-interface Node-84:xe-0/0/22.0;
 }
 }
 }
}
flow-specs {
 flowspec1 {
 unicast-ethernet-ipv4 {
 source-ip 10.244.10.100;
 destination-ip 10.244.10.110;
 destination-mac 0:0:80:8d:2d:15;
 source-l4-port 60;
 destination-l4-port 80;
 }
 }
 flowspec2 {
 unicast-ethernet-ipv4 {
 source-ip 10.244.10.100;
 destination-ip 10.244.10.111;
 destination-mac 0:0:80:8d:2d:17;
 source-l4-port 60;
 destination-l4-port 80;
 }
 }
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying Configuration of Fabric OAM Interfaces on page 6176](#)
- [Verifying Configuration of the Fabric OAM Flow Specification on page 6176](#)
- [Verifying Operation of the Unicast ping Command on page 6177](#)

### Verifying Configuration of Fabric OAM Interfaces

**Purpose** Verify that the interfaces are associated with the FMA.

**Action** Display the fabric OAM interfaces.

```
root@qfabric> show oam fabric interfaces
```

| Interface-name      | Fabric-Maintenance Association | VLAN | Interface state | MEP Identifier | MEP Name               |
|---------------------|--------------------------------|------|-----------------|----------------|------------------------|
| NW-NG-0:NULL        | fma-default                    | *    | up              | 32790          | fmeop-default-BBAK8793 |
| NW-NG-0:NULL        | fma-default                    | *    | up              | 32791          | fmeop-default-BBAK8798 |
| RSNG-8384:NULL      | fma-default                    | *    | up              | 32794          | fmeop-default-BBAK8748 |
| RSNG-8384:NULL      | fma-default                    | *    | up              | 32795          | fmeop-default-BBAK0465 |
| Node-84:xe-0/0/22.0 | fma1                           | V10  | up              | 2              |                        |
| SNG-60:NULL         | fma-default                    | *    | up              | 32793          | fmeop-default-BBAK9603 |
| SNG-80:NULL         | fma-default                    | *    | up              | 32792          | fmeop-default-BBAK8728 |
| Node-80:xe-0/0/22.0 | fma1                           | V10  | up              | 1              |                        |

**Meaning** Interfaces Node-84:xe-0/0/22.0 and Node-80:xe-0/0/22.0 are associated with the FMA, FMEP, and VLAN as configured.

Interfaces that are not configured for fabric OAM monitoring (NW-NG-0:NULL, RSNG-8384:NULL, SNG-60:NULL, and SNG-80:NULL) display an asterisk (\*) in the VLAN column, and the default FMA and FMEP names in the Fabric-Maintenance Association and MEP Name columns.

### Verifying Configuration of the Fabric OAM Flow Specification

**Purpose** Verify the configuration of the flow specifications.

**Action** Display the fabric OAM flow specification configuration.

```
root@qfabric> show oam fabric flow-specification
Flow specification name : flowspec1 Type : Ethernet Unicast IPV4
Ethernet frame size : Unspecified
Source-IP : 10.244.10.100
Source-IP Mask : Unspecified
Destination-IP : 10.244.10.110
Destination-IP Mask : Unspecified
Destination-mac : 0:0:80:8d:2d:15
IP-protocol : Unspecified
Source-L4-port : 60
Destination-L4-port : 80
Flow specification name : flowspec2 Type : Ethernet Unicast IPV4
Ethernet frame size : Unspecified
Source-IP : 10.244.10.100
Source-IP Mask : Unspecified
```

```

Destination-IP : 10.244.10.111
Destination-IP Mask : Unspecified
Destination-mac : 0:0:80:8d:2d:17
IP-protocol : Unspecified
Source-L4-port : 60
Destination-L4-port : 80

```

**Meaning** The output shows that flow specifications were configured as the user intended.

### *Verifying Operation of the Unicast ping Command*

**Purpose** Verify that the fabric OAM unicast **ping** command works.

**Action** Issue the **ping fabric unicast-flow** command.

```

root@qfabric> ping fabric unicast-flow source-fmep-id 1 dest-fmep-id 2 flow-spec-name
flowspec1 fma-name fma1 count 3

```

```

Fabric flow ping between source Node-80 destination Node-84
received response from fmep-id 2...
received response from fmep-id 2...
received response from fmep-id 2...
sent 3 requests, received 3 responses

```

```

root@qfabric> ping fabric unicast-flow source-fmep-id 1 dest-fmep-id 2 fma-name
fma1 flow-spec-name flowspec2

```

```

Fabric flow ping between source Node-80 destination Node-84
received response from fmep-id 2...
sent 1 requests, received 1 responses

```

**Meaning** The fabric OAM unicast ping request from the source Node device (Node-80) to the destination Node device (Node-84) has been successful.

- Related Documentation**
- [Overview of Internal Fabric Monitoring on page 6099](#)
  - [fabric-maintenance-associations on page 6258](#)
  - [fabric-maintenance-end-points on page 6259](#)
  - [flow-specs on page 6260](#)
  - [show oam fabric flow-specification on page 6432](#)
  - [show oam fabric interfaces on page 6435](#)

## Configuration Tasks for Network Management

- [Configuring Console and Auxiliary Port Properties on page 6177](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 6178](#)
- [Configuring Telnet Service for Remote Access to a Switch on page 6180](#)

### Configuring Console and Auxiliary Port Properties

The console port and auxiliary port on a switch provide out-of-band remote access to the switch. You can configure the console and auxiliary ports so that an external data

terminal may be connected to the switch. The console port is enabled by default, and its speed is 9600 baud. The auxiliary port is disabled by default.

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

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

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

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

2. To specify the auxiliary port terminal type:

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

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

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

3. To check the configuration:

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

- Related Documentation
- [auxiliary on page 242](#)
  - [console \(Physical Port\) on page 251](#)
  - [ports on page 281](#)

## Configuring SSH Service for Remote Access to the Router or Switch

To configure the router or switch to accept SSH as an access service, include the **ssh** statement at the **[edit system services]** hierarchy level:

```
[edit system services]
ssh {
 ciphers [cipher-1 cipher-2 cipher-3 ...]
 client-alive-count-max number;
 client-alive-interval seconds;
 connection-limit limit;
 hostkey-algorithm <algorithm | no-algorithm>;
 key-exchange algorithm;
 macs algorithm;
 max-sessions-per-connection number;
 no-tcp-forwarding;
 protocol-version [v1 v2];
```

```

 rate-limit limit;
 root-login <allow | deny | deny-password>;
}

```

By default, the router or switch supports a limited number of simultaneous SSH sessions and connection attempts per minute. Use the following statements to change the defaults:

- **connection-limit *limit***—Maximum number of simultaneous connections per protocol (IPv4 and IPv6). The range is a value from 1 through 250. The default is 75. When you configure a connection limit, the limit is applicable to the number of SSH sessions per protocol (IPv4 and IPv6). For example, a connection limit of 10 allows 10 IPv6 SSH sessions and 10 IPv4 SSH sessions.
- **max-sessions-per-connection *number***—Include this statement to specify the maximum number of SSH sessions allowed per single SSH connection. This allows you to limit the number of cloned sessions tunneled within a single SSH connection. The default value is 10.
- **rate-limit *limit***—Maximum number of connection attempts accepted per minute (a value from 1 through 250). The default is 150. When you configure a rate limit, the limit is applicable to the number of connection attempts per protocol (IPv4 and IPv6). For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

For information about other configuration settings, see the following topics:

- [Configuring the Root Login Through SSH on page 6179](#)
- [Configuring the SSH Protocol Version on page 6179](#)
- [Configuring the Client Alive Mechanism on page 6180](#)

### Configuring the Root Login Through SSH

By default, users are allowed to log in to the router or switch as **root** through SSH. To control user access through SSH, include the **root-login** statement at the **[edit system services ssh]** hierarchy level:

```

[edit system services ssh]
root-login (allow | deny | deny-password);

```

**allow**—Allows users to log in to the router or switch as root through SSH. The default is **allow**.

**deny**—Disables users from logging in to the router or switch as root through SSH.

**deny-password**—Allows users to log in to the router or switch as root through SSH when the authentication method (for example, RSA) does not require a password.

### Configuring the SSH Protocol Version

By default, both version 1 and version 2 of the SSH protocol are enabled. To configure the router or switch to use only version 1 of the SSH protocol, include the **protocol-version** statement and specify **v1** at the **[edit system services ssh]** hierarchy level:

```

[edit system services ssh]

```

```
protocol-version [v1];
```

To configure the router or switch to use only version 2 of the SSH protocol, include the **protocol-version** statement and specify **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [v2];
```

To explicitly configure the router or switch to use version 1 and 2 of the SSH protocol, include the **protocol-version** statement and specify **v1** and **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [v1 v2];
```

For J Series Services Routers, the export license software supports SSH version 1 only.

---

### Configuring the Client Alive Mechanism

The client alive mechanism is valuable when the client or server depends on knowing when a connection has become inactive. It differs from the standard keepalive mechanism because the client alive messages are sent through the encrypted channel. The client alive mechanism is not enabled at default. To enable it, configure the **client-alive-count-max** and the **client-alive-interval**. This option applies to SSH protocol version 2 only.

In the following example, unresponsive SSH clients will be disconnected after approximately 100 seconds (20 x 5).

```
[edit system services ssh]
client-alive-count-max 5;
client-alive-interval 20;
```

### Configuring Telnet Service for Remote Access to a Switch

Telnet provides unencrypted access to network devices. Configuring Telnet service for a switch enables in-band remote access to the switch.

By default, the switch supports a limited number of simultaneous Telnet sessions and connection attempts per minute. Optionally, you can change the default Telnet settings by configuring the connection limit and rate limit at the **[edit system services telnet]** hierarchy level.

The connection limit is the maximum number of simultaneous connections per protocol (IPv4). The range is from 1 through 250. The default is 75.

The rate limit is the maximum number of connection attempts accepted per minute per protocol. The range is from 1 through 250. The default is 150.

To configure Telnet service:

1. To specify the connection limit:

```
[edit system services]
user@switch# set telnet connection-limit connection-limit
```



2. To specify the rate limit:

```
[edit system services]
user@switch# set telnet rate-limit rate-limit
```

3. Check that the Telnet connection limit and rate limit show the values you specified:

```
[edit system services]
user@switch# show
telnet {
 connection-limit 50;
 rate-limit 100;
}
```

## Configuration Tasks for Automation Scripts

- [Controlling the Execution of Commit Scripts on page 6181](#)

### Controlling the Execution of Commit Scripts

This document describes the tasks that affect the way commit scripts are executed. In the QFabric system, commit scripts are stored in the in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory that is shared among Director devices in a Director group.

To determine which commit scripts are currently enabled on the QFabric system, use the **show** command to display the files included at the `[edit system scripts commit]` hierarchy level. To ensure that the enabled files are on the device, list the contents of the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory using the **file list** operational mode command.

See the following tasks:

- [Enabling Commit Scripts to Execute on page 6181](#)
- [Removing Commit Scripts from the Configuration on page 6182](#)
- [Deactivating Commit Scripts on page 6183](#)
- [Activating Inactive Commit Scripts on page 6183](#)

### Enabling Commit Scripts to Execute

The commit operation requires that all scripts be included in configuration at the `[edit system scripts commit file]` hierarchy level for all QFabric Director devices.

If you need to temporarily remove a script from a commit operation but do not want to remove it from the configuration permanently, you may configure the **optional** statement at the `[edit system scripts commit file filename]` hierarchy level to enable the commit operation to succeed even if a script is missing from the commit script directory.



**CAUTION:** When you include the **optional** statement at the `[edit system scripts commit file filename]` hierarchy level, no error message is generated during

the commit operation if the file does not exist. As a result, you might not be aware that a script has not been executed as expected.

.....

The filename of a commit script written in SLAX must include the `.slax` extension for the script to be executed.

To enable a commit script to execute during a commit operation:

1. Ensure that the commit script is located in the correct directory:  
`/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory on the Director device.

2. Configure the commit script.

```
[edit system scripts commit]
user@switch# set file filename <optional>
```

3. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

## Removing Commit Scripts from the Configuration

---

You can prevent commit scripts from executing during a commit operation by removing the scripts from the commit directory in the configuration.



**NOTE:** You can also deactivate a script using the `deactivate` statement instead of removing it from the configuration. Deactivated scripts may be reactivated later.

.....

To prevent a commit script from executing during a commit operation:

1. Delete the commit script file from the commit directory in the configuration.

```
[edit system scripts commit]
user@switch# delete file filename
```

2. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

3. Remove the commit script from the `/pbdata/mgd_shared/` directory on the Director device.
- .....



**BEST PRACTICE:** Although removing the commit script is not necessary, we recommend deleting unused files from the system.

.....

### Deactivating Commit Scripts

---

Deactivating a commit script results in its being marked as inactive in the configuration. The script is not executed during the commit operation, but you can reactivate the script by using the **activate** statement.

To deactivate the commit script:

1. Deactivate the script.

```
[edit]
user@switch deactivate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

3. Verify that the commit script is deactivated.

```
[edit]
user@switch# show system scripts commit
inactive: file mycommit.slax
```

### Activating Inactive Commit Scripts

---

Deactivating a commit script results in its being marked as inactive in the configuration and is therefore not executed during the commit operation.

To activate an inactive commit script:

1. Activate the script.

```
[edit]
user@switch# activate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

## Configuration Tasks for Fabric OAM

---

- [Configuring a Fabric Maintenance Association on page 6184](#)
- [Configuring Flow Specifications on page 6185](#)

## Configuring a Fabric Maintenance Association

On a QFabric system, enabling internal fabric monitoring using the ping and traceroute operations requires that a fabric maintenance association (FMA) be configured in addition to a flow specification. The FMA associates a particular VLAN with a set of source and destination interfaces called fabric maintenance endpoints (FMEPs) that are used for tracing the flow path between the endpoints. Configuration of the VLAN name is mandatory.

Before configuring the FMA, configure a VLAN name (and not a VLAN ID) for the VLAN in which the internal fabric monitoring occurs. If you already configured a VLAN ID instead of a VLAN name, you must first delete the VLAN ID and then specify a VLAN name before you configure the FMA.



**NOTE:** Because each FMA and VLAN have a one-to-one correspondence, an FMA in each QFabric system must be given a unique name.

To configure an FMA:

1. Configure an FMA name that is unique within the QFabric system:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name
```

2. (Optional) Configure the FMA description:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name description string
```

3. Configure the VLAN name:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name vlan-name vlan-name
```

4. Configure an FMEP, including the identifier, name, description, and associated interface:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name
 fabric-maintenance-end-points fmeop-id fmeop-name name description string
 fmeop-interface interface-name
```

5. Repeat Step 4 to configure additional FMEPs.

### Related Documentation

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring Flow Specifications on page 6185](#)
- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

## Configuring Flow Specifications

On a QFabric system, enabling internal fabric monitoring using the ping and traceroute operations requires that a flow specification be configured in addition to a fabric maintenance association (FMA). A flow specification is used to determine the contents of the packet that is used for a particular operation. For example, an Ethernet type flow specification describes the contents of a pure Layer 2 packet, whereas an Ethernet IPv4 type flow specification describes an Ethernet IPv4 packet.



**NOTE:** In the flow specification configuration, some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets:

- Unicast Ethernet—Source MAC address, destination MAC address, Ethertype
- Unicast IPv4—Source IP address, destination IP address, source L4 port, destination L4 port, destination MAC address

The configuration for each flow specification type includes parameters specific to that particular type. Use one of the following procedures to configure the flow specification for your flow type:

- [Configuring a Unicast Ethernet Flow Specification on page 6185](#)
- [Configuring a Unicast Ethernet IPv4 Flow Specification on page 6186](#)
- [Configuring a Multicast IPv4 Flow Specification on page 6187](#)
- [Configuring a Multicast VLAN Flood Flow Specification on page 6188](#)

### Configuring a Unicast Ethernet Flow Specification

This procedure describes how to configure a flow specification for a unicast Ethernet flow type. This configuration is applicable for the unicast ping and traceroute operations.



**NOTE:** Some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets: source MAC address, destination MAC address, and Ethertype.

To configure a unicast Ethernet flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. (Optional) Set the Ethernet frame size:

```
[edit protocols oam fabric]
user@host# set ethernet-frame-size ethernet-frame-size
```

3. Set the flow specification type as unicast Ethernet and specify the EtherType:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set unicast-ethernet ethertype ethertype
```

4. (Optional) Set the MAC address and the MAC address range of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet]
user@host# set source-mac source-mac-address source-mac-mask source-mac-mask
```

---



**NOTE:** Configuration of the `source-mac` and `source-mac-mask` parameters is optional, but the `source-mac-mask` parameter can be specified only if the `source-mac` parameter is also configured.

---

5. (Optional) Set the MAC address or address range mask for the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet]
user@host# set destination-mac destination-mac-address destination-mac-mask destination-mac-mask
```

---



**NOTE:** Configuration of the `destination-mac` and `destination-mac-mask` parameters is optional, but the `destination-mac-mask` parameter can be specified only if the `destination-mac` parameter is also configured.

---

## Configuring a Unicast Ethernet IPv4 Flow Specification

---

This procedure describes how to configure a flow specification for a unicast Ethernet IPv4 flow type. This configuration is applicable for the unicast ping and traceroute operations.

---



**NOTE:** All parameters except the flow specification name are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets: source IP address, destination IP address, source L4 port, destination L4 port, destination MAC address.

---

To configure a unicast Ethernet IPv4 flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. (Optional) Set the Ethernet frame size:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set ethernet-frame-size size
```

3. Set the flow specification type as unicast Ethernet IPv4:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set unicast-ethernet-ipv4
```

4. (Optional) Set the IPv4 address or the address range of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set source-ip source-ip-address
```

or

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set source-ip-mask source-ip-mask
```

5. (Optional) Set the IPv4 address or the address range of the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set destination-ip destination-ip-address
```

or

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set destination-ip-mask destination-ip-mask
```

6. (Optional) Set the IPv4 protocol type:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# ip-proto protocol
```

7. (Optional) Set the L4 port (TCP or UDP port number) of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# source-l4-port source-l4-port-number
```

8. (Optional) Set the L4 port of the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# destination-l4-port destination-l4-port-number
```

### Configuring a Multicast IPv4 Flow Specification

This procedure describes how to configure a flow specification for a multicast IPv4 flow type.



**NOTE:** Some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

To configure a multicast IPv4 flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. Set the flow specification type as multicast IPv4:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set multicast-ipv4
```

3. (Optional) Set the IPv4 address of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name multicast-ipv4]
user@host# set source-ip ipv4-address
```

4. Set the IPv4 multicast group address:

```
[edit protocols oam fabric flow-specs flow-specification-name multicast-ipv4]
user@host# set dest-ip-multicast-group ipv4-mcast-address
```

---

### Configuring a Multicast VLAN Flood Flow Specification

This procedure describes how to configure a flow specification for a multicast VLAN flood flow type.

To configure a multicast VLAN flood flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. Set the flow specification type as multicast VLAN flood:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set multicast-vlan-flood
```

#### Related Documentation

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring a Fabric Maintenance Association on page 6184](#)
- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

---

### Configuration Tasks for sFlow Technology

- [Configuring sFlow Technology on page 6189](#)



## Configuring sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a QFX Series device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

On the QFabric system, the sFlow monitoring global configuration that is defined on the Director device is distributed to Node groups that have sFlow sampling configured on the interfaces.

To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@host# set collector ip-address udp-port port-number
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@host# set interfaces interface-name
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@host# set polling-interval seconds
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets are sampled at the global level. For example, configuring a *number* of 1000 sets a sample rate of 1 in 1000 packets.

```
[edit protocols sflow]
user@host# set sample-rate number
```

5. (Optional) You can also configure the polling interval and sample rate at the interface level:

```
[edit protocols sflow]
```

```
user@host# set interfaces interface-name polling-interval seconds sample-rate number
```



**NOTE:** The interface-level configuration overrides the global configuration for the specified interface.

**Related  
Documentation**

- [Example: Monitoring Network Traffic Using sFlow Technology on page 6165](#)
- [Overview of sFlow Technology on page 6105](#)

---

## Configuration Tasks for SNMP

- [Configuring SNMP on page 6190](#)
- [Configuring the SNMP Community String on page 6194](#)
- [Configuring SNMP Trap Groups on page 6195](#)
- [Adding a Group of Clients to an SNMP Community on page 6196](#)
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 6197](#)
- [Configuring MIB Views on page 6197](#)
- [Configuring RMON Alarms and Events on page 6198](#)
- [Configuring Health Monitoring on page 6200](#)
- [Creating SNMPv3 Users on page 6201](#)
- [Configuring Access Privileges for a Group on page 6202](#)
- [Assigning a Security Name to a Group on page 6204](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6204](#)
- [Configuring SNMP Informs on page 6206](#)

### Configuring SNMP

SNMP is implemented in the Junos OS Software running on the QFX Series products. By default, SNMP is not enabled. To enable SNMP, you must include the SNMP configuration statements at the **[edit]** hierarchy level.

To configure the minimum requirements for SNMP, include the following statements at the **[edit]** hierarchy level of the configuration:

```
[edit]
snmp {
 community public;
}
```

To configure complete SNMP features, include the following statements at the **[edit]** hierarchy level of the configuration:

```
snmp {
 client-list client-list-name {
 ip-addresses;
 }
}
```

```

community community-name {
 authorization authorization;
 client-list-name client-list-name;
 clients {
 address restrict;
 }
 logical-system logical-system-name {
 routing-instance routing-instance-name {
 clients {
 addresses;
 }
 }
 }
 routing-instance routing-instance-name {
 clients {
 addresses;
 }
 }
 view view-name;
}
contact contact;
description description;
filter-duplicates;
filter-interfaces;
health-monitor {
 falling-threshold integer;
 interval seconds;
 rising-threshold integer;
}
interface [interface-names];
location location;
name name;
nonvolatile {
 commit-delay seconds;
}
rmon {
 alarm index {
 description description;
 falling-event-index index;
 falling-threshold integer;
 falling-threshold-interval seconds;
 interval seconds;
 request-type;
 rising-event-index index;
 rising-threshold integer;
 sample-type (absolute-value | delta-value);
 startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
 syslog-subtag syslog-subtag;
 variable oid-variable;
 }
 event index {
 community community-name;
 description description;
 type type;
 }
 history history-index {

```

```
 bucket-size number;
 interface interface-name;
 interval seconds;
 owner owner-name;
 }
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable> <match
 regular-expression>;
 flag flag;
}
trap-group group-name {
 categories {
 category;
 }
 destination-port port-number;
 routing-instance routing-instance-name;
 targets {
 address;
 }
 version (all | v1 | v2);
}
trap-options {
 agent-address outgoing-interface;
 source-address address;
}
v3 {
 notify name {
 tag tag-name;
 type trap;
 }
 notify-filter profile-name {
 oid object-identifier (include | exclude);
 }
 snmp-community community-index {
 community-name community-name;
 security-name security-name;
 tag tag-name;
 }
 target-address target-address-name {
 address address;
 address-mask address-mask;
 logical-system logical-system;
 port port-number;
 retry-count number;
 routing-instance routing-instance-name;
 tag-list tag-list;
 target-parameters target-parameters-name;
 timeout seconds;
 }
 target-parameters target-parameters-name {
 notify-filter profile-name;
 parameters {
 message-processing-model (v1 | v2c | V3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 }
 }
}
```

```

 security-name security-name;
 }
}
usm {
 local-engine {
 user username {
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none;
 }
 }
 remote-engine engine-id {
 user username {
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
 }
}
vacm {
 access {
 group group-name {
 (default-context-prefix | context-prefix context-prefix) {
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 }
 }
 }
 }
 }
}

```

```
 read-view view-name;
 write-view view-name;
 }
}
}
}
}
security-to-group {
 security-model (usm | v1 | v2c) {
 security-name security-name {
 group group-name;
 }
 }
}
}
}
view view-name {
 oid object-identifier (include | exclude);
}
}
```

**Related  
Documentation**

- [Understanding the Implementation of SNMP on page 6107](#)
- [snmp on page 1810](#)

## Configuring the SNMP Community String

The SNMP community string defines the relationship between an SNMP server system and the client systems. This string acts like a password to control the clients' access to the server. To configure a community string in a Junos OS configuration, include the **community** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
community name {
 authorization authorization;
 clients {
 default restrict;
 address restrict;
 }
 view view-name;
}
```

If the community name contains spaces, enclose it in quotation marks (" ").

The default authorization level for a community is **read-only**. To allow **Set** requests within a community, you need to define that community as **authorization read-write**. For **Set** requests, you also need to include the specific MIB objects that are accessible with read-write privileges using the **view** statement. The default view includes all supported MIB objects that are accessible with read-only privileges; no MIB objects are accessible with read-write privileges. For more information about the **view** statement, see [“Configuring MIB Views” on page 6197](#).

The **clients** statement lists the IP addresses of the clients (community members) that are allowed to use this community. If no **clients** statement is present, all clients are

allowed. For **address**, you must specify an IPv4 address, not a hostname. Include the **default restrict** option to deny access to all SNMP clients for which access is not explicitly granted. We recommend that you always include the **default restrict** option to limit SNMP client access to the local switch.



**NOTE:** Community names must be unique within each SNMP system.

**Related Documentation**

- [Configuring SNMP on page 1703](#)

## Configuring SNMP Trap Groups

Before any SNMP traps can be sent, you must configure a trap group, the categories of traps the group can receive, and the targets (systems) that will receive the traps. To create and name an SNMP trap group, include the **trap-group** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
trap-group group-name {
 categories {
 category;
 }
 destination-port port-number;
 targets {
 address;
 }
 version (all | v1 | v2);
}
```

The trap group name can be any string and is embedded in the community name field of the trap. To configure your own trap group port, include the **destination-port** statement. The default destination port is port 162.

For each trap group that you define, you must include the **target** statement to define at least one system as the recipient of the SNMP traps in the trap group. Specify the IPv4 address of each recipient and not its hostname.

Specify the types of traps the trap group can receive in the **categories** statement.

A trap group can receive the following categories of traps:

- **authentication**—Authentication failures
- **chassis**—Chassis or environment notifications
- **configuration**—Configuration notifications
- **link**—Link-related notifications such as up-down transitions
- **remote-operations**—Remote operation notifications
- **startup**—System warm and cold starts

The **version** statement allows you to specify the SNMP version of the traps sent to targets of the trap group. If you specify **v1** only, SNMPv1 traps are sent. If you specify **v2** only, SNMPv2 traps are sent. If you specify **all**, both an SNMPv1 and an SNMPv2 trap are sent for every trap condition. For more information about the **version** statement, see [version](#).

**Related  
Documentation**

- *Standard SNMP Version 1 Traps*
- *Standard SNMP Version 2 Traps*
- *Juniper Networks Enterprise-Specific SNMP Version 1 Traps*
- *Juniper Networks Enterprise-Specific SNMP Version 2 Traps*

## Adding a Group of Clients to an SNMP Community

Junos OS enables you to add one or more groups of clients to an SNMP community. You can include the **client-list-name** *name* statement at the **[edit snmp community community-name]** hierarchy level to add all the members of the client list or prefix list to an SNMP community.

To define a list of clients, include the **client-list** statement followed by the IP addresses of the clients at the **[edit snmp]** hierarchy level:

```
[edit snmp]
 client-list client-list-name {
 ip-addresses;
 }
```

You can configure a prefix list at the **[edit policy options]** hierarchy level. Support for prefix lists in the SNMP community configuration enables you to use a single list to configure the SNMP and routing policies. For more information about the **prefix-list** statement, see the *Routing Policy Configuration Guide*.

To add a client list or prefix list to an SNMP community, include the **client-list-name** statement at the **[edit snmp community community-name]** hierarchy level:

```
[edit snmp community community-name]
 client-list-name client-list-name;
```



**NOTE:** The client list and prefix list must not have the same name.

---

The following example shows how to define a client list:

```
[edit]
snmp {
 client-list clentlist1 {
 10.1.1.1/32;
 10.2.2.2/32;
 }
}
```

The following example shows how to add a client list to an SNMP community:

```
[edit]
```



```
snmp {
 community community1 {
 authorization read-only;
 client-list-name clientlist1;
 }
}
```

The following example shows how to add a prefix list to an SNMP community:

```
[edit]
policy-options {
 prefix-list prefixlist {
 10.3.3.3/32;
 10.5.5.5/32;
 }
}
snmp {
 community community2 {
 client-list-name prefixlist;
 }
}
```

- Related Documentation**
- [client-list on page 1763](#)
  - [client-list-name on page 1764](#)

## Configuring the Interfaces on Which SNMP Requests Can Be Accepted

By default, all router or switch interfaces have SNMP access privileges. To limit the access through certain interfaces only, include the **interface** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
interface [interface-names];
```

Specify the names of any logical or physical interfaces that should have SNMP access privileges. Any SNMP requests entering the router or switch from interfaces not listed are discarded.

- Related Documentation**
- *Configuring SNMP on a Device Running Junos OS*
  - *Configuration Statements at the [edit snmp] Hierarchy Level*
  - *Example: Configuring Secured Access List Checking*
  - [Configuring SNMP on page 1703](#)

## Configuring MIB Views

By default, an SNMP community grants read access and denies write access to all supported MIB objects (even communities configured as **authorization read-write**). To restrict or grant read or write access to a set of MIB objects, you must configure a MIB view and associate the view with a community.

To configure MIB views, include the **view** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
view view-name {
 oid object-identifier (include | exclude);
}
```

The **view** statement defines a MIB view and identifies a group of MIB objects. Each MIB object of a view has a common object identifier (OID) prefix. Each object identifier represents a subtree of the MIB object hierarchy. The subtree can be represented either by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as interfaces). A configuration statement uses a view to specify a group of MIB objects on which to define access. You can also use a wildcard character asterisk (\*) to include OIDs that match a particular pattern in the SNMP view. To enable a view, you must associate the view with a community.



**NOTE:** To remove an OID completely, use the **delete view all oid oid-number** command but omit the include parameter.

To associate MIB views with a community, include the **view** statement at the **[edit snmp community community-name]** hierarchy level:

```
[edit snmp community community-name]
view view-name;
```

For more information about the Ping MIB, see RFC 2925 and the *PING MIB* topic in the *SNMP MIBs and Traps Reference*.

#### Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
- [Example: Ping Proxy MIB](#)
- [view \(Configuring a MIB View\)](#)
- [view \(Associating MIB View with a Community\) on page 6364](#)
- [oid](#)

## Configuring RMON Alarms and Events

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

To configure RMON alarms and events using the CLI, perform these tasks:

1. [Configuring SNMP on page 6199](#)
2. [Configuring an Event on page 6199](#)
3. [Configuring an Alarm on page 6200](#)

## Configuring SNMP

To configure SNMP:

1. Grant read-only access to all SNMP clients:

```
[edit snmp]
user@switch# set community community-name authorization authorization
For example:
```

```
[edit snmp]
user@switch# set community public authorization read-only
```

2. Grant read-write access to the RMON and jnx-rmon MIBs:

```
[edit snmp]
user@switch# set view view-name oid object-identifier include
user@switch# set view view-name oid object-identifier include
user@switch# set community community-name authorization authorization view view-name
For example:
```

```
[edit snmp]
user@switch# set view rmon-mib-view oid .1.3.6.1.2.1.16 include
user@switch# set view rmon-mib-view oid .1.3.6.1.4.1.2636.13 include
user@switch# set community private authorization read-write view rmon-mib-view
OIDs 1.3.6.1.2.1.16 and 1.3.6.1.4.1.2636.13 correspond to the RMON and jnxRmon MIBs.
```

3. Configure an SNMP trap group:

```
[edit snmp]
user@switch# set trap-group group-name categories category
user@switch# set trap-group group-name targets address
For example:
```

```
[edit snmp]
user@switch# set trap-group rmon-trap-group categories rmon-alarm
user@switch# set trap-group rmon-trap-group targets 192.168.5.5
```

The trap group **rmon-trap-group** is configured to send RMON traps to 192.168.5.5.

## Configuring an Event

To configure an event:

1. Configure an event index, community name, and type:

```
[edit snmp rmon]
user@switch# set event index community community-name typetype
For example:
```

```
[edit snmp rmon]
user@switch# set event 1 community rmon-trap-group type log-and-trap
```

The event community corresponds to the SNMP trap group and is not the same as an SNMP community. This event generates an SNMP trap and adds an entry to the **logTable** in the RMON MIB.

2. Configure a description for the event:

```
[edit snmp rmon]
user@switch# set event index description description
```

For example:

```
[edit snmp rmon]
user@switch# set event 1 description "rmon event"
```

## Configuring an Alarm

---

To configure an alarm:

1. Configure an alarm index, the variable to monitor, the rising and falling thresholds, and the corresponding rising and falling events:

```
[edit snmp rmon]
user@switch# set alarm index variable oid-variable falling-threshold integer rising-threshold
integer rising-event-index index falling-event-index index
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 variable .1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0 falling-threshold 75
rising-threshold 90 rising-event-index 1 falling-event-index 1
```

The variable .1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0 corresponds to the `jnxRmon` MIB object `jnxOperatingCPU`, which represents the CPU utilization of the Routing Engine. The falling and rising threshold integers are 75 and 90. The rising and falling events both generate the same event (event index 1).

2. Configure the sample interval and type and the alarm type:

```
[edit snmp rmon]
user@switch# set alarm index interval seconds sample-type (absolute-value | delta-value)
startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm)
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 interval 30 sample-type absolute-value
startup-alarm rising-or-falling-alarm
```

The absolute value of the monitored variable is sampled every 30 seconds. The initial alarm can occur because of rising above the rising threshold or falling below the falling threshold.

### Related Documentation

- [Configuring SNMP on page 1703](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [Monitoring RMON MIB Tables on page 6389](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6121](#)
- [Understanding RMON on page 6119](#)

## Configuring Health Monitoring

This topic describes how to configure the health monitor feature for QFX Series devices.

The health monitor feature extends the SNMP RMON alarm infrastructure to provide predefined monitoring for a selected set of object instances (such as file system usage, CPU usage, and memory usage) and dynamic object instances (such as Junos OS processes).

To configure health monitoring:

1. Configure the health monitor:

```
[edit snmp]
user@switch# set health-monitor
```

2. Configure the falling threshold:

```
[edit snmp]
user@switch# set health-monitor falling-threshold percentage
```

For example:

```
user@switch# set health-monitor falling-threshold 85
```

3. Configure the rising threshold:

```
[edit snmp]
user@switch# set health-monitor rising-threshold percentage
```

For example:

```
user@switch# set health-monitor rising-threshold 75
```

4. Configure the interval:

```
[edit snmp]
user@switch# set health-monitor interval seconds
```

For example:

```
user@switch# set health-monitor interval 600
```

#### Related Documentation

- [Understanding Health Monitoring on page 6123](#)
- [falling-threshold on page 1772](#)
- [interval \(Health Monitor\) on page 1777](#)
- [rising-threshold \(Health Monitor\) on page 1807](#)

## Creating SNMPv3 Users

For each SNMPv3 user, you can specify the username, authentication type, authentication password, privacy type, and privacy password. After a user enters a password, a key based on the engine ID and password is generated and is written to the configuration file. After the generation of the key, the password is deleted from this configuration file.



**NOTE:** You can configure only one encryption type for each SNMPv3 user.

To create users, include the **user** statement at the **[edit snmp v3 usm local-engine]** hierarchy level:

```
[edit snmp v3 usm local-engine]
user username;
```

**username** is the name that identifies the SNMPv3 user.

To configure user authentication and encryption, include the following statements at the `[edit snmp v3 usm local-engine user username]` hierarchy level:

```
[edit snmp v3 usm local-engine user username]
authentication-md5 {
 authentication-password authentication-password;
}
authentication-sha {
 authentication-password authentication-password;
}
authentication-none;
privacy-aes128 {
 privacy-password privacy-password;
}
privacy-des {
 privacy-password privacy-password;
}
privacy-3des {
 privacy-password privacy-password;
}
privacy-none;
```

**Related  
Documentation**

- [Complete SNMPv3 Configuration Statements](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6118](#)
- [Example: Creating SNMPv3 Users Configuration](#)
- [Example: SNMPv3 Configuration](#)

## Configuring Access Privileges for a Group

In SNMPv3, you can configure a group that sets the same access privileges for one or more users. Configuring a group includes defining the security model and security level, and associating one or more MIB view permissions for the group.



**NOTE:** You must associate at least one MIB view with the group. You can associate multiple MIB views (read, notify, write) to authorize different permissions based on the view. The view name cannot exceed 32 characters.

To configure access privileges for a group:

1. To configure the group:

```
[edit snmp v3 vacm access]
user@switch# edit group group-name
```

2. To configure the context prefix of the SNMP instance for the group:

```
[edit snmp v3 vacm access group group-name]
user@switch# edit (default-context-prefix | context-prefix context-prefix)
```

For example, to configure the default context prefix:

```
[edit snmp v3 vacm access group group-name]
```

```
user@switch# edit default-context-prefix
```

3. To configure the security model:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model (any | usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model usm
```

4. To configure the security level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level (authentication | none | privacy)
```

For example, to configure a security level requiring user authentication and encryption:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level privacy
```



**NOTE:** Access privileges are granted to all packets with a security level equal to or greater than that configured. If you are configuring the SNMPv1 or v2c security model, use *none* as your security level. If you are configuring the SNMPv3 security model (USM), use the *authentication*, *none*, or *privacy* security level.

5. (Optional) To associate a read-only MIB view with an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit read-view view-name
```

6. (Optional) To associate a MIB view with an SNMP notification permission for an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit notify-view view-name
```

7. (Optional) To associate a MIB view with write permission for an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit write-view view-name
```

## Assigning a Security Name to a Group

In SNMPv3, each username is associated with a security name. The security name, together with the SNMP engine ID, is included in SNMP messages to ensure messaging security.

Before you assign a security name to a group, first create the security name. For an SNMPv3 client, the security name is the username configured at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level. For SNMPv1 or v2c clients, the security name is the community string configured at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level.

Assigning a security name to a group includes configuring a security model for the group, assigning the security name to the group, and configuring the group.

To assign an SNMP security name to a group:

1. To configure a security model for the group:

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model (usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model usm
```

2. To associate the security name with a group:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]
user@switch# edit security-name security-name
```

3. To configure a group of SNMPv3 security names with the same security policy:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name
security-name]
user@switch# edit group group-name
```

### Related Documentation

- [Creating SNMPv3 Users on page 6201](#)
- [group \(Associating a Security Name\) on page 6300](#)
- [security-model \(Group\) on page 6333](#)
- [security-name \(Community String\) on page 6335](#)
- [security-name \(Security Group\) on page 6336](#)

## Configuring SNMPv3 Traps on a Device Running Junos OS

In SNMPv3, you create traps and informs by configuring the **notify**, **target-address**, and **target-parameters** parameters. Traps are unconfirmed notifications, whereas informs are confirmed notifications. This section describes how to configure SNMP traps. For information about configuring SNMP informs, see [“Configuring SNMP Informs” on page 6206](#).



The target address defines a management application's address and parameters to be used in sending notifications. Target parameters define the message processing and security parameters that are used in sending notifications to a particular management target. SNMPv3 also lets you define SNMPv1 and SNMPv2c traps.



**NOTE:** When you configure SNMP traps, make sure your configured access privileges allow the traps to be sent. Access privileges are configured at the `[edit snmp v3 vacm access]` and `[edit snmp v3 vacm security-to-group]` hierarchy levels.

To configure SNMP traps, include the following statements at the `[edit snmp v3]` hierarchy level:

```
[edit snmp v3]
 notify name {
 tag tag-name;
 type trap;
 }
 notify-filter name {
 oid object-identifier (include | exclude);
 }
 target-address target-address-name {
 address address;
 address-mask address-mask;
 logical-system (SNMP) logical-system;
 port port-number;
 routing-instance instance;
 tag-list tag-list;
 target-parameters target-parameters-name;
 }
 target-parameters target-parameters-name {
 notify-filter profile-name;
 parameters {
 message-processing-model (v1 | v2c | v3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 security-name security-name;
 }
 }
}
```

#### Related Documentation

- [Configuring the SNMPv3 Trap Notification](#)
- [Configuring the Trap Notification Filter](#)
- [Configuring the Trap Target Address](#)
- [Defining and Configuring the Trap Target Parameters](#)
- [Configuring SNMP Informs on page 6206](#)
- [Configuring the Remote Engine and Remote User](#)
- [Configuring the Inform Notification Type and Target Address](#)
- [Complete SNMPv3 Configuration Statements](#)

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6118](#)

## Configuring SNMP Informs

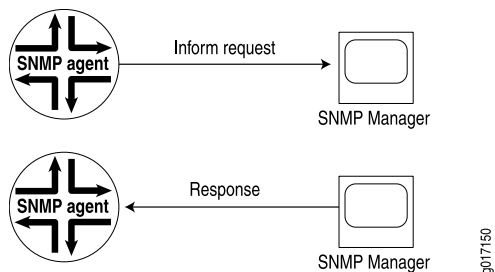
Junos OS supports two types of notifications: traps and informs. With traps, the receiver does not send any acknowledgment when it receives a trap. Therefore, the sender cannot determine if the trap was received. A trap may be lost because a problem occurred during transmission. To increase reliability, an inform is similar to a trap except that the inform is stored and retransmitted at regular intervals until one of these conditions occurs:

- The receiver (target) of the inform returns an acknowledgment to the SNMP agent.
- A specified number of unsuccessful retransmissions have been attempted and the agent discards the inform message.

If the sender never receives a response, the inform can be sent again. Thus, informs are more likely to reach their intended destination than traps are. Informs use the same communications channel as traps (same socket and port) but have different protocol data unit (PDU) types.

Informs are more reliable than traps, but they consume more network, router, and switch resources (see [Figure 226 on page 6206](#)). Unlike a trap, an inform is held in memory until a response is received or the timeout is reached. Also, traps are sent only once, whereas an inform may be retried several times. Use informs when it is important that the SNMP manager receive all notifications. However, if you are more concerned about network traffic, or router and switch memory, use traps.

**Figure 226: Inform Request and Response**



For information about configuring SNMP traps, see [“Configuring SNMPv3 Traps on a Device Running Junos OS” on page 6204](#).

### Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6204](#)
- [Configuring the Remote Engine and Remote User](#)
- [Configuring the Inform Notification Type and Target Address](#)
- [Complete SNMPv3 Configuration Statements](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6118](#)

## Configuration Tasks for System Log Messages

- [Junos OS Minimum System Logging Configuration on page 6207](#)
- [Junos OS System Log Configuration Statements on page 6208](#)
- [Adding a Text String to System Log Messages on page 6209](#)
- [Directing System Log Messages to a Log File on page 6209](#)
- [Directing System Log Messages to a Remote Machine on page 6210](#)
- [Directing System Log Messages to a User Terminal on page 6211](#)
- [Directing System Log Messages to the Console on page 6211](#)
- [Disabling the System Logging of a Facility on page 6212](#)
- [Displaying a Log File from a Single-Chassis System on page 6212](#)
- [Including Priority Information in System Log Messages on page 6213](#)
- [Including the Year or Millisecond in Timestamps on page 6215](#)
- [Logging Messages in Structured-Data Format on page 6215](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6216](#)
- [Interpreting Messages Generated in Standard Format on page 6219](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6220](#)
- [Specifying the Facility and Severity of Messages to Include in the Log on page 6222](#)
- [Junos OS System Logging Facilities and Message Severity Levels on page 6222](#)
- [System Log Default Facilities for Messages Directed to a Remote Destination on page 6224](#)
- [Junos OS System Log Alternate Facilities for Remote Logging on page 6224](#)
- [Changing the Alternative Facility Name for Remote System Log Messages on page 6225](#)
- [Using Regular Expressions to Refine the Set of Logged Messages on page 6227](#)

### Junos OS Minimum System Logging Configuration

To record or view system log messages, you must include the **syslog** statement at the **[edit system]** hierarchy level. Specify at least one destination for the messages, as described in [Table 603 on page 6207](#). For more information about the configuration statements, see *Single-Chassis System Logging Configuration Overview*.

**Table 603: Minimum Configuration Statements for System Logging**

| Destination                                    | Minimum Configuration Statements                                             |
|------------------------------------------------|------------------------------------------------------------------------------|
| File                                           | <pre>[edit system syslog] file filename {   facility severity; }</pre>       |
| Terminal session of one, several, or all users | <pre>[edit system syslog] user (username   *) {   facility severity; }</pre> |

**Table 603: Minimum Configuration Statements for System Logging (*continued*)**

| Destination                                                        | Minimum Configuration Statements                                                                            |
|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Router or switch console                                           | [edit system syslog]<br>console {<br><i>facility severity</i> ;<br>}                                        |
| Remote machine or the other Routing Engine on the router or switch | [edit system syslog]<br>host ( <i>hostname</i>   other-routing-engine) {<br><i>facility severity</i> ;<br>} |

**Related Documentation**

- [Junos OS System Log Configuration Overview](#)
- [Overview of Junos OS System Log Messages on page 6154](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

**Junos OS System Log Configuration Statements**

To configure the switch to log system messages, include the **syslog** statement at the **[edit system]** hierarchy level:

```
[edit system]
syslog {
 archive <files number> <size size> <world-readable | no-world-readable>;
 console {
 facility severity;
 }
 file filename {
 facility severity;
 archive <archive-sites (ftp-url <password password>)> <files number> <size size>
 <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
 no-world-readable>;
 explicit-priority;
 match "regular-expression";
 structured-data {
 brief;
 }
 }
 host hostname {
 facility severity;
 explicit-priority;
 facility-override facility;
 log-prefix string
 match "regular-expression";
 }
 source-address source-address;
 time-format (year | millisecond | year millisecond);
 user (username | *) {
 facility severity;
 match "regular-expression";
 }
}
```

```
 }
}
```

**Related  
Documentation**

- [Overview of Junos OS System Log Messages on page 6154](#)

## Adding a Text String to System Log Messages

To add a text string to every system log message directed to a remote machine or to the other Routing Engine, include the **log-prefix** statement at the **[edit system syslog host]** hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]
 facility severity;
 log-prefix string;
```

The string can contain any alphanumeric or special character except the equal sign ( = ) and the colon ( : ). It also cannot include the space character; do not enclose the string in quotation marks ( " ") in an attempt to include spaces in it.

The Junos OS system logging utility automatically appends a colon and a space to the specified string when the system log messages are written to the log. The string is inserted after the identifier for the Routing Engine that generated the message.

The following example shows how to add the string M120 to all messages to indicate that the router is an M120 router, and direct the messages to the remote machine hardware-logger.mycompany.com:

```
[edit system syslog]
 host hardware-logger.mycompany.com {
 any info;
 log-prefix M120;
 }
```

When these configuration statements are included on an M120 router called origin1, a message in the system log on hardware-logger.mycompany.com looks like the following:

```
Mar 9 17:33:23 origin1 M120:mgd[477]: UI_CMDLINE_READ_LINE: user 'root', command 'run
show version'
```

**Related  
Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6220](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Directing System Log Messages to a Log File

To direct system log messages to a file in the **/var/log** directory of the local Routing Engine, include the **file** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
 file filename {
 facility severity;
```

```
archive <archive-sites (ftp-url <password password>) > <files number> <size size>
 <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
 no-world-readable>;
explicit-priority;
match "regular-expression";
structured-data {
 brief;
}
}
```

For the list of facilities and severity levels, see *Specifying the Facility and Severity of Messages to Include in the Log*.

To prevent log files from growing too large, the Junos OS system logging utility by default writes messages to a sequence of files of a defined size. By including the **archive** statement, you can configure the number of files, their maximum size, and who can read them, either for all log files or for a certain log file. For more information, see [“Specifying Log File Size, Number, and Archiving Properties”](#) on page 6220.

For information about the following statements, see the indicated sections:

- **explicit-priority**—See [“Including Priority Information in System Log Messages”](#) on page 6213
- **match**—See [“Using Regular Expressions to Refine the Set of Logged Messages”](#) on page 6227
- **structured-data**—See *Logging Messages in Structured-Data Format*

#### Related Documentation

- *Single-Chassis System Logging Configuration Overview*
- [Overview of Junos OS System Log Messages on page 6154](#)
- [Logging Messages in Structured-Data Format on page 6215](#)
- *Examples: Configuring System Logging*
- [Examples: Configuring System Logging on page 6159](#)

## Directing System Log Messages to a Remote Machine

To direct system log messages to a remote machine, include the **host** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
host (hostname | other-routing-engine) {
 facility severity;
 explicit-priority;
 facility-override facility;
 log-prefix string;
 match "regular-expression";
}
source-address source-address;
```

To direct system log messages to a remote machine, include the **host hostname** statement to specify the remote machine's IP version 4 (IPv4) address or fully qualified hostname. The remote machine must be running the standard **syslogd** utility. We do not recommend

directing messages to another Juniper Networks switch. In each system log message directed to the remote machine, the hostname of the local Routing Engine appears after the timestamp to indicate that it is the source for the message.

For the list of logging facilities and severity levels to configure under the **host** statement, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6222](#).

To record facility and severity level information in each message, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 6213](#).

For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6227](#).

When directing messages to remote machines, you can include the **source-address** statement to specify the IP address of the switch that is reported in the messages as their source. In each **host** statement, you can also include the **facility-override** statement to assign an alternative facility and the **log-prefix** statement to add a string to each message.

**Related  
Documentation**

- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Directing System Log Messages to a User Terminal

To direct system log messages to the terminal session of one or more specific users (or all users) when they are logged in to the local Routing Engine, include the **user** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
user (username | *) {
 facility severity;
 match "regular-expression";
}
```

Specify one or more Junos OS usernames, separating multiple values with spaces, or use the asterisk (\*) to indicate all users who are logged in to the local Routing Engine.

For the list of logging facilities and severity levels, see *Specifying the Facility and Severity of Messages to Include in the Log*. For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6227](#).

**Related  
Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6159](#)

## Directing System Log Messages to the Console

To direct system log messages to the console of the local Routing Engine, include the **console** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
console {
 facility severity;
}
```

For the list of logging facilities and severity levels, see *Specifying the Facility and Severity of Messages to Include in the Log*.

**Related  
Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6159](#)

## Disabling the System Logging of a Facility

To disable the logging of messages that belong to a particular facility, include the **facility none** statement in the configuration. This statement is useful when, for example, you want to log messages that have the same severity level and belong to all but a few facilities. Instead of including a statement for each facility you want to log, you can include the **any severity** statement and then a **facility none** statement for each facility that you do not want to log. For example, the following logs all messages at the **error** level or higher to the console, except for messages from the **daemon** and **kernel** facilities. Messages from those facilities are logged to the file `>/var/log/internals` instead:

```
[edit system syslog]
console {
 any error;
 daemon none;
 kernel none;
}
file internals {
 daemon info;
 kernel info;
}
```

**Related  
Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system such as the QFX3500 switch, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
```



```

Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysAppElemRunMemory.5.6.2293)
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysAppElemRunMemory.5.8.2292)
...
Nov 4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov 4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov 4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages

```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```

user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...

```

#### Related Documentation

- [Interpreting Messages Generated in Standard Format on page 6219](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6216](#)

## Including Priority Information in System Log Messages

The facility and severity level of a message are together referred to as its *priority*. By default, messages logged in the standard Junos OS format do not include information about priority. To include priority information in standard-format messages directed to a file, include the **explicit-priority** statement at the **[edit system syslog file *filename*]** hierarchy level:

```

[edit system syslog file filename]
 facility severity;
 explicit-priority;

```



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

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

To include priority information in messages directed to a remote machine or the other Routing Engine, include the `explicit-priority` statement at the `[edit system syslog host (hostname | other-routing-engine)]` hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]
 facility severity;
 explicit-priority;
```



**NOTE:** The `other-routing-engine` option does not apply to the QFX Series.

The priority recorded in a message always indicates the original, local facility name. If the `facility-override` statement is included for messages directed to a remote destination, the Junos OS system logging utility still uses the alternative facility name for the messages themselves when directing them to the remote destination. For more information, see [“Changing the Alternative Facility Name for Remote System Log Messages” on page 6225](#).

When the `explicit-priority` statement is included, the Junos OS logging utility prepends codes for the facility name and severity level to the message tag name, if the message has one:

`FACILITY-severity[-TAG]`

(The tag is a unique identifier assigned to some Junos OS system log messages; for more information, see the *Junos OS System Log Messages Reference*.)

In the following example, the `CHASSISD_PARSE_COMPLETE` message belongs to the `daemon` facility and is assigned severity `info` (6):

```
Aug 21 12:36:30 router1 chassisd[522]: %DAEMON-6-CHASSISD_PARSE_COMPLETE:
 Using new configuration
```

When the `explicit-priority` statement is not included, the priority does not appear in the message:

```
Aug 21 12:36:30 router1 chassisd[522]: CHASSISD_PARSE_COMPLETE: Using new
 configuration
```

For more information about message formatting, see the *Junos OS System Log Messages Reference*.

#### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Examples: Configuring System Logging](#)

### Including the Year or Millisecond in Timestamps

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

```
Aug 21 15:36:30
```

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

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

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

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2010):

```
Aug 21 15:36:30.401 2010
```



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

For information about the **structured-data** statement, see “[Logging Messages in Structured-Data Format](#)” on page 6215. For information about interpreting messages in a structured-data format, see “[Interpreting Messages Generated in Structured-Data Format](#)” on page 6216.

### Logging Messages in Structured-Data Format

You can log messages to a file in structured-data format instead of the standard Junos OS format. The structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.

The structured-data format complies with Internet draft draft-ietf-syslog-protocol-21.txt. The draft establishes a standard message format regardless of the source or transport protocol for logged messages.

To output messages to a file in structured-data format, include the **structured-data** statement at the `[edit system syslog file filename]` hierarchy level:

```
[edit system syslog file filename]
 facility severity;
 structured-data {
 brief;
 }
```

The optional **brief** statement suppresses the English-language text that appears by default at the end of a message to describe the error or event. For information about the fields in a structured-data-format message, see [“Interpreting Messages Generated in Structured-Data Format” on page 6216](#).

The structured format is used for all messages logged to the file that are generated by a Junos OS process or software library.



**NOTE:** If you include either or both of the **explicit-priority** and **time-format** statements along with the **structured-data** statement, they are ignored. These statements apply to the standard Junos OS system log format, not to structured-data format.

---

## Interpreting Messages Generated in Structured-Data Format

By default, Junos OS processes and software libraries write messages to the system log file in structured-data format. For information about the **structured-data** statement, see *Logging Messages in Structured-Data Format*.

Structured-format makes it easier for automated applications to extract information from the message. In particular, the standardized format for reporting the value of variables (elements in the English-language message that vary depending on the circumstances that triggered the message) makes it easy for an application to extract those values.

The structured-data format for a message includes the following fields (which appear here on two lines only for legibility):

```
<priority code>version timestamp hostname process processID TAG [junos@2636.platform
variable-value-pairs] message-text
```

[Table 604 on page 6217](#) describes the fields. If the system logging utility cannot determine the value in a particular field, a hyphen ( - ) appears instead.

Table 604: Fields in Structured-Data Messages

| Field                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Examples                                                                                                                                                                     |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>&lt;priority code&gt;</b> | Number that indicates the facility and severity of a message. It is calculated by multiplying the facility number by 8 and then adding the numerical value of the severity. For a mapping of the numerical codes to facility and severity, see <i>Specifying the Facility and Severity of Messages to Include in the Log</i> .                                                                                                                                                                        | <b>&lt;165&gt;</b> for a message from the <b>pfe</b> facility (facility=20) with severity <b>notice</b> (severity=5).                                                        |
| <b>version</b>               | Version of the Internet Engineering Task Force (IETF) system logging protocol specification.                                                                                                                                                                                                                                                                                                                                                                                                          | 1 for the initial version                                                                                                                                                    |
| <b>timestamp</b>             | Time when the message was generated, in one of two representations: <ul style="list-style-type: none"> <li><b>YYYY-MM-DDTHH:MM:SS.MSZ</b> is the year, month, day, hour, minute, second and millisecond in Universal Coordinated Time (UTC)</li> <li><b>YYYY-MM-DDTHH:MM:SS.MS+/-HH:MM</b> is the year, month, day, hour, minute, second and millisecond in local time; the hour and minute that follows the plus sign (+) or minus sign (-) is the offset of the local time zone from UTC</li> </ul> | 2007-02-15T09:17:15.719Z is 9:17 AM UTC on 15 February 2007.<br>2007-02-15T01:17:15.719-08:00 is the same timestamp expressed as Pacific Standard Time in the United States. |
| <b>hostname</b>              | Name of the host that originally generated the message.                                                                                                                                                                                                                                                                                                                                                                                                                                               | switch1                                                                                                                                                                      |
| <b>process</b>               | Name of the Junos OS process that generated the message.                                                                                                                                                                                                                                                                                                                                                                                                                                              | mgd                                                                                                                                                                          |
| <b>processID</b>             | UNIX process ID (PID) of the Junos process that generated the message.                                                                                                                                                                                                                                                                                                                                                                                                                                | 3046                                                                                                                                                                         |
| <b>TAG</b>                   | Junos OS system log message tag, which uniquely identifies the message.                                                                                                                                                                                                                                                                                                                                                                                                                               | UI_DBASE_LOGOUT_EVENT                                                                                                                                                        |
| <b>junos@2636.platform</b>   | An identifier for the type of hardware platform that generated the message. The junos@2636 prefix indicates that the platform runs the Junos OS. It is followed by a dot-separated numerical identifier for the platform type.                                                                                                                                                                                                                                                                        | junos@2636.1.1.1.2.18                                                                                                                                                        |
| <b>variable-value-pairs</b>  | A variable-value pair for each element in the <i>message-text</i> string that varies depending on the circumstances that triggered the message. Each pair appears in the format <b>variable = "value"</b> .                                                                                                                                                                                                                                                                                           | username="regress"                                                                                                                                                           |

Table 604: Fields in Structured-Data Messages (*continued*)

| Field               | Description                                                                                                                                                                       | Examples                                  |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <i>message-text</i> | English-language description of the event or error (omitted if the brief statement is included at the [edit system syslog file <i>filename</i> structured-data] hierarchy level). | User 'regress' exiting configuration mode |

By default, the structured-data version of a message includes English text at the end, as in the following example (which appears on multiple lines only for legibility):

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"] User 'regress' exiting configuration mode
```

When the brief statement is included at the [edit system syslog file *filename* structured-data ] hierarchy level, the English text is omitted, as in this example:

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"]
```

Table 605 on page 6218 maps the codes that appear in the *priority-code* field to facility and severity level.



**NOTE:** Not all of the facilities and severities listed in Table 605 on page 6218 can be included in statements at the [edit system syslog] hierarchy level (some are used by internal processes). For a list of the facilities and severity levels that can be included in the configuration, see “Specifying the Facility and Severity of Messages to Include in the Log” on page 6222.

Table 605: Facility and Severity Codes in the priority-code Field

| Facility (number) | Severity<br>emergency | alert | critical | error | warning | notice | info | debug |
|-------------------|-----------------------|-------|----------|-------|---------|--------|------|-------|
| kernel (0)        | 1                     | 1     | 2        | 3     | 4       | 5      | 6    | 7     |
| user (1)          | 8                     | 9     | 10       | 11    | 12      | 13     | 14   | 15    |
| mail (2)          | 16                    | 17    | 18       | 19    | 20      | 21     | 22   | 23    |
| daemon (3)        | 24                    | 25    | 26       | 27    | 28      | 29     | 30   | 31    |
| authorization (4) | 32                    | 33    | 34       | 35    | 36      | 37     | 38   | 39    |
| syslog (5)        | 40                    | 41    | 42       | 43    | 44      | 45     | 46   | 47    |
| printer (6)       | 48                    | 49    | 50       | 51    | 52      | 53     | 54   | 55    |
| news (7)          | 56                    | 57    | 58       | 59    | 60      | 61     | 62   | 63    |
| uucp (8)          | 64                    | 65    | 66       | 67    | 68      | 69     | 70   | 71    |

**Table 605: Facility and Severity Codes in the priority-code Field (*continued*)**

| Facility (number)          | Severity<br>emergency | alert | critical | error | warning | notice | info | debug |
|----------------------------|-----------------------|-------|----------|-------|---------|--------|------|-------|
| clock (9)                  | 72                    | 73    | 74       | 75    | 76      | 77     | 78   | 79    |
| authorization-private (10) | 80                    | 81    | 82       | 83    | 84      | 85     | 86   | 87    |
| ftp (11)                   | 88                    | 89    | 90       | 91    | 92      | 93     | 94   | 95    |
| ntp (12)                   | 96                    | 97    | 98       | 99    | 100     | 101    | 102  | 103   |
| security (13)              | 104                   | 105   | 106      | 107   | 108     | 109    | 110  | 111   |
| console (14)               | 112                   | 113   | 114      | 115   | 116     | 117    | 118  | 119   |
| local0 (16)                | 128                   | 129   | 130      | 131   | 132     | 133    | 134  | 135   |
| dfc (17)                   | 136                   | 137   | 138      | 139   | 140     | 141    | 142  | 143   |
| local2 (18)                | 144                   | 145   | 146      | 147   | 148     | 149    | 150  | 151   |
| firewall (19)              | 152                   | 153   | 154      | 155   | 156     | 157    | 158  | 159   |
| pfe (20)                   | 160                   | 161   | 162      | 163   | 164     | 165    | 166  | 167   |
| conflict-log (21)          | 168                   | 169   | 170      | 171   | 172     | 173    | 174  | 175   |
| change-log (22)            | 176                   | 177   | 178      | 179   | 180     | 181    | 182  | 183   |
| interactive-commands (23)  | 184                   | 185   | 186      | 187   | 188     | 189    | 190  | 191   |

## Interpreting Messages Generated in Standard Format

The syntax of a standard-format message generated by a Junos OS process or subroutine library depends on whether it includes priority information:

- When the **explicit-priority** statement is included at the `[edit system syslog file filename]` or `[edit system syslog host hostname]` hierarchy level, a system log message has the following syntax:

```
timestamp message-source: %facility-severity-TAG: message-text
```

- When directed to the console or to users, or when the **explicit-priority** statement is not included for files or remote hosts, a system log message has the following syntax:

```
timestamp message-source: TAG: message-text
```

Table 606 on page 6220 describes the message fields.

Table 606: Fields in Standard-Format Messages

| Field                 | Description                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>timestamp</i>      | Time at which the message was logged.                                                                                                                                                                                                                                                                                                                                                           |
| <i>message-source</i> | Identifier of the process or component that generated the message and the routing platform on which the message was logged. This field includes two or more subfields: hostname, process and process ID (PID). If the process does not report its PID, the PID is not displayed. The message source subfields are displayed in the following format:<br><br><i>hostname process[process-ID]</i> |
| <i>facility</i>       | Code that specifies the facility to which the system log message belongs. For a mapping of codes to facility names, see Table: <b>Facility Codes Reported in Priority Information</b> in “Including Priority Information in System Log Messages” on page 6213.                                                                                                                                  |
| <i>severity</i>       | Numerical code that represents the severity level assigned to the system log message. For a mapping of codes to severity names, see Table: <b>Numerical Codes for Severity Levels Reported in Priority Information</b> in “Including Priority Information in System Log Messages” on page 6213.                                                                                                 |
| <i>TAG</i>            | Text string that uniquely identifies the message, in all uppercase letters and using the underscore (_) to separate words. The tag name begins with a prefix that indicates the generating software process or library. The entries in this reference are ordered alphabetically by this prefix.<br><br>Not all processes on a routing platform use tags, so this field does not always appear. |
| <i>message-text</i>   | Text of the message.                                                                                                                                                                                                                                                                                                                                                                            |

## Specifying Log File Size, Number, and Archiving Properties

To prevent log files from growing too large, by default the Junos system logging utility writes messages to a sequence of files of a defined size. The files in the sequence are referred to as *archive* files to distinguish them from the *active* file to which messages are currently being written. The default maximum size depends on the platform type:

- 128 kilobytes (KB) for EX Series switches and J Series routers
- 1 megabyte (MB) for M Series, MX Series, and T Series routers
- 10 MB for TX Matrix or TX Matrix Plus routers
- 1 MB for the QFX Series

When an active log file called **logfile** reaches the maximum size, the logging utility closes the file, compresses it, and names the compressed archive file **logfile.0.gz**. The logging utility then opens and writes to a new active file called **logfile**. This process is also known as file rotation. When the new **logfile** reaches the configured maximum size, **logfile.0.gz** is renamed **logfile.1.gz**, and the new **logfile** is closed, compressed, and renamed **logfile.0.gz**. By default, the logging utility creates up to 10 archive files in this manner. When the maximum number of archive files is reached and when the size of the active file reaches



the configured maximum size, the contents of the last archived file are overwritten by the current active file. The logging utility by default also limits the users who can read log files to the **root** user and users who have the Junos OS **maintenance** permission.

Junos OS provides a configuration statement **log-rotate-frequency** that configures the system log file rotation frequency by configuring the time interval for checking the log file size. The frequency can be set to a value of 1 minute through 59 minutes. The default frequency is 15 minutes.

To configure the log rotation frequency, include the **log-rotate-frequency** statement at the **[edit system syslog]** hierarchy level.

You can include the **archive** statement to change the maximum size of each file, how many archive files are created, and who can read log files.

To configure values that apply to all log files, include the **archive** statement at the **[edit system syslog]** hierarchy level:

```
archive <files number> <size size> <world-readable | no-world-readable>;
```

To configure values that apply to a specific log file, include the **archive** statement at the **[edit system syslog file *filename*]** hierarchy level:

```
archive <archive-sites (ftp-url <password password>)> <files number> <size size>
 <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
 no-world-readable>;
```

**archive-sites *site-name*** specifies a list of archive sites that you want to use for storing files. The ***site-name*** value is any valid FTP URL to a destination. If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the specified log filename. For information about how to specify valid FTP URLs, see [“Format for Specifying Filenames and URLs in Junos OS CLI Commands” on page 73](#).

**binary-data** Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems). To restore the default setting, include the **no-binary-data** statement.

**files *number*** specifies the number of files to create before the oldest file is overwritten. The value can be from 1 through 1000.

**size *size*** specifies the maximum size of each file. The value can be from 64 KB (64k) through 1 gigabyte (1g); to represent megabytes, use the letter **m** after the integer. There is no space between the digits and the **k**, **m**, or **g** units letter.

**start-time "YYYY-MM-DD.hh:mm"** defines the date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

**transfer-interval *interval*** defines the amount of time the current log file remains open (even if it has not reached the maximum possible size) and receives new statistics before

it is closed and transferred to an archive site. This interval value can be from 5 through 2880 minutes.

**world-readable** enables all users to read log files. To restore the default permissions, include the **no-world-readable** statement.

#### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Examples: Configuring System Logging](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

## Specifying the Facility and Severity of Messages to Include in the Log

Each system log message belongs to a *facility*, which is a group of messages that are either generated by the same software process or concern a similar condition or activity (such as authentication attempts). Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects router functions.

When you configure logging for a facility and destination, you specify a severity level for each facility. Messages from the facility that are rated at that level or higher are logged to the destination:

```
[edit system syslog]
(console | file filename | host destination | user username) {
 facility severity;
}
```

#### Related Documentation

- [Junos OS System Logging Facilities and Message Severity Levels on page 6222](#)
- [Single-Chassis System Logging Configuration Overview](#)
- [Examples: Configuring System Logging](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Junos OS System Logging Facilities and Message Severity Levels

Table 607 on page 6222 lists the Junos system logging facilities that you can specify in configuration statements at the **[edit system syslog]** hierarchy level.

Table 607: Junos OS System Logging Facilities

| Facility             | Type of Event or Error                                |
|----------------------|-------------------------------------------------------|
| <b>any</b>           | All (messages from all facilities)                    |
| <b>authorization</b> | Authentication and authorization attempts             |
| <b>change-log</b>    | Changes to the Junos OS configuration                 |
| <b>conflict-log</b>  | Specified configuration is invalid on the router type |

Table 607: Junos OS System Logging Facilities (*continued*)

| Facility                    | Type of Event or Error                                                                                                                            |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>daemon</b>               | Actions performed or errors encountered by system processes                                                                                       |
| <b>dfc</b>                  | Events related to dynamic flow capture                                                                                                            |
| <b>firewall</b>             | Packet filtering actions performed by a firewall filter                                                                                           |
| <b>ftp</b>                  | Actions performed or errors encountered by the FTP process                                                                                        |
| <b>interactive-commands</b> | Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client |
| <b>kernel</b>               | Actions performed or errors encountered by the Junos OS kernel                                                                                    |
| <b>pfe</b>                  | Actions performed or errors encountered by the Packet Forwarding Engine                                                                           |
| <b>user</b>                 | Actions performed or errors encountered by user-space processes                                                                                   |

Table 608 on page 6223 lists the severity levels that you can specify in configuration statements at the **[edit system syslog]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Unlike the other severity levels, the **none** level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see “Disabling the System Logging of a Facility” on page 6212.

Table 608: System Log Message Severity Levels

| Severity Level   | Description                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>any</b>       | Includes all severity levels                                                                                            |
| <b>none</b>      | Disables logging of the associated facility to a destination                                                            |
| <b>emergency</b> | System panic or other condition that causes the router to stop functioning                                              |
| <b>alert</b>     | Conditions that require immediate correction, such as a corrupted system database                                       |
| <b>critical</b>  | Critical conditions, such as hard errors                                                                                |
| <b>error</b>     | Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels |
| <b>warning</b>   | Conditions that warrant monitoring                                                                                      |
| <b>notice</b>    | Conditions that are not errors but might warrant special handling                                                       |

Table 608: System Log Message Severity Levels (*continued*)

| Severity Level | Description                               |
|----------------|-------------------------------------------|
| info           | Events or nonerror conditions of interest |

**Related Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Examples: Configuring System Logging](#)

## System Log Default Facilities for Messages Directed to a Remote Destination

Table 609 on page 6224 lists the default alternative facility name next to the Junos OS-specific facility name for which it is used. For facilities that are not listed, the default alternative name is the same as the local facility name.

Table 609: Default Facilities for Messages Directed to a Remote Destination

| Junos OS-specific Local Facility | Default Facility When Directed to Remote Destination |
|----------------------------------|------------------------------------------------------|
| change-log                       | local6                                               |
| conflict-log                     | local5                                               |
| dfc                              | local1                                               |
| firewall                         | local3                                               |
| interactive-commands             | local7                                               |
| pfe                              | local4                                               |

**Related Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Junos OS System Log Alternate Facilities for Remote Logging

Table 610 on page 6224 lists the facilities that you can specify in the **facility-override** statement.

Table 610: Facilities for the facility-override Statement

| Facility      | Description                                                 |
|---------------|-------------------------------------------------------------|
| authorization | Authentication and authorization attempts                   |
| daemon        | Actions performed or errors encountered by system processes |

Table 610: Facilities for the facility-override Statement (*continued*)

| Facility      | Description                                                     |
|---------------|-----------------------------------------------------------------|
| <b>ftp</b>    | Actions performed or errors encountered by the FTP process      |
| <b>kernel</b> | Actions performed or errors encountered by the Junos OS kernel  |
| <b>local0</b> | Local facility number 0                                         |
| <b>local1</b> | Local facility number 1                                         |
| <b>local2</b> | Local facility number 2                                         |
| <b>local3</b> | Local facility number 3                                         |
| <b>local4</b> | Local facility number 4                                         |
| <b>local5</b> | Local facility number 5                                         |
| <b>local6</b> | Local facility number 6                                         |
| <b>local7</b> | Local facility number 7                                         |
| <b>user</b>   | Actions performed or errors encountered by user-space processes |

We do not recommend including the **facility-override** statement at the **[edit system syslog host other-routing-engine]** hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

#### Related Documentation

- *Examples: Assigning an Alternative Facility*
- *Single-Chassis System Logging Configuration Overview*
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)

## Changing the Alternative Facility Name for Remote System Log Messages

Some facilities assigned to messages logged on the local router or switch have Junos OS-specific names (see [Table 607 on page 6222](#)). In the recommended configuration, a remote machine designated at the **[edit system syslog host *hostname*]** hierarchy level is not a Juniper Networks router or switch, so its syslogd utility cannot interpret the Junos OS-specific names. To enable the standard syslogd utility to handle messages from these facilities when messages are directed to a remote machine, a standard **localX** facility name is used instead of the Junos OS-specific facility name.

[Table 609 on page 6224](#) lists the default alternative facility name next to the Junos OS-specific facility name it is used for.

The syslogd utility on a remote machine handles all messages that belong to a facility in the same way, regardless of the source of the message (the Juniper Networks router or switch or the remote machine itself). For example, the following statements in the configuration of the router called **local-router** direct messages from the **authorization** facility to the remote machine *monitor.mycompany.com*:

```
[edit system syslog]
host monitor.mycompany.com {
 authorization info;
}
```

The default alternative facility for the local **authorization** facility is also **authorization**. If the syslogd utility on **monitor** is configured to write messages belonging to the **authorization** facility to the file */var/log/auth-attempts*, then the file contains the messages generated when users log in to **local-router** and the messages generated when users log in to **monitor**. Although the name of the source machine appears in each system log message, the mixing of messages from multiple machines can make it more difficult to analyze the contents of the **auth-attempts** file.

To make it easier to separate the messages from each source, you can assign an alternative facility to all messages generated on **local-router** when they are directed to **monitor**. You can then configure the syslogd utility on **monitor** to write messages with the alternative facility to a different file from messages generated on **monitor** itself.

To change the facility used for all messages directed to a remote machine, include the **facility-override** statement at the **[edit system syslog host *hostname*]** hierarchy level:

```
[edit system syslog host hostname]
facility severity;
facility-override facility;
```

In general, it makes sense to specify an alternative facility that is not already in use on the remote machine, such as one of the **localX** facilities. On the remote machine, you must also configure the syslogd utility to handle the messages in the desired manner.

[Table 610 on page 6224](#) lists the facilities that you can specify in the **facility-override** statement.

We do not recommend including the **facility-override** statement at the **[edit system syslog host *other-routing-engine*]** hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

The following example shows how to log all messages generated on the local router at the error level or higher to the local0 facility on the remote machine called *monitor.mycompany.com*:

```
[edit system syslog]
host monitor.mycompany.com {
 any error;
 facility-override local0;
}
```

The following example shows how to configure routers located in California and routers located in New York to send messages to a single remote machine called `central-logger.mycompany.com`. The messages from California are assigned to alternative facility `local0` and the messages from New York are assigned to alternative facility `local2`.

- Configure California routers to aggregate messages in the `local0` facility:

```
[edit system syslog]
host central-logger.mycompany.com {
 change-log info;
 facility-override local0;
}
```

- Configure New York routers to aggregate messages in the `local2` facility:

```
[edit system syslog]
host central-logger.mycompany.com {
 change-log info;
 facility-override local2;
}
```

On `central-logger`, you can then configure the system logging utility to write messages from the `local0` facility to the file `change-log` and the messages from the `local2` facility to the file `new-york-config`.

#### Related Documentation

- [Table 609 on page 6224](#)
- [Junos OS System Log Alternate Facilities for Remote Logging on page 6224](#)
- [Examples: Assigning an Alternative Facility](#)
- [Examples: Assigning an Alternative Facility on page 6161](#)

## Using Regular Expressions to Refine the Set of Logged Messages

The predefined facilities group together related messages, but you can also use regular expression matching to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination.

To specify the text string that must (or must not) appear in a message for the message to be logged to a destination, include the `match` statement and specify the regular expression which the text string must match:

```
match "regular-expression";
```

You can include this statement at the following hierarchy levels:

- `[edit system syslog file filename]` (for a file)
- `[edit system syslog user (username | *)]` (for a specific user session or for all user sessions on a terminal)
- `[edit system syslog host (hostname | other-routing-engine)]` (for a remote destination)

In specifying the regular expression, use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions. Explaining regular expression syntax

is beyond the scope of this document, but POSIX standards are available from the Institute of Electrical and Electronics Engineers (IEEE, <http://www.ieee.org>).

Table 611 on page 6228 specifies which character or characters are matched by some of the regular expression operators that you can use in the match statement. In the descriptions, the term term refers to either a single alphanumeric character or a set of characters enclosed in square brackets, parentheses, or braces.



**NOTE:** The match statement is not case-sensitive.

**Table 611: Regular Expression Operators for the match Statement**

| Operator                     | Matches                                                                                                                                                                                                                              |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| . (period)                   | One instance of any character except the space.                                                                                                                                                                                      |
| * (asterisk)                 | Zero or more instances of the immediately preceding term.                                                                                                                                                                            |
| + (plus sign)                | One or more instances of the immediately preceding term.                                                                                                                                                                             |
| ? (question mark)            | Zero or one instance of the immediately preceding term.                                                                                                                                                                              |
| (pipe)                       | One of the terms that appears on either side of the pipe operator.                                                                                                                                                                   |
| ! (exclamation point)        | Any string except the one specified by the expression, when the exclamation point appears at the start of the expression. Use of the exclamation point is Junos OS-specific.                                                         |
| ^ (caret)                    | Start of a line, when the caret appears outside square brackets.<br><br>One instance of any character that does not follow it within square brackets, when the caret is the first character inside square brackets.                  |
| \$ (dollar sign)             | End of a line.                                                                                                                                                                                                                       |
| [ ] (paired square brackets) | One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen ( - ) to separate the beginning and ending characters of the range. For example, [a-z0-9] matches any letter or number. |
| ( ) (paired parentheses)     | One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.                                                                                        |

#### Using Regular Expressions

Filter messages that belong to the **interactive-commands** facility, directing those that include the string **configure** to the terminal of the root user:

```
[edit system syslog]
user root {
 interactive-commands any;
 match ".*configure.*";
}
```



Messages like the following appear on the **root** user's terminal when a user issues a **configure** command to enter configuration mode:

```
timestamp router-name mgd[PID]: UI_CMDLINE_READ_LINE: User 'user', command
'configure private'
```

Filter messages that belong to the **daemon** facility and have a severity of **error** or higher, directing them to the file **/var/log/process-errors**. Omit messages generated by the SNMP process (**snmpd**), instead directing them to the file **/var/log/snmpd-errors**:

```
[edit system syslog]
file process-errors {
 daemon error;
 match "!(.*snmpd.*)";
}
file snmpd-errors {
 daemon error;
 match ".*snmpd.*";
}
```

#### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6155](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6159](#)

## Configuration Statements for Network Management

- [connection-limit on page 6230](#)
- [destination-override on page 6231](#)
- [no-remote-trace on page 6231](#)
- [protocol-version on page 6232](#)
- [rate-limit on page 6233](#)
- [ssh on page 6234](#)
- [telnet on page 6235](#)
- [tracing on page 6236](#)

## connection-limit

---

|                            |                                                                                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | connection-limit <i>limit</i> ;                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit system services finger],<br>[edit system services ftp],<br>[edit system services netconf ssh],<br>[edit system services ssh],<br>[edit system services telnet],<br>[edit system services xnm-clear-text],<br>[edit system services xnm-ssl] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                    |
| <b>Description</b>         | Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).                                                                  |
| <b>Options</b>             | <b>limit</b> —(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).<br><b>Range:</b> 1 through 250<br><b>Default:</b> 75                                                                                       |



**NOTE:** The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured connection-limit value if the system resources are limited.

---

|                           |                                                            |
|---------------------------|------------------------------------------------------------|
| <b>Required Privilege</b> | system—To view this statement in the configuration.        |
| <b>Level</b>              | system-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</a></li><li>• <a href="#">Configuring DTCP-over-SSH Service for the Flow-Tap Application</a></li><li>• <a href="#">Configuring Finger Service for Remote Access to the Router</a></li><li>• <a href="#">Configuring FTP Service for Remote Access to the Router or Switch</a></li><li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1709</a></li><li>• <a href="#">Configuring Telnet Service for Remote Access to a Router or Switch</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## destination-override

---

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | destination-override {<br>syslog host <i>ip-address</i> ;<br>}                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system tracing]                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Override the system-wide configuration of the switch at the <b>[edit system tracing]</b> hierarchy level. This statement has no effect if system tracing is not configured.                                                                                                                                |
| <b>Options</b>                  | <p><b>syslog</b>—System process log files to send to the remote tracing host.</p> <ul style="list-style-type: none"> <li>• <b>syslog</b>—System process log files to send to the remote tracing host.</li> <li>• <b>host <i>ip-address</i></b>—IP address to which to send tracing information.</li> </ul> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Tracing and Logging Operations on page 6081</a></li> <li>• <a href="#">tracing on page 312</a></li> </ul>                                                                                                                               |

## no-remote-trace

---

|                                 |                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-remote-trace                                                                                                              |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                            |
| <b>Description</b>              | Configure the switch to disable remote tracing after remote tracing has been enabled.                                        |
| <b>Default</b>                  | Remote tracing is disabled.                                                                                                  |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">tracing on page 312</a></li> </ul>                                      |

## protocol-version

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol-version <i>version</i>;</code>                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system services ssh]                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Specify the secure shell (SSH) protocol version.                                                                                                                                               |
| <b>Default</b>                  | <b>v2</b> —SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.                                                                                                         |
| <b>Options</b>                  | <b><i>version</i></b> —SSH protocol version: <b>v1</b> , <b>v2</b> , or both.                                                                                                                  |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the SSH Protocol Version on page 1710</a></li></ul>                                                                            |

## rate-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rate-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit system services finger],</code><br><code>[edit system services ftp],</code><br><code>[edit system services netconf ssh],</code><br><code>[edit system services ssh],</code><br><code>[edit system services telnet],</code><br><code>[edit system services xnm-clear-text],</code><br><code>[edit system services xnm-ssl]</code> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                |
| <b>Description</b>              | Configure the maximum number of connections attempts per protocol (either IPv6 or IPv4) on an access service.                                                                                                                                                                                                                                |
| <b>Default</b>                  | 150 connections                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><code>rate-limit <i>limit</i></code></b>—(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6).</p> <p><b>Range:</b> 1 through 250</p> <p><b>Default:</b> 150</p>                                                                                                                 |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li> </ul>                                                                                                                                                                                                    |

## ssh

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ssh {<br/>  ciphers [ <i>cipher-1 cipher-2 cipher-3 ...</i>];<br/>  client-alive-count-max <i>seconds</i>;<br/>  client-alive-interval <i>seconds</i>;<br/>  connection-limit <i>limit</i>;<br/>  hostkey-algorithm &lt;<i>algorithm</i> no-<i>algorithm</i>&gt;;<br/>  key-exchange &lt;<i>algorithm</i>&gt;;<br/>  macs &lt;<i>algorithm</i>&gt;;<br/>  max-sessions-per-connection &lt;<i>number</i>&gt;;<br/>  no-tcp-forwarding;<br/>  protocol-version [<i>v1 v2</i>];<br/>  rate-limit <i>limit</i>;<br/>  root-login (<i>allow</i>   <i>deny</i>   <i>deny-password</i>);<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system services]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>client-alive-interval</b> and <b>client-alive-max-count</b> statements introduced in Junos OS Release 12.2.</p>                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Allow SSH requests from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1709</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

---


## telnet

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | telnet {<br>connection-limit <i>limit</i> ;<br>rate-limit <i>limit</i> ;<br>}                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system services]                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Provide Telnet connections from remote systems to the local router or switch.<br><br>The remaining statements are explained separately.                                                        |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i></li></ul>                                                                    |

## tracing

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                             | <pre>tracing {<br/>    destination-override syslog host <i>ip-address</i>;<br/>}</pre>                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    | [edit system]                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                        | Configure the switch to enable remote tracing to a specified host IP address.                                                                                                          |
| <hr/>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                        |
| <div> <b>NOTE:</b> The tracing statement is not supported on the QFX3000 QFabric system.</div> <hr/>                                                                                                                                                                                                                                                                     |                                                                                                                                                                                        |
| <p>The following processes are supported:</p> <ul style="list-style-type: none"><li>• <b>chassisd</b>—Chassis-control process</li><li>• <b>eventd</b>—Event-processing process</li><li>• <b>cosd</b>—Class-of-service process</li></ul> <p>If you enabled remote tracing but wish to disable it for specific processes on the switch, use the <b>no-remote-trace</b> statement at the [edit system <i>process-name</i> traceoptions] hierarchy level.</p> |                                                                                                                                                                                        |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                            | Remote tracing is disabled by default.                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>destination-override syslog host <i>ip-address</i></b> —Overrides the global configuration for system tracing and has no effect if the <b>tracing</b> statement is not configured.  |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                           | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Understanding Tracing and Logging Operations on page 6081</a></li><li>• <a href="#">destination-override on page 256</a></li></ul> |

## Configuration Statements for Automation Scripts

---

- [allow-transients on page 6237](#)
- [apply-macro on page 6238](#)
- [checksum on page 6239](#)
- [command on page 6240](#)
- [commit on page 6241](#)



- [description on page 6242](#)
- [direct-access on page 6242](#)
- [file \(Commit Scripts\) on page 6243](#)
- [file \(Op Scripts\) on page 6244](#)
- [no-allow-url on page 6245](#)
- [op on page 6246](#)
- [optional on page 6247](#)
- [refresh \(Commit Scripts\) on page 6248](#)
- [refresh \(Op Scripts\) on page 6249](#)
- [refresh-from \(Commit Scripts\) on page 6250](#)
- [refresh-from \(Op Scripts\) on page 6251](#)
- [scripts on page 6252](#)
- [source \(Commit Scripts\) on page 6253](#)
- [source \(Op Scripts\) on page 6254](#)

## allow-transients

---

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-transients;                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit system scripts commit]                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| <b>Description</b>              | For Junos OS commit scripts, enable transient configuration changes to be committed.                                                                                                                                                                   |
| <b>Default</b>                  | Transient changes are disabled by default. If you do not include the <b>allow-transients</b> statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails. |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Generating a Persistent or Transient Change</i></li> <li>• <i>Creating a Macro to Read the Custom Syntax and Generate Related Configuration Statements</i></li> </ul>                                      |

## apply-macro

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>apply-macro <i>apply-macro-name</i> {<br/>    <i>parameter-name parameter-value</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | All hierarchy levels                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>With commit script macros, use custom syntax in your configuration.</p> <p>Macros work by locating <b>apply-macro</b> statements that you include in the candidate configuration and using the values specified in the <b>apply-macro</b> statement as parameters to a set of instructions (the macro) defined in a commit script. The commit script alters your configuration from one that contains custom syntax into a full configuration containing standard Junos OS statements.</p> <p>In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides data (or <i>hooks</i>) that are used by commit script macros.</p> <p>You can include the <b>apply-macro</b> statement at any level of the configuration hierarchy. You can include multiple <b>apply-macro</b> statements at each level of the configuration hierarchy; however, each must have a unique name.</p> |
| <b>Options</b>                  | <p><b><i>apply-macro-name</i></b>—Name of the <b>apply-macro</b> statement.</p> <p><b><i>parameter-name</i></b>—One or more parameters. Parameters can be any text you want to include in your configuration.</p> <p><b><i>parameter-value</i></b>—A value that corresponds to the parameter name. Parameter values can be any text you want to include in your configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Overview of Creating Custom Configuration Syntax with Macros</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## checksum

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>checksum (md5   sha-256   sha1) hash;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit event-options event-script file <i>filename</i> ],<br>[edit system <a href="#">scripts commit file filename</a> ],<br>[edit system <a href="#">scripts op file filename</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | For Junos OS commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event, commit, or op script, Junos OS verifies the authenticity of the script by using the configured checksum hash.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>md5 hash</b> —MD5 checksum of this script.<br><br><b>sha-256 hash</b> —SHA-256 checksum of this script.<br><br><b>sha1 hash</b> —SHA-1 checksum of this script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <b>maintenance</b> —To view this statement in the configuration.<br><b>maintenance-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i></li> <li>• <i>Executing an Op Script from a Remote Site</i></li> <li>• <a href="#">file checksum md5 on page 339</a> command in the <i>System Basics and Services Command Reference</i></li> <li>• <a href="#">file checksum sha-256 on page 341</a> command in the <i>System Basics and Services Command Reference</i></li> <li>• <a href="#">file checksum sha1 on page 340</a> command in the <i>System Basics and Services Command Reference</i></li> </ul> |

## command

---

|                                 |                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>command filename-alias;</code>                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts op file filename</a> ]                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                          |
| <b>Description</b>              | For Junos OS op scripts, configure a filename alias for the script file. This allows you to run the script by referencing either the script filename or the filename alias. |
| <b>Options</b>                  | <i>filename-alias</i> —Alias for the script file.                                                                                                                           |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling an Op Script and Defining a Script Alias</i></li></ul>                                                                  |

## commit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> commit {   allow-transients;   direct-access;   file <i>filename</i> {     checksum (md5   sha-256   sha1) <i>hash</i>;     optional;     refresh;     refresh-from <i>url</i>;     source <i>url</i>;   }   max-datasize   refresh;   refresh-from <i>url</i>;   traceoptions {     file &lt;<i>filename</i>&gt; &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | For Junos OS commit scripts, configure the commit-time scripting mechanism.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Storing and Enabling Scripts</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                            |

## description

---

|                                 |                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>descriptive-text</i>;</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts op file filename</a> ]<br>[edit system <a href="#">scripts op file filename</a> arguments <i>argument-name</i> ]                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                        |
| <b>Description</b>              | For Junos OS op scripts, provide a help-text string that appears in the command-line interface (CLI).                                                                                                     |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Help Text for Op Scripts</i></li><li>• <i>Declaring Arguments in Op Scripts</i></li><li>• <a href="#">file (Op Scripts) on page 6244</a></li></ul> |

## direct-access

---

|                                 |                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>direct-access;</code>                                                                                                 |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts commit</a> ]                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.          |
| <b>Description</b>              | Specify that commit scripts read input configurations directly from the database when inspecting these scripts for errors.  |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Executing Large Commit Scripts</i></li></ul>                                     |

## file (Commit Scripts)

|                                 |                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>file <i>filename</i> {     checksum (md5   sha-256   sha1) <i>hash</i>;     optional;     refresh;     refresh-from <i>url</i>;     source <i>url</i>; }</pre>                                                                        |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts commit</a> ]                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                              |
| <b>Description</b>              | For Junos OS commit scripts, enable a commit script that is located in the <code>/var/db/scripts/commit</code> directory.                                                                                                                  |
| <b>Options</b>                  | <p><b><i>filename</i></b>—Name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing a commit script.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Controlling Execution of Commit Scripts During Commit Operations</i></li> </ul>                                                                                                                  |

## file (Op Scripts)

---

**Syntax**    file *filename* {  
              arguments {  
                  *argument-name* {  
                    **description** *descriptive-text*;  
                  }  
              }  
              **checksum** (md5 | sha-256 | sha1) *hash*;  
              **command** *filename-alias*;  
              **description** *descriptive-text*;  
              **refresh**;  
              **refresh-from** *url*;  
              **source** *url*;  
              }

**Hierarchy Level**    [edit system **scripts op**]

**Release Information**    Statement introduced in Junos OS Release 7.6.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    For Junos OS op scripts, enable an op script that is located in the `/var/db/scripts/op` directory.

**Options**    **filename**—The name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing an op script.  
  
              The statements are explained separately.

**Required Privilege Level**    maintenance—To view this statement in the configuration.  
                                  maintenance-control—To add this statement to the configuration.

**Related Documentation**    • *Enabling an Op Script and Defining a Script Alias*



---

## no-allow-url

---

|                                 |                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-allow-url;                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system <a href="#">scripts op</a> ]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                  |
| <b>Description</b>              | For Junos OS op scripts, prohibit the remote execution of scripts. When you include this configuration statement, the <b>op url</b> operational mode command generates an error and does not permit you to execute the op script from a remote site. |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">file (Op Scripts) on page 6244</a></li><li>• <i>Executing an Op Script from a Remote Site</i></li></ul>                                                                                          |

## op

---

**Syntax**    **op** {  
    **file** *filename* {  
        arguments {  
            *argument-name* {  
                **description** *descriptive-text*;  
            }  
        }  
    }  
    **checksum** (md5 | sha-256 | sha1) *hash*;  
    **command** *filename-alias*;  
    **description** *descriptive-text*;  
    max-datasize  
    **refresh**;  
    **refresh-from** *url*;  
    **source** *url*;  
    }  
    **no-allow-url**  
    **refresh**;  
    **refresh-from** *url*;  
    traceoptions {  
        file <*filename*> <files *number*> <size *size*> <world-readable | no-world-readable>;  
        flag *flag*;  
        no-remote-trace;  
    }  
}

**Hierarchy Level**    [edit system **scripts**]

**Release Information**    Statement introduced in Junos OS Release 7.6.  
    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    For Junos OS op scripts, configure an operation scripting mechanism.


**Options**    The statements are explained separately.

**Required Privilege Level**    maintenance—To view this statement in the configuration.  
    maintenance-control—To add this statement to the configuration.

**Related Documentation**    • *Storing and Enabling Scripts*


## optional

---


|                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                              | optional;                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                     | [edit system <b>scripts commit</b> file <i>filename</i> ]                                                                                                                                             |
| <b>Release Information</b>                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                    |
| <b>Description</b>                                                                                                                                                                                                                                                         | For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the <b>file</b> statement is missing from the <b>/var/db/scripts/commit</b> directory on the device. |
| <div>  <p><b>NOTE:</b> On the QFabric system, commit scripts are stored in the <b>/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/</b> directory on the Director device.</p> </div> |                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                            | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                                                               | <ul style="list-style-type: none"> <li><i>Controlling Execution of Commit Scripts During Commit Operations</i></li> </ul>                                                                             |

## refresh (Commit Scripts)

---


|                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                                                                                                                                                                                                                                                   | refresh;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level                                                                                                                                                                                                                                                          | [edit system <a href="#">scripts commit</a> ],<br>[edit system <a href="#">scripts commit</a> file <i>filename</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Release Information                                                                                                                                                                                                                                                      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Description                                                                                                                                                                                                                                                              | <p>For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p> |
| <div> <b>NOTE:</b> On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/</code> directory on the Director device.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Required Privilege Level                                                                                                                                                                                                                                                 | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"><li>• <i>Using a Master Source Location for a Script</i></li><li>• <a href="#">refresh-from (Commit Scripts)</a> on page 6250</li><li>• <a href="#">source (Commit Scripts)</a> on page 6253</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                        |

## refresh (Op Scripts)


|                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                             | refresh;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                    | [edit system <a href="#">scripts op</a> ],<br>[edit system <a href="#">scripts op file filename</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                                                                                                                                                                                                                                                | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 11.1 on the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                                        | <p>For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at the source URL, specified in the <b>source</b> statement at the same hierarchy level.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p> |
| <div>  <p><b>NOTE:</b> On the QFabric system, op scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                           | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• <i>Using a Master Source Location for a Script</i></li> <li>• <a href="#">refresh-from (Op Scripts) on page 6251</a></li> <li>• <a href="#">source (Op Scripts) on page 6254</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                             |

## refresh-from (Commit Scripts)

---

|                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                                                                                                                                           | <code>refresh-from url;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Hierarchy Level                                                                                                                                                  | [edit system <a href="#">scripts commit</a> ],<br>[edit system <a href="#">scripts commit</a> file <i>filename</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Release Information                                                                                                                                              | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Description                                                                                                                                                      | <p>For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at a URL other than the URL specified in the <b>source</b> statement.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p> |
| <div> <b>NOTE:</b> This statement is not supported on the QFabric system.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                                                                                                                                                          | <b>url</b> —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Required Privilege Level                                                                                                                                         | <b>maintenance</b> —To view this statement in the configuration.<br><b>maintenance-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Related Documentation                                                                                                                                            | <ul style="list-style-type: none"><li>• <i>Using an Alternate Source Location for a Script</i></li><li>• <a href="#">refresh (Commit Scripts) on page 6248</a></li><li>• <a href="#">source (Commit Scripts) on page 6253</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                      |

## refresh-from (Op Scripts)

|                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                               | <code>refresh-from url;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                      | [edit system <a href="#">scripts op</a> ],<br>[edit system <a href="#">scripts op file filename</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                                                                                                                                                  | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 11.1 on the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                          | <p>For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at a URL other than the URL specified in the <b>source</b> statement.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p> |
| <div>  <b>NOTE:</b> This statement is not supported on the QFabric system.         </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                              | <b>url</b> —Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                             | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                | <ul style="list-style-type: none"> <li>• <i>Using an Alternate Source Location for a Script</i></li> <li>• <a href="#">refresh (Op Scripts) on page 6249</a></li> <li>• <a href="#">source (Op Scripts) on page 6254</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              |

## scripts

```
Syntax scripts {
 commit {
 allow-transients;
 direct-access;
 file filename {
 checksum (md5 | sha-256 | sha1) hash;
 optional;
 refresh;
 refresh-from url;
 source url;
 }
 max-datasize
 refresh;
 refresh-from url;
 traceoptions {
 file <filename> <files number> <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }
 }
 op {
 file filename {
 arguments {
 argument-name {
 description descriptive-text;
 }
 }
 checksum (md5 | sha-256 | sha1) hash;
 command filename-alias;
 description descriptive-text;
 max-datasize
 refresh;
 refresh-from url;
 source url;
 }
 no-allow-url
 refresh;
 refresh-from url;
 traceoptions {
 file <filename> <files number> <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }
 }
 }
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** For Junos OS commit or op scripts, configure scripting mechanisms.





**NOTE:** The `traceoptions` statement is not supported on QFabric systems.

|                                 |                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The statements are explained separately.                                                                                    |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Storing and Enabling Scripts</i></li> </ul>                                     |

## source (Commit Scripts)

|                            |                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>source url;</code>                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit system <a href="#">scripts commit file filename</a> ]                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                         |
| <b>Description</b>         | For Junos OS commit scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/commit</code> directory. When you include the <b>refresh</b> statement at the same hierarchy level and commit the configuration, the local copy is overwritten by the version stored at the specified URL. |




**NOTE:** On the QFabric system, commit scripts are stored in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/op/` directory on the Director device.

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>url</i> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Using a Master Source Location for a Script</i></li> <li>• <i>Overview of Updating Scripts from a Remote Source</i></li> <li>• <a href="#">refresh (Commit Scripts) on page 6248</a></li> <li>• <a href="#">refresh-from (Commit Scripts) on page 6250</a></li> </ul> |

## source (Op Scripts)

---

|                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                             | <code>source url;</code>                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                    | [edit system <a href="#">scripts op file filename</a> ]                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                                                                                                                                | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                                                                                                        | For Junos OS op scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/op</code> directory. When you include the <b>refresh</b> statement at the same hierarchy level, the local copy is overwritten by the version stored at the specified URL. |
| <hr/>                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                       |
| <div> <b>NOTE:</b> On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device.</div> <hr/> |                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                                                                                                                                                                                                                                            | <b>url</b> —Master source file for an op script specified as an HTTP URL, FTP URL, or scp-style remote file specification.                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                           | <b>maintenance</b> —To view this statement in the configuration.<br><b>maintenance-control</b> —To add this statement to the configuration.                                                                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Using a Master Source Location for a Script</a></li><li>• <a href="#">refresh (Op Scripts) on page 6249</a></li><li>• <a href="#">refresh-from (Op Scripts) on page 6251</a></li></ul>                                                            |

## Configuration Statements for Fabric OAM

---

- [ethernet-frame-size on page 6255](#)
- [fabric \(OAM\) on page 6256](#)
- [fabric-maintenance-associations on page 6258](#)
- [fabric-maintenance-end-points on page 6259](#)
- [flow-specs on page 6260](#)
- [unicast-ethernet-ipv4 on page 6262](#)
- [multicast-ipv4 on page 6264](#)
- [multicast-vlan-flood on page 6265](#)
- [unicast-ethernet on page 6266](#)

---

## ethernet-frame-size

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ethernet-frame-size <i>ethernet-frame-size</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols oam fabric flow-specs <i>flow-specification-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the Ethernet frame size in the flow specification parameters for internal fabric monitoring.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b><i>ethernet-frame-size</i></b> —Integer defining the size (in bytes) of the Ethernet frame.<br><b>Range:</b> 256 to 9116 bytes                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li><li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li><li>• <a href="#">Configuring Flow Specifications on page 6185</a></li><li>• <a href="#">fabric (OAM) on page 6256</a></li><li>• <a href="#">ping fabric multicast-flow on page 6419</a></li><li>• <a href="#">ping fabric unicast-flow on page 6421</a></li><li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li></ul> |

## fabric (OAM)

---

```
Syntax fabric {
 fabric-maintenance-associations {
 fma-name {
 description string;
 fabric-maintenance-end-points {
 fmep-id {
 description string;
 fmep-interface interface-name;
 fmep-name name;
 }
 }
 vlan-name vlan-name;
 }
 }
 flow-specs {
 flow-specification-name {
 ethernet-frame-size size;
 multicast-ipv4 {
 dest-ip-multicast-group ipv4-mcast-address;
 source-ip ipv4-address;
 }
 multicast-vlan-flood;
 unicast-ethernet {
 destination-mac destination-mac-address destination-mac-mask
 destination-mac-mask;
 ethertype ethertype;
 source-mac mac-address source-mac-mask source-mac-mask;
 }
 unicast-ethernet-ipv4 {
 destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
 destination-l4-port destination-l4-port-number;
 destination-mac destination-mac-address;
 ip-proto protocol;
 source-ip source-ip-address source-ip-mask source-ip-mask;
 source-l4-port source-l4-port-number;
 }
 }
 }
 faboam-trace-options {
 file filename{
 files number;
 no-stamp;
 no-world-readable;
 replace;
 size size;
 world-readable;
 }
 flag {
 all;
 debug-all;
 ffping;
 generic;
 }
 }
 }
```

```

 netio;
 packets;
 trace-route;
 }
}

```

**Hierarchy Level** [edit protocols oam]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Configure the fabric maintenance association (FMA) and flow specification parameters for internal fabric monitoring.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring a Fabric Maintenance Association on page 6184](#)
- [Configuring Flow Specifications on page 6185](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

## fabric-maintenance-associations

```
Syntax fabric-maintenance-associations {
 fma-name {
 description string;
 fabric-maintenance-end-points {
 fmep-id {
 description string;
 fmep-interface interface-name;
 fmep-name name;
 }
 }
 vlan-name vlan-name;
 }
}
```

**Hierarchy Level** [edit protocols oam fabric]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Configure the fabric maintenance association (FMA) parameters for internal fabric monitoring. The FMA defines a set of QFabric system interfaces (fabric maintenance endpoints [FMEPs]) that are members of a given VLAN which are used in internal fabric monitoring operations. The FMA associates the set of FMEPs with a VLAN. Configuration of the VLAN is mandatory.

The remaining statements under this hierarchy level are explained separately.

**Options** **description *string***—Description of the FMA.  
**Syntax:** 1 to 128 alphanumeric characters

**fma-name**—Name of the FMA.  
**Syntax:** 1 to 20 alphanumeric characters

**vlan-name *vlan-name***—Name of the VLAN for which the FMA is defined.



**NOTE:** Configuring the VLAN option is mandatory.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring a Fabric Maintenance Association on page 6184](#)
- [Configuring Flow Specifications on page 6185](#)
- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)

- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

## fabric-maintenance-end-points

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> fabric-maintenance-end-points {     fmep-id {         description <i>string</i>;         fmep-interface <i>interface-name</i>;         fmep-name <i>name</i>;     } } </pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols oam fabric fabric-maintenance-associations <i>fma-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the fabric maintenance association endpoint (FMEP) parameters within a fabric maintenance association (FMA). The FMA associates the FMEP with a VLAN for purposes of internal fabric monitoring.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>description <i>string</i></b>—Description of the FMEP.<br/> <b>Syntax:</b> 1 to 128 alphanumeric characters</p> <p><b><i>fmep-id</i></b>—32-bit integer that is used to identify an FMEP.<br/> <b>Range:</b> 1 to 64,000</p> <p><b><i>fmep-name name</i></b>—Name of the fabric maintenance association endpoint.<br/> <b>Syntax:</b> 1 to 20 alphanumeric characters</p> <p><b><i>fmep-interface interface-name</i></b>—Interface on a QFabric system Node device.<br/> <b>Syntax:</b> <i>node-device-name:interface-type-fpc/pic/slot.unit</i></p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> <li>• <a href="#">ping fabric multicast-flow on page 6419</a></li> <li>• <a href="#">ping fabric unicast-flow on page 6421</a></li> <li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li> </ul>       |

## flow-specs

```
Syntax flow-specs {
 flow-specification-name {
 ethernet-frame-size ethernet-frame-size;
 multicast-ipv4 {
 dest-ip-multicast-group ipv4-mcast-address;
 source-ip ipv4-address;
 }
 multicast-vlan-flood;
 unicast-ethernet {
 destination-mac destination-mac-address destination-mac-mask
 destination-mac-mask;
 ethertype ethertype;
 source-mac mac-address source-mac-mask source-mac-mask;
 }
 unicast-ethernet-ipv4 {
 destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
 destination-l4-port destination-l4-port-number;
 destination-mac destination-mac-address;
 ip-proto protocol;
 source-ip source-ip-address source-ip-mask source-ip-mask;
 source-l4-port source-l4-port-number;
 }
 }
 }
```

**Hierarchy Level** [edit protocols oam fabric]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Configure the fabric flow specification parameters for internal fabric monitoring. The following conditions apply:

- You must specify a protocol type for each flow specification: **unicast-ethernet**, **unicast-ethernet-ipv4**, **multicast-ipv4**, or **multicast-vlan-flood**.
- You can specify only one type for each flow specification.
- All other parameters are optional.

The remaining statements under this hierarchy level are explained separately.

**Options** *flow-specification-name*—Name of the flow specification.

**Syntax:** 1 to 20 alphanumeric characters

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring a Fabric Maintenance Association on page 6184](#)
- [Configuring Flow Specifications on page 6185](#)



- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

## unicast-ethernet-ipv4

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>unicast-ethernet-ipv4 {<br/>    source-ip <i>source-ip-address</i> source-ip-mask <i>source-ip-mask</i>;<br/>    destination-ip <i>destination-ip-address</i> destination-ip-mask <i>destination-ip-mask</i>;<br/>    ip-proto <i>protocol</i>;<br/>    source-l4-port <i>source-l4-port-number</i>;<br/>    destination-l4-port <i>destination-l4-port-number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit protocols oam fabric flow-specs <i>flow-specification-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | Configure unicast parameters for the Ethernet IPv4 protocol type in the internal fabric monitoring flow specifications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>destination-ip <i>destination-ip-address</i></b>—Unicast IPv4 address of the destination fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).<br/><b>Syntax:</b> Dotted decimal notation without the prefix</p> <p><b>destination-l4-port <i>destination-l4-port-number</i></b>—L4 TCP or UDP port for the destination fabric maintenance endpoint (FMEP).<br/><b>Values:</b> 2-byte decimal or hexadecimal notation</p> <p><b>ip-proto <i>protocol</i></b>—IPv4 protocol type.<br/><b>Values:</b> 1-byte value in decimal or hexadecimal notation</p> <p><b>destination-ip-mask <i>destination-ip-mask</i></b>—IPv4 address range mask for the destination FMEP that is configured in the flow specification.<br/>The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.<br/><b>Syntax:</b> Dotted decimal notation without the prefix<br/><b>Values:</b> Maximum of 6 contiguous masked bits, or a mask range of 64</p> <p><b>source-ip <i>source-ip-address</i></b>—Unicast IPv4 address of the source fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).<br/><b>Syntax:</b> Dotted decimal notation without the prefix</p> <p><b>source-l4-port <i>source-l4-port-number</i></b>—TCP or UDP port for the source fabric maintenance endpoint (FMEP).<br/><b>Values:</b> 2-byte decimal or hexadecimal notation</p> <p><b>source-ip-mask <i>source-ip-mask</i></b>—IPv4 address range mask for the source FMEP that is configured in the flow specification.</p> |

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.

**Syntax:** Dotted decimal notation without the prefix

**Values:** Maximum of 6 contiguous masked bits, or a mask range of 64

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li><li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li><li>• <a href="#">Configuring Flow Specifications on page 6185</a></li><li>• <a href="#">fabric (OAM) on page 6256</a></li><li>• <a href="#">ping fabric multicast-flow on page 6419</a></li><li>• <a href="#">ping fabric unicast-flow on page 6421</a></li><li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## multicast-ipv4

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multicast-ipv4 {<br/>    source-ip <i>source-ip-address</i>;<br/>    dest-ip-multicast-group <i>ipv4-mcast-address</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols oam fabric flow-specs <i>flow-specification-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure multicast parameters for the IPv4 Ethernet protocol type in the internal fabric monitoring flow specifications.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>dest-ip-multicast-group <i>ipv4-mcast-address</i></b>—IPv4 multicast group address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification. Configuration of this parameter is mandatory.</p> <p><b>Syntax:</b> Dotted decimal notation without the prefix</p> <p><b>source-ip <i>source-ip-address</i></b>—(Optional) IPv4 multicast address of the source fabric maintenance endpoint (FMEP) that is configured in the flow specification.</p> <p><b>Syntax:</b> Dotted decimal notation without the prefix</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li><li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li><li>• <a href="#">Configuring Flow Specifications on page 6185</a></li><li>• <a href="#">fabric (OAM) on page 6256</a></li><li>• <a href="#">ping fabric multicast-flow on page 6419</a></li><li>• <a href="#">ping fabric unicast-flow on page 6421</a></li><li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li></ul>                       |

---

## multicast-vlan-flood

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multicast-vlan-flood                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols oam fabric flow-specs <i>flow-specification-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure parameters of the multicast VLAN flood protocol type in the internal fabric monitoring flow specification.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | There are no configuration options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li><li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li><li>• <a href="#">Configuring Flow Specifications on page 6185</a></li><li>• <a href="#">fabric (OAM) on page 6256</a></li><li>• <a href="#">ping fabric multicast-flow on page 6419</a></li><li>• <a href="#">ping fabric unicast-flow on page 6421</a></li><li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li></ul> |

## unicast-ethernet

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>unicast-ethernet {     destination-mac <i>destination-mac-address</i> destination-mac-mask <i>destination-mac-mask</i>;     ethertype <i>ethertype</i>;     source-mac <i>source-mac-address</i> source-mac-mask <i>source-mac-mask</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols oam fabric flow-specs <i>flow-specification-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure unicast Ethernet parameters for a flow specification used in internal fabric monitoring.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>destination-mac</b> <i>destination-mac-address</i>—(Optional) MAC address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification.</p> <p><b>destination-mac-mask</b> <i>destination-mac-mask</i>—(Optional) MAC address range mask for the destination FMEP that is configured in the flow specification. You cannot configure the <b>destination-mac-mask</b> parameter unless you also configure the <b>destination-mac</b> parameter.</p> <p>The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.</p> <p><b>Values:</b> Maximum of 6 contiguous masked bits, or a mask range of 64</p> <p><b>ethertype</b> <i>ethertype</i>—Valid Ethernet type other than the IPv4 Ethertype.</p> <p><b>Values:</b> 2 bytes in decimal or hexadecimal notation</p> <p><b>source-mac</b> <i>source-mac-address</i>—(Optional) MAC address of the source FMEP or interface that is configured in the flow specification.</p> <p><b>source-mac-mask</b> <i>source-mac-mask</i>—(Optional) MAC address range mask for the source FMEP that is configured in the flow specification. You cannot configure the <b>source-mac-mask</b> parameter unless you also configure the <b>source-mac</b> parameter.</p> <p>The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.</p> <p><b>Values:</b> Maximum of 6 contiguous masked bits, or a mask range of 64</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)
- [traceroute fabric unicast-flow on page 6467](#)

## Configuration Statements for sFlow Technology

---

- [agent-id on page 6267](#)
- [collector on page 6268](#)
- [interfaces \(sFlow\) on page 6268](#)
- [polling-interval on page 6269](#)
- [sample-rate on page 6270](#)
- [sflow on page 6271](#)
- [source-ip on page 6272](#)
- [traceoptions \(sFlow Technology\) on page 6273](#)
- [udp-port on page 6274](#)

### agent-id

---

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>agent-id <i>ip-address</i>;</code>                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols sflow]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                      |
| <b>Description</b>              | Configure the IP address of the sFlow agent. If you do not configure the sFlow agent ID, the IP address for the agent is dynamically created using the IP address of an interface configured on the QFX Series device. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> <li>• <a href="#">sflow on page 6271</a></li> </ul>                                                            |

## collector

---

|                                 |                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>collector <i>ip-address</i> {<br/>    <code>udp-port</code> <i>port-number</i>;<br/>}</code>                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <code>protocols sflow</code> ]                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the configured collector for analysis. You can configure up to four collectors on the device. You specify the IP address for each collector you configure.</p> <p>The remaining statement is explained separately.</p> |
| <b>Options</b>                  | <i>ip-address</i> —IP address of the collector.                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6189</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li></ul>                                                                                                                        |

## interfaces (sFlow)

---

|                                 |                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interfaces <i>interface-name</i> {<br/>    <code>polling-interval</code> <i>seconds</i>;<br/>    <code>sample-rate</code> <i>number</i>;<br/>}</code>                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <code>protocols sflow</code> ]                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Configure sFlow network traffic monitoring on the specified interface on the device. You can configure sFlow parameters (polling interval, sample rate) with different values on different interfaces.</p> <p>The remaining statements are explained separately.</p> |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface on which to configure sFlow parameters.                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6189</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li></ul>                                                           |



## polling-interval

---

|                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>polling-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit <a href="#">protocols sflow</a>],</code><br><code>[edit <a href="#">protocols sflow interfaces</a> <i>interface-name</i>]</code>                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the rate (in seconds) at which successive samples of interface statistics (counters) are taken.                                                                                                                                                                                            |
| <b>Default</b>                  | If no polling interval is configured for a particular interface, the device uses the global polling interval configured at the <code>[edit <a href="#">protocols sflow</a>]</code> hierarchy level. If no global interval is configured, the device uses the default polling interval of 20 seconds. |
| <b>Options</b>                  | <b><i>seconds</i></b> —Number of seconds between successive samples of interface statistics.<br>Specifying a value of <b>0</b> (zero) disables the polling.<br><b>Range:</b> 0 through 3600 seconds                                                                                                  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> </ul>                                                                                     |

## sample-rate

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>sample-rate <i>number</i>;</code>                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols sflow</a> ],<br>[edit <a href="#">protocols sflow interfaces</a> <i>interface-name</i> ]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the denominator ( <i>number</i> ) of the ratio that is the sample rate in sFlow traffic monitoring. For example, to configure a sample rate of 1 in 1000 packets, you specify a <i>number</i> of 1000.                                                              |
| <b>Default</b>                  | If no sample rate is configured for a particular interface, the device uses the global sample rate configured at the <a href="#">[edit protocols sflow]</a> hierarchy level. If no global rate is configured, the device uses the default sample rate of 1 in 2000 packets. |
| <b>Options</b>                  | <i>number</i> —Denominator of the ratio representing the sample rate (one packet out of <i>number</i> ).<br><b>Range:</b> 1 through 16,777,215                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6189</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li></ul>                                                               |

## sflow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> sflow {   agent-id <i>ip-address</i>;   collector <i>ip-address</i> {     udp-port <i>port-number</i>;   }   interfaces <i>interface-name</i> {     polling-interval <i>number</i>;     sample-rate {       egress <i>number</i>;       ingress <i>number</i>;     }   }   polling-interval <i>number</i>;   sample-rate {     egress <i>number</i>;     ingress <i>number</i>;   }   source-ip <i>ip-address</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i>;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure sFlow technology to monitor traffic continuously on specified interfaces simultaneously. sFlow data can be used to characterize network activity.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | The sFlow protocol is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                              |

## source-ip

---

|                                 |                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-ip <i>ip-address</i> ;                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols sflow]                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                            |
| <b>Description</b>              | Configure the source IP address to be used for sFlow datagrams. If you do not configure a source IP address, it is dynamically created based on the IP address of an Ethernet interface configured on the QFX Series device. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6189</a></li><li>• <a href="#">sflow on page 6271</a></li></ul>                                                                     |

## traceoptions (sFlow Technology)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">sflow</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | Define tracing operations for sFlow technology.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Output files are located in the <code>/var/log/</code> directory.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>. Incoming trace file data is logged in the now empty <b>trace-file</b>. When <b>trace-file</b> again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify the maximum number of files, you must also specify the maximum file size using the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>all—Trace all sFlow monitoring events.</li> <li>client-server—Trace sFlow monitoring client-server events.</li> <li>configuration—Trace sFlow monitoring configuration events.</li> <li>interface—Trace sFlow monitoring interface events.</li> <li>rtsock—Trace routing socket code events.</li> </ul> <p><b>no-stamp</b>—(Optional) Do not place timestamp information at the beginning of each line in the trace file.</p> <p><b>no-world-readable</b>—(Optional) Prevent any user from reading the trace file.</p> <p><b>replace</b>—(Optional) Replace an existing trace file if there is one.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches its maximum size, it</p> |

is renamed **trace-file.0**. Incoming trace file data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size of 4 GB

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the trace file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of sFlow Technology on page 6105](#)

---

## udp-port

---

**Syntax** `udp-port port-number;`

**Hierarchy Level** [edit protocols [sflow collector](#)]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the UDP port for a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the collector for analysis.

**Default** Port 6343

**Options** *port-number*—UDP port number for this collector.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring sFlow Technology on page 6189](#)
- [Example: Monitoring Network Traffic Using sFlow Technology on page 6165](#)

---

## Configuration Statements for SNMP

---

- [access \(SNMP\) on page 6278](#)
- [address \(SNMP\) on page 6278](#)
- [address-mask on page 6279](#)
- [agent-address on page 6279](#)
- [alarm \(SNMP RMON\) on page 6280](#)

- [authentication-md5 on page 6281](#)
- [authentication-none on page 6282](#)
- [authentication-password on page 6283](#)
- [authentication-sha on page 6284](#)
- [authorization on page 6285](#)
- [bucket-size on page 6286](#)
- [categories on page 6286](#)
- [client-list on page 6287](#)
- [client-list-name on page 6287](#)
- [clients on page 6288](#)
- [commit-delay on page 6288](#)
- [community \(SNMP\) on page 6289](#)
- [community \(RMON\) on page 6290](#)
- [community-name \(SNMP\) on page 6291](#)
- [contact on page 6292](#)
- [description \(SNMP\) on page 6292](#)
- [description \(RMON\) on page 6293](#)
- [destination-port \(SNMP\) on page 6293](#)
- [engine-id on page 6294](#)
- [event on page 6295](#)
- [falling-event-index \(RMON\) on page 6296](#)
- [falling-threshold \(Health Monitor\) on page 6297](#)
- [falling-threshold \(RMON\) on page 6298](#)
- [falling-threshold-interval on page 6299](#)
- [filter-duplicates on page 6299](#)
- [filter-interfaces on page 6300](#)
- [group \(Associating a Security Name\) on page 6300](#)
- [group \(Configuring Access Privileges\) on page 6301](#)
- [health-monitor on page 6302](#)
- [history on page 6303](#)
- [interface \(SNMP\) on page 6304](#)
- [interface \(RMON\) on page 6305](#)
- [interval \(Health Monitor\) on page 6305](#)
- [interval \(RMON\) on page 6306](#)
- [local-engine on page 6307](#)
- [location on page 6308](#)
- [logical-system \(SNMP\) on page 6309](#)

- [message-processing-model](#) on page 6310
- [name](#) on page 6310
- [nonvolatile](#) on page 6311
- [notify](#) on page 6311
- [notify-filter \(Applying to the Management Target\)](#) on page 6312
- [notify-filter \(Configuring the Profile Name\)](#) on page 6312
- [notify-view](#) on page 6313
- [oid](#) on page 6313
- [oid \(SNMPv3\)](#) on page 6314
- [owner](#) on page 6314
- [parameters](#) on page 6315
- [port \(SNMP\)](#) on page 6315
- [privacy-3des](#) on page 6316
- [privacy-aes128](#) on page 6317
- [privacy-des](#) on page 6318
- [privacy-none](#) on page 6318
- [privacy-password](#) on page 6319
- [read-view](#) on page 6320
- [remote-engine](#) on page 6321
- [request-type](#) on page 6322
- [retry-count \(SNMPv3\)](#) on page 6323
- [rising-event-index](#) on page 6324
- [rising-threshold \(Health Monitor\)](#) on page 6325
- [rising-threshold \(RMON\)](#) on page 6326
- [rmon](#) on page 6327
- [routing-instance \(SNMP\)](#) on page 6328
- [sample-type](#) on page 6329
- [security-level \(Defining Access Privileges\)](#) on page 6330
- [security-level \(Generating SNMP Notifications\)](#) on page 6331
- [security-model \(Access Privileges\)](#) on page 6332
- [security-model \(Group\)](#) on page 6333
- [security-model \(SNMP Notifications\)](#) on page 6334
- [security-name \(Community String\)](#) on page 6335
- [security-name \(Security Group\)](#) on page 6336
- [security-name \(SNMP Notifications\)](#) on page 6337
- [security-to-group](#) on page 6338
- [snmp](#) on page 6339



- [snmp-community](#) on page 6343
- [source-address \(SNMP\)](#) on page 6343
- [startup-alarm](#) on page 6344
- [syslog-subtag](#) on page 6345
- [tag \(Configuring Notification Targets\)](#) on page 6345
- [tag \(Configuring the SNMP Community\)](#) on page 6346
- [tag-list](#) on page 6346
- [target-address](#) on page 6347
- [target-parameters](#) on page 6348
- [targets](#) on page 6349
- [timeout](#) on page 6349
- [traceoptions \(SNMP\)](#) on page 6350
- [trap-group](#) on page 6352
- [trap-options](#) on page 6353
- [type \(RMON Notification\)](#) on page 6354
- [type \(SNMPv3\)](#) on page 6355
- [user](#) on page 6355
- [usm](#) on page 6356
- [v3](#) on page 6358
- [vacm](#) on page 6360
- [variable](#) on page 6361
- [version](#) on page 6362
- [view \(Configuring a MIB View\)](#) on page 6363
- [view \(Associating MIB View with a Community\)](#) on page 6364
- [write-view](#) on page 6364

## access (SNMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>access {<br/>  group group-name {<br/>    (default-context-prefix   context-prefix context-prefix) {<br/>      security-model (any   usm   v1   v2c) {<br/>        security-level (authentication   none   privacy) {<br/>          notify-view view-name;<br/>          read-view view-name;<br/>          write-view view-name;<br/>        }<br/>      }<br/>    }<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm]                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Set SNMP access limits.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                  |

## address (SNMP)

---

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>address address;</pre>                                                                                          |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address target-address-name]                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                    |
| <b>Description</b>              | Specify the SNMP target address for receiving traps or informs.                                                      |
| <b>Options</b>                  | <b>address</b> —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.        |

## address-mask

---

|                                 |                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address-mask <i>address-mask</i>;</code>                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-address <i>target-address-name</i>]</code>                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 on the QFX Series. |
| <b>Description</b>              | Define and verify the source addresses for a group of target addresses for SNMP traps and informs.                                                                                            |
| <b>Options</b>                  | <b><i>address-mask</i></b> —Define a range of addresses.                                                                                                                                      |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Address Mask</i></li> </ul>                                                                                                       |

## agent-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>agent-address outgoing-interface;</code>                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-options]</code>                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is <b>outgoing-interface</b> , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.                                                                                        |
| <b>Options</b>                  | <b>outgoing-interface</b> —Value of the agent address of all SNMPv1 traps generated by this router or switch. The <b>outgoing-interface</b> option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.<br><b>Default:</b> Disabled (the agent address is not specified in SNMPv1 traps). |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Agent Address for SNMP Traps</i></li> </ul>                                                                                                                                                                                                                                |

## alarm (SNMP RMON)

---

**Syntax**    alarm *index* {  
              description *description*;  
              falling-event-index *index*;  
              falling-threshold *integer*;  
              falling-threshold-interval *seconds*;  
              interval *seconds*;  
              request-type (get-next-request | get-request | walk-request);  
              rising-event-index *index*;  
              rising-threshold *integer*;  
              sample-type (absolute-value | delta-value);  
              startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);  
              syslog-subtag *syslog-subtag*;  
              variable *oid-variable*;  
              }

**Hierarchy Level**    [edit snmp rmon]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure RMON alarm entries.

**Options**    *index*—Identifies this alarm entry as an integer.  
  
              The remaining statements are explained separately.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                  snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring an Alarm Entry and Its Attributes](#)
- [event \(SNMP\)](#)
- [Configuring RMON Alarms and Events on page 6198](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6121](#)
- [Monitoring RMON MIB Tables on page 6389](#)
- [Understanding RMON on page 6119](#)
- [Junos OS Network Management Configuration Guide](#)

## authentication-md5

|                            |                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | authentication-md5 {<br>authentication-password authentication-password;<br>}                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | Configure MD5 as the authentication type for the SNMPv3 user.                                                                                                                                  |



**NOTE:** You can only configure one authentication type for each SNMPv3 user.

The remaining statement is explained separately.

|                                 |                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring MD5 Authentication</li> </ul>                              |

## authentication-none

---

|                            |                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | authentication-none;                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | Configure that there should be no authentication for the SNMPv3 user.                                                                                                                          |



**NOTE:** You can configure only one authentication type for each SNMPv3 user.

---

|                              |                                                                                        |
|------------------------------|----------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | snmp—To view this statement in the configuration.                                      |
| <b>Level</b>                 | snmp-control—To add this statement to the configuration.                               |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring No Authentication</i></li></ul> |

## authentication-password

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-password <i>authentication-password</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit snmp v3 usm local-engine user <i>username</i> authentication-md5],<br>[edit snmp v3 usm local-engine user <i>username</i> authentication-sha],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-md5],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-sha]                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the password for user authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>authentication-password</i></b>—Password that a user enters. The password is then converted into a key that is used for authentication.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> <li>• The password must be at least eight characters long.</li> <li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li> </ul> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring MD5 Authentication</i></li> <li>• <i>Configuring SHA Authentication</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                  |

## authentication-sha

---

|                            |                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>authentication-sha {<br/>    <b>authentication-password</b> <i>authentication-password</i>;<br/>}</code>                                                                                 |
| <b>Hierarchy Level</b>     | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | Configure the secure hash algorithm (SHA) as the authentication type for the SNMPv3 user.                                                                                                      |



**NOTE:** You can configure only one authentication type for each SNMPv3 user.

---


The remaining statement is explained separately.

|                                 |                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SHA Authentication</i></li></ul>                       |



## authorization

---

|                                 |                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authorization <i>authorization</i>;</code>                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Set the access authorization for SNMP <b>Get</b> , <b>GetBulk</b> , <b>GetNext</b> , and <b>Set</b> requests.                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>authorization</i></b>—Access authorization level:</p> <ul style="list-style-type: none"> <li>• <b>read-only</b>—Enable <b>Get</b>, <b>GetNext</b>, and <b>GetBulk</b> requests.</li> <li>• <b>read-write</b>—Enable all requests, including <b>Set</b> requests. You must configure a view to enable <b>Set</b> requests.</li> </ul> |
|                                 | <div>  <p><b>NOTE:</b> The <b>read-write</b> option is not supported on the QFX3000 QFabric system.</p> </div>                                                                                                                                               |
|                                 | <b>Default:</b> read-only                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the SNMP Community String on page 6194</a></li> </ul>                                                                                                                                                                                                                        |

## bucket-size

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bucket-size <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon history <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | 50                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <i>number</i> —Number of discrete samples of Ethernet statistics requested.                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## categories

---

|                                 |                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>categories {<br/>    <i>category</i>;<br/>}</code>                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-group <i>group-name</i>]</code>                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                      |
| <b>Description</b>              | Define the types of traps that are sent to the targets of the named trap group.                                                                                                        |
| <b>Default</b>                  | If you omit the <b>categories</b> statement, all trap types are included in trap notifications.                                                                                        |
| <b>Options</b>                  | <i>category</i> —Name of a trap type: <b>authentication</b> , <b>chassis</b> , <b>configuration</b> , <b>link</b> , <b>remote-operations</b> , <b>rmon-alarm</b> , or <b>startup</b> . |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6195</a></li></ul>                                                                            |

## client-list

---

|                                 |                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>client-list <i>client-list-name</i> {<br/>    <i>ip-addresses</i>;<br/>}</code>                                                               |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                   |
| <b>Description</b>              | Define a list of SNMP clients.                                                                                                                      |
| <b>Options</b>                  | <p><i>client-list-name</i>—Name of the client list.</p> <p><i>ip-addresses</i>—IP addresses of the SNMP clients to be added to the client list,</p> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6196</a></li> </ul>                     |

## client-list-name

---

|                                 |                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>client-list-name <i>client-list-name</i>;</code>                                                                          |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                               |
| <b>Description</b>              | Add a client list or prefix list to an SNMP community.                                                                          |
| <b>Options</b>                  | <i>client-list-name</i> —Name of the client list or prefix list.                                                                |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6196</a></li> </ul> |

## clients

---


|                                 |                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clients {<br/>    address &lt;restrict&gt;;<br/>}</pre>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.                                                                                                                                                                                                          |
| <b>Default</b>                  | If you omit the <b>clients</b> statement, all SNMP clients using this community string are authorized to access the switch.                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>address</b>—Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <b>address</b> options.</p> <p><b>restrict</b>—(Optional) Do not allow the specified SNMP client to access the switch.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the SNMP Community String</i></li></ul>                                                                                                                                                                                                                  |

## commit-delay

---

|                                 |                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>commit-delay seconds;</pre>                                                                                                             |
| <b>Hierarchy Level</b>          | [edit snmp nonvolatile]                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                            |
| <b>Description</b>              | Configure the timer for the SNMP <b>Set</b> reply and start of the commit.                                                                   |
| <b>Options</b>                  | <p><b>seconds</b>—Delay between an affirmative SNMP <b>Set</b> reply and start of the commit operation.</p> <p><b>Default:</b> 5 seconds</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li></ul>                                                  |

## community (SNMP)


|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>community <i>community-name</i> {   authorization <i>authorization</i>;   client-list-name <i>client-list-name</i>;   clients {     address restrict;   }   view <i>view-name</i>; }</pre>                                                                             |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.</p> |
|                                 | <p> <b>NOTE:</b> The <b>authorization read-write</b> option is not supported on the QFX3000 QFabric system.</p>                                                                           |
|                                 | <p>The SNMP client application specifies an SNMP community name in <b>Get</b>, <b>GetBulk</b>, <b>GetNext</b>, and <b>Set</b> SNMP requests.</p>                                                                                                                            |
| <b>Default</b>                  | If you omit the <b>community</b> statement, all SNMP requests are denied.                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>community-name</i></b>—Community string. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>                                                                                           |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring the SNMP Community String on page 6194</a></li> </ul>                                                                                                                                                        |

## community (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>community <i>community-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp rmon event <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the SNMP trap group that is used when generating a trap (if the <b>eventType</b> object is configured to send traps). If that trap group has the <b>rmon-alarm</b> trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of <b>eventCommunity</b>). If nothing is configured, traps are sent to each group that has the <b>rmon-alarm</b> category configured.</p> <p>The event community is not the same as an SNMP community.</p> |
| <b>Options</b>                  | <b><i>community-name</i></b> —Name of the trap group that is used when generating a trap if the event is configured to send traps.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>                                                                                                                                                            |

## community-name (SNMP)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <code>community-name <i>community-name</i>;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <code>[edit snmp v3 snmp-community <i>community-index</i>]</code>                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11. for the QFX Series.                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Define an SNMP community to authorize SNMPv1 or SNMPv2c clients in an SNMPv3 system. When you configure a community in SNMPv3, you can also specify a security name. The access privileges associated with the security name determine which MIB objects are available and which operations (read, write, or notify) are allowed on those objects. |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b><i>community-name</i></b> —Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose the name in quotation marks (" ").                                                                                                                                 |
| <div>  <p><b>NOTE:</b> Community names must be unique. You cannot configure the same community name at the <code>[edit snmp community]</code> and <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy levels.</p> <p>The community name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level is encrypted and not displayed in the command-line interface (CLI).</p> </div> |                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"> <li><i>Configuring the SNMPv3 Community</i></li> </ul>                                                                                                                                                                                                                                                          |

## contact

---

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>contact <i>contact</i>;</code>                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                    |
| <b>Description</b>              | Define the value of the MIB II <b>sysContact</b> object, which is the contact person for the managed system.         |
| <b>Options</b>                  | <b>contact</b> —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").        |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Contact on a Device Running Junos OS</i></li></ul> |

## description (SNMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>description</i>;</code>                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Define the value of the MIB II <b>sysDescription</b> object, which is the description of the system being managed.                                                                             |
| <b>Options</b>                  | <b>description</b> —System description. If the name includes spaces, enclose it in quotation marks (" ").                                                                                      |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Description on a Device Running Junos OS</i></li></ul>                                                                       |



## description (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>description</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ],<br>[edit snmp rmon event <i>index</i> ]                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Text description of alarm or event.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b><i>description</i></b> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |


## destination-port (SNMP)

---

|                                 |                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>destination-port <i>port-number</i>;</code>                                                             |
| <b>Hierarchy Level</b>          | [edit snmp trap-group]                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                             |
| <b>Description</b>              | Assign a trap port number other than the default.                                                             |
| <b>Default</b>                  | If you omit this statement, the default port is 162.                                                          |
| <b>Options</b>                  | <b><i>port-number</i></b> —SNMP trap port number.                                                             |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Trap Groups on page 6195</a></li> </ul> |

## engine-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | engine-id {<br>(local <i>engine-id-suffix</i>   use-default-ip-address   use-mac-address);<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Define a unique identifier for an SNMPv3 engine by configuring the suffix of the engine ID. The engine ID is used for identification only and not for addressing. There are two parts of an engine ID: the prefix and the suffix. The prefix is formatted according to the specifications defined in RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> and cannot be configured. The suffix is configured here.</p> <div><b>NOTE:</b> SNMPv3 authentication and encryption keys are generated based on the associated user passwords and the engine ID. If you configure or change the engine ID, you must commit the user passwords and new engine ID before you configure SNMPv3 users, or the authentication will fail.</div> <p>By default, the engine ID suffix is configured with the MAC address of the management interface (the <i>use-mac-address</i> option) on the QFX Series. You can override this configuration by using the local <i>engine-id-suffix</i> or <i>use-default-ip-address</i> option.</p> |
| <b>Default</b>                  | use-mac-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>local engine-id-suffix</i>—The engine ID suffix is set based on the data entered.</p> <p><i>use-default-ip-address</i>—The engine ID suffix is generated from the default IP address.</p> <p><i>use-mac-address</i>—The engine ID suffix is generated from the MAC address of the management interface on the switch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## event

---

|                                 |                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>event <i>index</i> {     community <i>community-name</i>;     description <i>description</i>;     type (RMON Notification) <i>type</i>; }</pre>                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit snmp rmon]                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure RMON event entries.                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><i>index</i>—Identifier for a specific event entry.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## falling-event-index (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-event-index <i>index</i>;</code>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon alarm <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set the index number of the event entry that is used when a falling threshold is crossed. You specify the falling-event index when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.                                                                                                                                                                                                    |
| <b>Options</b>                  | <i>index</i> —Index of the event entry that is used when a falling threshold is crossed.<br><b>Range:</b> 0 through 65,535<br><b>Default:</b> 0                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

---

## falling-threshold (Health Monitor)

---

|                                 |                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold <i>percentage</i>;</code>                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp health-monitor]                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                      |
| <b>Description</b>              | Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range. |
| <b>Options</b>                  | <p><b><i>percentage</i></b>—Lower threshold for the alarm entry.</p> <p><b>Range:</b> 1 through 100</p> <p><b>Default:</b> 70 percent of the maximum possible value</p>                                                                                                |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">rising-threshold on page 1807</a></li><li>• <a href="#">Configuring Health Monitoring on page 6200</a></li></ul>                                                                                                   |

## falling-threshold (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold <i>integer</i>;</code>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon alarm <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.                                                                                                                                           |
| <b>Options</b>                  | <p><i>integer</i>—Lower threshold for the alarm entry.</p> <p><b>Range:</b> -2,147,483,648 through 2,147,483,647</p> <p><b>Default:</b> 20 percent less than the <b>rising-threshold</b> value</p>                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## falling-threshold-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon alarm <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Set the interval between samples after the rising threshold is exceeded and the value of the sample starts to drop. If the value of the sample drops and exceeds the falling threshold, the regular sampling interval is used.                                                                                                                                                                                              |
| <b>Options</b>                  | <p><i>interval</i>—Time between samples, in seconds.</p> <p><b>Range:</b> 1 through 2,147,483,647 seconds</p> <p><b>Default:</b> 60 seconds</p>                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## filter-duplicates

---

|                                 |                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter-duplicates;</code>                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit snmp]</code>                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                        |
| <b>Description</b>              | Filter duplicate <b>Get</b> , <b>GetNext</b> , or <b>GetBulk</b> SNMP requests.                                          |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p> |

## filter-interfaces

---

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>filter-interfaces {<br/>    all-internal-interfaces;<br/>    interfaces <i>interface</i><br/>}</pre>                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                        |
| <b>Description</b>              | Filter out information related to specific interfaces from the output of SNMP <b>Get</b> and <b>GetNext</b> requests performed on interface-related MIBs.                                                                                                                |
| <b>Options</b>                  | <p><b>all-internal-interfaces</b>—Filter out information from SNMP <b>Get</b> and <b>GetNext</b> requests for all internal interfaces.</p> <p><b>interfaces</b>—Filter out information from SNMP <b>Get</b> and <b>GetNext</b> requests for the specified interface.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Filtering Interface Information Out of SNMP Get and GetNext Output</i></li></ul>                                                                                                                                              |

## group (Associating a Security Name)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group <i>group-name</i>;</pre>                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm security-to-group security-model (usm   v1   v2c)<br><i>security-name security-name</i> ]                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Associate a security name with a group composed of users with the same access privileges. The security name is used during authentication of SNMP messages, and is mapped to a username.       |
| <b>Options</b>                  | <b>group-name</b> —Collection of SNMP security names that share the same SNMPv3 access privileges.                                                                                             |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Group</i></li></ul>                                                                                                                 |



## group (Configuring Access Privileges)

```
Syntax group group-name {
 (default-context-prefix | context-prefix context-prefix){
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }
```

**Hierarchy Level** [edit snmp v3 vacm access]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

**Options** *group-name*—SNMPv3 group name created for the SNMPv3 group.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Group*

## health-monitor

---

|                                 |                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>health-monitor {<br/>    falling-threshold <i>percentage</i>;<br/>    interval <i>seconds</i>;<br/>    rising-threshold <i>percentage</i>;<br/>}</pre>                         |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                   |
| <b>Description</b>              | <p>Configure health monitoring.</p> <p>The remaining statements are explained separately.</p>                                                                                       |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6200</a></li><li>• <a href="#">Understanding Health Monitoring on page 6123</a></li></ul> |

## history

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>history <i>history-index</i> {     <i>bucket-size</i> <i>number</i>;     interface <i>interface-name</i>;     interval <i>seconds</i>;     owner <i>owner-name</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit snmp rmon]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent on the network to monitor all the traffic flowing among devices on all connected LAN segments. The RMON history feature collects statistics in accordance with user-configurable parameters.</p> <p>The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. If you use the <b>history</b> statement, you must also configure the <b>interface <i>interface-name</i></b> statement.</p> |
| <b>Options</b>                  | <p><b>history-index</b>—Provide a number for this history entry.</p> <p><b>Range:</b> 1 through 65535</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>                                                                                                                                                                                                                                                                                                                               |

## interface (SNMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface [ <i>interface-names</i> ];</code>                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit snmp]</code>                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the interfaces on which SNMP requests can be accepted.                                                                                                                               |
| <b>Default</b>                  | If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.                                                                                    |
| <b>Options</b>                  | <i>interface-names</i> —Names of one or more logical interfaces.                                                                                                                               |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 6197</a></li></ul>                                               |

## interface (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon history <i>history-index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Specify the interface to be monitored in the specified RMON history entry.</p> <p>Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.</p>                                                                                              |
| <b>Options</b>                  | <i>interface-name</i> —Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## interval (Health Monitor)

---

|                                 |                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interval <i>seconds</i>;</code>                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit snmp health-monitor]</code>                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                            |
| <b>Description</b>              | Configure the interval between sampling of the object being monitored by the health monitor.                                                 |
| <b>Options</b>                  | <p><i>seconds</i>—Time between samples, in seconds.</p> <p><b>Range:</b> 1 through 2147483647 seconds</p> <p><b>Default:</b> 300 seconds</p> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Health Monitoring on page 6200</a></li> </ul>                               |

## interval (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ],<br>[edit snmp rmon history <i>index</i> ]                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the interval over which data is to be sampled for the specified alarm or interface.                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | 60 sec for alarm sampling.<br><br>1800 sec for history sampling.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>seconds</i> —Interval at which data is to be sampled for the specified alarm or interface.                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## local-engine

**Syntax**

```
local-engine {
 user username {
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 authentication-sha {
 authentication-password authentication-password;
 }
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
}
```

**Hierarchy Level** [edit snmp v3 [usm](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure local engine information for the user-based security model (USM).  
  
The remaining statements are explained separately.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Creating SNMPv3 Users on page 6201](#)


## location

---

|                                 |                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>location <i>location</i>;</code>                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                      |
| <b>Description</b>              | Define the value of the MIB II <b>sysLocation</b> object, which is the physical location of the managed system.        |
| <b>Options</b>                  | <b><i>location</i></b> —Location of the local system. You must enclose the name within quotation marks (" ").          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Location for a Device Running Junos OS</i></li></ul> |



## logical-system (SNMP)

|                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                    | <pre>logical-system <i>logical-system-name</i> {     <i>routing-instance routing-instance-name</i>; }</pre>                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                           | <pre>[edit snmp community <i>community-name</i>], [edit snmp <i>trap-group</i>], [edit snmp trap-options] [edit snmp v3target-address <i>target-address-name</i>]</pre>                                                                                                                                                                                                                             |
| <b>Release Information</b>                                                                                                                                                                                                                                                       | <p>Statement introduced in Junos OS Release 9.3</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                            |
| <div>  <p><b>NOTE:</b> The <code>logical-system</code> statement replaces the <code>logical-router</code> statement, and is backward-compatible with Junos OS Release 8.3 and later.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                               | <p>Specify a logical system name for SNMP v1 and v2c clients.</p> <p>Include at the <code>[edit snmp trap-options]</code> hierarchy level to specify a logical-system address as the source address of an SNMP trap.</p> <p>Include at the <code>[edit snmp v3 target-address]</code> hierarchy level to specify a logical-system name as the destination address for an SNMPv3 trap or inform.</p> |
| <b>Options</b>                                                                                                                                                                                                                                                                   | <p><b><i>logical-system-name</i></b>—Name of the logical system.</p> <p><b><i>routing-instance routing-instance-name</i></b>—Statement to specify a routing instance associated with the logical system.</p>                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                  | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                     | <ul style="list-style-type: none"> <li>• <i>Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</i></li> <li>• <i>Configuring the Trap Target Address</i></li> </ul>                                                                                                                                                                                                                    |

## message-processing-model

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | message-processing-model (v1   v2c   v3);                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the message processing model to be used when generating SNMP notifications.                                                                                                          |
| <b>Options</b>                  | v1—SNMPv1 message process model.<br><br>v2c—SNMPv2c message process model.<br><br>v3—SNMPv3 message process model.                                                                             |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Message Processing Model</i></li></ul>                                                                                              |

## name

---

|                                 |                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | name <i>name</i> ;                                                                                            |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                             |
| <b>Description</b>              | Set the system name from the command-line interface.                                                          |
| <b>Options</b>                  | <i>name</i> —System name override.                                                                            |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Name</i></li></ul>                          |

## nonvolatile

---

|                                 |                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | nonvolatile {<br><b>commit-delay</b> <i>seconds</i> ;<br>}                                                                   |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                            |
| <b>Description</b>              | Configure options for SNMP <b>Set</b> requests.<br><br>The statement is explained separately.                                |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Commit Delay Timer</i></li> <li>• <i>commit-delay</i></li> </ul> |

## notify

---

|                                 |                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | notify <i>name</i> {<br>tag <i>tag-name</i> ;<br><b>type</b> (trap   inform);<br>}                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>type inform</b> option added in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                   |
| <b>Description</b>              | Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.                                                                                                                                                       |
| <b>Options</b>                  | <p><b>name</b>—Name assigned to the notification.</p> <p><b>tag-name</b>—Notifications are sent to all targets configured with this tag.</p> <p><b>type</b>—Notification type is <b>trap</b> or <b>inform</b>. Traps are unconfirmed notifications. Informs are confirmed notifications.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Inform Notification Type and Target Address</i></li> <li>• <i>Configuring the SNMPv3 Trap Notification</i></li> </ul>                                                                                                            |

## notify-filter (Applying to the Management Target)

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-filter <i>profile-name</i>;</code>                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 <b>target-parameters</b> <i>target-parameters-name</i>]</code>                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.           |
| <b>Description</b>              | Configure the notify filter applied to a specific set of SNMPv3 target parameters. Target parameters are the message processing and security parameters for notifications sent to a target SNMP manager. |
| <b>Options</b>                  | <b><i>profile-name</i></b> —Name of the notify filter to apply to notifications.                                                                                                                         |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Applying the Trap Notification Filter</i></li></ul>                                                                                                           |

## notify-filter (Configuring the Profile Name)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-filter <i>profile-name</i> {<br/>oid <i>oid</i> (include   exclude);<br/>}</code>                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3]</code>                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.                                       |
| <b>Options</b>                  | <b><i>profile-name</i></b> —Name assigned to the notify filter.<br><br>The remaining statement is explained separately.                                                                        |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Trap Notification Filter</i></li><li>• <i>oid (SNMP)</i></li></ul>                                                                  |

## notify-view

---

|                                 |                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-view <i>view-name</i>;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i> ) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.             |
| <b>Description</b>              | Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                           |
| <b>Options</b>                  | <b><i>view-name</i></b> —Name of the view to which the SNMP user group has access.                                                                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6197</a></li> <li>• <a href="#">Configuring the Notify View</a></li> </ul>                                              |

## oid

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oid <i>object-identifier</i> (exclude  include);</code>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp view <i>view-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify an object identifier (OID) used to represent a subtree of MIB objects.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>exclude</b> —Exclude the subtree of MIB objects represented by the specified OID.<br><br><b>include</b> —Include the subtree of MIB objects represented by the specified OID.<br><br><b><i>object-identifier</i></b> —OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6197</a></li> </ul>                                                                                                                                                                                                                                                                                                                                           |

## oid (SNMPv3)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oid <i>oid</i> (include   exclude);</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 notify-filter <i>profile-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>exclude</b>—Exclude the subtree of MIB objects represented by the specified OID.</p> <p><b>include</b>—Include the subtree of MIB objects represented by the specified OID.</p> <p><b>oid</b>—Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p> |
| <b>Required Privilege Level</b> | <p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                            |

## owner

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>owner <i>owner-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon history <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Specify the user or group responsible for this RMON history configuration.                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>owner-name</b>—User or group responsible for this configuration.</p> <p><b>Range:</b> 0 through 32 alphanumeric characters</p>                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## parameters

---

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>parameters {   message-processing-model (v1   v2c   v3);   security-level (none   authentication   privacy);   security-model (usm   v1   v2c);   security-name security-name; }</pre>                   |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-parameters <i>target-parameters-name</i> ]                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | <p>Configure a set of target parameters for message processing and security.</p> <p>The remaining statements are explained separately.</p>                                                                    |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Defining and Configuring the Trap Target Parameters</i></li> </ul>                                                                                                |

## port (SNMP)

---

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port port-number;</code>                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | Configure a UDP port number for an SNMP target.                                                                                                                                                               |
| <b>Default</b>                  | If you omit this statement, the default port is 162.                                                                                                                                                          |
| <b>Options</b>                  | <i>port-number</i> —Port number for the SNMP target.                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Port</i></li> </ul>                                                                                                                               |

## privacy-3des

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>privacy-3des {<br/>    <b>privacy-password</b> <i>privacy-password</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the triple Data Encryption Standard (3DES) as the privacy type for the SNMPv3 user.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>privacy-password</b> <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"><li>• The password must be at least eight characters long.</li><li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li></ul> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Encryption Type</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                       |



## privacy-aes128

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>privacy-aes128 {   privacy-password <i>privacy-password</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <pre>[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]</pre>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure the Advanced Encryption Standard encryption algorithm (CFB128-AES-128 Privacy Protocol) for the SNMPv3 user.</p>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>privacy-password <i>privacy-password</i></b>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> <li>• The password must be at least eight characters long.</li> <li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li> </ul> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Encryption Type</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                        |

## privacy-des

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>privacy-des {<br/>    <b>privacy-password</b> <i>privacy-password</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the Data Encryption Standard (DES) as the privacy type for the SNMPv3 user.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>privacy-password</b> <i>privacy-password</i> —Password that a user enters. The password is then converted into a key that is used for encryption.<br><br>SNMPv3 has special requirements when you create plain-text passwords on a router or switch: <ul style="list-style-type: none"><li>• The password must be at least eight characters long.</li><li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li></ul> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Encryption Type</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                 |

## privacy-none

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>privacy-none;</code>                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit snmp v3 usm local-engine user <i>username</i> ],<br>[edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure that no encryption be used for the SNMPv3 user.                                                                                                                                      |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Encryption Type</i></li></ul>                                                                                                       |

## privacy-password

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>privacy-password <i>privacy-password</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit snmp v3 usm local-engine user <i>username</i> privacy-3des],<br/>         [edit snmp v3 usm local-engine user <i>username</i> privacy-aes128],<br/>         [edit snmp v3 usm local-engine user <i>username</i> privacy-des],<br/>         [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-3des],<br/>         [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-aes128],<br/>         [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-des]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure a privacy password for the SNMPv3 user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>privacy-password</i></b>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> <li>• The password must be at least eight characters long.</li> <li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li> </ul>                                                                 |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.<br/>         snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Encryption Type</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## read-view

---

|                                 |                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>read-view view-name;</code>                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i> ) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.             |
| <b>Description</b>              | Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                        |
| <b>Options</b>                  | <i>view-name</i> —The name of the view to which the SNMP user group has access.                                                                                                                            |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Read View</i></li><li>• <a href="#">Configuring MIB Views on page 6197</a></li></ul>                                                            |

## remote-engine

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> remote-engine <i>engine-id</i> {   user <i>username</i> {     authentication-md5 {       authentication-password <i>authentication-password</i>;     }     authentication-none;     authentication-sha {       authentication-password <i>authentication-password</i>;     }     privacy-aes128 {       privacy-password <i>privacy-password</i>;     }     privacy-des {       privacy-password <i>privacy-password</i>;     }     privacy-3des {       privacy-password <i>privacy-password</i>;     }     privacy-none {       privacy-password <i>privacy-password</i>;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit snmp v3 usm]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the remote engine information for the user-based security model (USM). To send inform messages to an SNMPv3 user on a remote device, you must configure the engine identifier for the SNMP agent on the remote device where the user resides.                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>engine-id</i></b>—Engine identifier. Used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Remote Engine and Remote User</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## request-type

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | request-type (get-next-request   get-request   walk-request);                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Extend monitoring to a specific SNMP object instance ( <b>get-request</b> ), to all object instances belonging to a MIB branch ( <b>walk-request</b> ), or to the next object instance after the instance specified in the configuration ( <b>get-next-request</b> ).                                                                                                                                                 |
| <b>Default</b>                  | walk-request                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>get-next-request</b> —Perform an SNMP get next request.<br><br><b>get-request</b> —Perform an SNMP get request.<br><br><b>walk-request</b> —Perform an SNMP walk request.                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

---

## retry-count (SNMPv3)

---

|                                 |                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | retry-count <i>number</i> ;                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                           |
| <b>Description</b>              | Configure the retry count for SNMP informs.                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>number</i> —Maximum number of times the inform is transmitted if no acknowledgment is received. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded.<br><b>Default:</b> 3 times |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Informs on page 6206</a></li><li>• <i>timeout</i></li></ul>                                                                                                                             |

## rising-event-index

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rising-event-index <i>index</i>;</code>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">snmp rmon alarm index</a> ]                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set the index of the event entry that is used when a rising alarm threshold is exceeded. The rising-event index is specified when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>index</i></b> —Index of the event entry that is used when a rising threshold is exceeded.<br><b>Range:</b> 0 through 65,535<br><b>Default:</b> 0                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |



---

## rising-threshold (Health Monitor)

---

|                                 |                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rising-threshold <i>percentage</i>;</code>                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit snmp health-monitor]                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                    |
| <b>Description</b>              | Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range. |
| <b>Options</b>                  | <p><b><i>percentage</i></b>—Upper threshold for the alarm entry.</p> <p><b>Range:</b> 1 through 100</p> <p><b>Default:</b> 80 percent of the maximum possible value</p>                                                                                              |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6200</a></li><li>• <a href="#">falling-threshold on page 1772</a></li></ul>                                                                                                |

## rising-threshold (RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | rising-threshold <i>integer</i> ;                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set the upper threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.                                                                                                                                           |
| <b>Options</b>                  | <i>integer</i> —Upper threshold for the alarm entry.<br><b>Range:</b> -2,147,483,648 through 2,147,483,647                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## rmon

```
Syntax rmon {
 alarm index {
 description description;
 falling-event-index index;
 falling-threshold integer;
 falling-threshold-interval seconds;
 interval seconds;
 request-type;
 rising-event-index index;
 rising-threshold integer;
 sample-type (absolute-value | delta-value);
 startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
 syslog-subtag syslog-subtag;
 variable oid-variable;
 }
 event index {
 community community-name;
 description description;
 type (RMON Notification) type;
 }
 history history-index {
 bucket-size number;
 interface interface-name;
 interval seconds;
 owner owner-name;
 }
 }
```

**Hierarchy Level** [edit snmp]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Provide comprehensive network fault diagnosis, planning, and performance tuning information. RMON delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.

Junos OS supports the RMON statistics, history, alarm, and event groups.

The remaining statements are explained separately.

**Default** Disabled.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6121](#)
- [Monitoring RMON MIB Tables on page 6389](#)
- [Understanding RMON on page 6119](#)

- [Junos OS Network Management Configuration Guide](#)

---

## routing-instance (SNMP)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <code>routing-instance <i>routing-instance-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Hierarchy Level          | <code>[edit snmp community <i>community-name</i>],</code><br><code>[edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>],</code><br><code>[edit snmp <b>trap-group</b> <i>group</i>]</code>                                                                                                                                                                                                                                                                                                                  |
| Release Information      | Statement introduced in Junos OS Release 8.3.<br>Added to the <code>[edit snmp community <i>community-name</i>]</code> hierarchy level in Junos OS Release 8.4.<br>Added to the <code>[edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>]</code> hierarchy level in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                     |
| Description              | Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance.<br><br>If the routing instance is defined within a logical system, include the <b>logical-system <i>logical-system-name</i></b> statement at the <code>[edit snmp community <i>community-name</i>]</code> hierarchy level and specify the <b>routing-instance</b> statement under the <code>[edit snmp community <i>community-name</i> logical-system <i>logical system-name</i>]</code> hierarchy level. |
| Options                  | <b><i>routing-instance-name</i></b> —Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Required Privilege Level | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups</a></li><li>• <a href="#">Configuring the Source Address for SNMP Traps</a></li><li>• <a href="#">Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</a></li></ul>                                                                                                                                                                                                                                                                             |

## sample-type

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | sample-type (absolute-value   delta-value);                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the method of sampling the selected variable (monitored object). When you configure an SNMP RMON alarm, you can specify the sample type.                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>absolute-value</b>—Actual value of the selected variable is used when comparing against the thresholds.</p> <p><b>delta-value</b>—Difference between samples of the selected variable is used when comparing against the thresholds.</p>                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## security-level (Defining Access Privileges)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>security-level (authentication   none   privacy) {<br/>    notify-view view-name;<br/>    read-view view-name;<br/>    write-view view-name;<br/>}</pre>                                  |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i> ) security-model (any   usm   v1   v2c)]                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Define the security level used for access privileges.                                                                                                                                          |
| <b>Default</b>                  | none                                                                                                                                                                                           |
| <b>Options</b>                  | <b>authentication</b> —Provide authentication but no encryption.<br><br><b>none</b> —No authentication and no encryption.<br><br><b>privacy</b> —Provide authentication and encryption.        |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Level</i></li></ul>                                                                                                        |

---

## security-level (Generating SNMP Notifications)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | security-level (authentication   none   privacy);                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security level to use when generating SNMP notifications.                                                                                                                        |
| <b>Default</b>                  | none                                                                                                                                                                                           |
| <b>Options</b>                  | <b>authentication</b> —Provide authentication but no encryption.<br><br><b>none</b> —No authentication and no encryption.<br><br><b>privacy</b> —Provide authentication and encryption.        |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Level</i></li></ul>                                                                                                        |

## security-model (Access Privileges)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-model (usm   v1   v2c);</code>                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>)]</code>                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.                                                                     |
| <b>Options</b>                  | <code>usm</code> —SNMPv3 security model.<br><br><code>v1</code> —SNMPv1 security model.<br><br><code>v2c</code> —SNMPv2c security model.                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>                                                                                                        |



## security-model (Group)

---

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>security-model (usm   v1   v2c) {   security-name security-name {     group group-name;   } }</pre>                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm <a href="#">security-to-group</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | Define a security model for an SNMPv3 group and associate the security name of a user with the group. All users in the group have the same access privileges.                                                 |
| <b>Options</b>                  | <p><b>usm</b>—SNMPv3 security model.</p> <p><b>v1</b>—SNMPv1 security model.</p> <p><b>v2c</b>—SNMPv2c security model.</p>                                                                                    |
| <b>Required Privilege Level</b> | <p><b>snmp</b>—To view this statement in the configuration.</p> <p><b>snmp-control</b>—To add this statement to the configuration.</p>                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Security Model</i></li> </ul>                                                                                                                       |

## security-model (SNMP Notifications)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-model (usm   v1   v2c);</code>                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.                                                                                           |
| <b>Options</b>                  | <code>usm</code> —SNMPv3 security model.<br><br><code>v1</code> —SNMPv1 security model.<br><br><code>v2c</code> —SNMPv2c security model.                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>                                                                                                        |

## security-name (Community String)

|                            |                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>security-name <i>security-name</i>;</code>                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <code>[edit snmp v3 <i>snmp-community</i> <i>community-index</i>]</code>                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                 |
| <b>Description</b>         | Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level. |
| <b>Options</b>             | <i>security-name</i> —Name that is used for messaging security and user access control.                                                                                                                                                                        |



**NOTE:** The security name must match the configured security name at the `[edit snmp v3 target-parameters target-parameters-name parameters]` hierarchy level when you configure traps or informs.


|                              |                                                                                           |
|------------------------------|-------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | snmp—To view this statement in the configuration.                                         |
| <b>Level</b>                 | snmp-control—To add this statement to the configuration.                                  |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Configuring the Security Names</i></li> </ul> |

## security-name (Security Group)

---

|                                 |                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-name security-name {<br/>    group group-name;<br/>}</code>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm security-to-group <b>security-model</b> (usm   v1   v2c)]                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                       |
| <b>Description</b>              | Associate the security name of a user (for SNMPv3 clients) or a community string (for SNMPv1 and SNMPv2c clients) with a configured security group.                                                                                                                                                                                  |
| <b>Options</b>                  | <b>security-name</b> —SNMPv3 secure username configured at the [edit snmp v3 usm local-engine <b>user username</b> ] hierarchy level that is used for messaging security. For SNMPv1 and SNMPv2c, the security name is the community string configured at the [edit snmp v3 snmp-community <b>community-index</b> ] hierarchy level. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Assigning Security Names to Groups</i></li><li>• <a href="#">Assigning a Security Name to a Group on page 6204</a></li></ul>                                                                                                                                                              |

## security-name (SNMP Notifications)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <code>security-name <i>security-name</i>;</code>                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Configure the security name used when generating SNMP notifications.                                                                                                                                                                          |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b><i>security-name</i></b> —If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification. |
| <div>  <p><b>NOTE:</b> The access privileges for the group associated with this security name must allow this notification to be sent.</p> <p>If you are using the v1 or v2 security models, the security name at the <code>[edit snmp v3 vacm security-to-group]</code> hierarchy level must match the security name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.</p> </div> |                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li>• <i>Configuring the Security Name</i></li> </ul>                                                                                                                                                      |

## security-to-group

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>security-to-group {<br/>  security-model (usm   v1   v2c) {<br/>    group group-name;<br/>    security-name security-name;<br/>  }<br/>}</pre>                                            |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm]                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.<br><br>The remaining statements are explained separately.              |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Assigning Security Model and Security Name to a Group</i></li></ul>                                                                                 |

## snmp

```

Syntax snmp {
 client-list client-list-name {
 ip-addresses;
 }
 community community-name {
 authorization authorization;
 client-list-name client-list-name;
 clients {
 address restrict;
 }
 logical-system logical-system-name {
 routing-instance routing-instance-name {
 clients {
 addresses;
 }
 }
 }
 routing-instance routing-instance-name {
 clients {
 addresses;
 }
 }
 view view-name;
 }
 contact contact;
 description description;
 filter-duplicates;
 filter-interfaces;
 health-monitor {
 falling-threshold integer;
 interval seconds;
 rising-threshold integer;
 }
 interface [interface-names];
 location location;
 name name;
 nonvolatile {
 commit-delay seconds;
 }
 rmon {
 alarm index {
 description description;
 falling-event-index index;
 falling-threshold integer;
 falling-threshold-interval seconds;
 interval seconds;
 request-type;
 rising-event-index index;
 rising-threshold integer;
 sample-type (absolute-value | delta-value);
 startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
 syslog-subtag syslog-subtag;
 }
 }
}

```

```
 variable oid-variable;
 }
 event index {
 community community-name;
 description description;
 type type;
 }
 history history-index {
 bucket-size number;
 interface interface-name;
 interval seconds;
 owner owner-name;
 }
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable> <match
 regular-expression>;
 flag flag;
}
trap-group group-name {
 categories {
 category;
 }
 destination-port port-number;
 routing-instance routing-instance-name;
 targets {
 address;
 }
 version (all | v1 | v2);
}
trap-options {
 agent-address outgoing-interface;
 source-address address;
}
v3 {
 notify name {
 tag tag-name;
 type trap;
 }
 notify-filter profile-name {
 oid object-identifier (include | exclude);
 }
 snmp-community community-index {
 community-name community-name;
 security-name security-name;
 tag tag-name;
 }
 target-address target-address-name {
 address address;
 address-mask address-mask;
 logical-system logical-system;
 port port-number;
 retry-count number;
 routing-instance routing-instance-name;
 tag-list tag-list;
 target-parameters target-parameters-name;
```



```

 timeout seconds;
 }
 target-parameters target-parameters-name {
 notify-filter profile-name;
 parameters {
 message-processing-model (v1 | v2c | V3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 security-name security-name;
 }
 }
 usm {
 local-engine {
 user username {
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none;
 }
 }
 remote-engine engine-id {
 user username {
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
 }
 }
}

```

```
}
vacm {
 access {
 group group-name {
 (default-context-prefix | context-prefix context-prefix) {
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }
 }
}
security-to-group {
 security-model (usm | v1 | v2c) {
 security-name security-name {
 group group-name;
 }
 }
}
view view-name {
 oid object-identifier (include | exclude);
}
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure SNMP.

The remaining statements are explained separately.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding the Implementation of SNMP on page 6107](#)
- [Configuring SNMP on page 1703](#)

## snmp-community

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | snmp-community <i>community-index</i> {<br><i>community-name</i> <i>community-name</i> ;<br><i>security-name</i> <i>security-name</i> ;<br>tag <i>tag-name</i> ;<br>}                          |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the SNMP community which authorizes SNMPv1 or SNMPv2c clients in an SNMPv3 system.                                                                                                   |
| <b>Options</b>                  | <i>community-index</i> —(Optional) String that identifies an SNMP community.<br><br>The remaining statements are explained separately.                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the SNMPv3 Community</i></li> </ul>                                                                                                    |

## source-address (SNMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-address <i>address</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp trap-options]                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Set the source address of every SNMP trap packet sent by this switch to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>address</i>—Source address of SNMP traps. You can configure the source address of trap packets two ways: <b>lo0</b> or a valid IPv4 address configured on one of the interfaces. The value <b>lo0</b> indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface <b>lo0</b>.</p> <p><b>Default:</b> Disabled. (The source address is the address of the outgoing interface.)</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Source Address for SNMP Traps</i></li> </ul>                                                                                                                                                                                                                                                                                                                                             |

## startup-alarm

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | startup-alarm (falling-alarm   rising-alarm   rising-or-falling-alarm);                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Set an initial alarm that is sent after the configured SNMP RMON alarm becomes active.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | rising-or-falling-alarm                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>falling-alarm</b>—Generated if the first sample after the alarm becomes active is equal to or greater than the falling threshold.</p> <p><b>rising-alarm</b>—Generated if the first sample after the alarm becomes active is equal to or greater than the rising threshold.</p> <p><b>rising-or-falling-alarm</b>—Generated if the first sample after the alarm entry becomes active is equal to or greater than either the rising threshold or the falling threshold.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>                                                            |

## syslog-subtag

---

|                                 |                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>syslog-subtag <i>syslog-subtag</i>;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Add the <b>syslog-subtag</b> tag to the system log message. The tag should not exceed 80 uppercase characters.                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## tag (Configuring Notification Targets)

---

|                                 |                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag <i>tag-name</i>;</code>                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit snmp v3 notify <i>name</i> ]                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                       |
| <b>Description</b>              | Configure a set of target addresses to receive SNMP traps or informs (for IPv4 packets only).                                                                                                                                           |
| <b>Options</b>                  | <b>tag-name</b> —Define the target addresses to which an SNMP notification is sent. Target addresses containing the same tag in their tag list are sent the same notification. The <b>tag-name</b> is not included in the notification. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                           |

## tag (Configuring the SNMP Community)

---

|                                 |                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag tag-name;</code>                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp v3 snmp-community <i>community-index</i> ]                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                             |
| <b>Description</b>              | Configure a set of SNMP managers that are authorized to use a community string.                               |
| <b>Options</b>                  | <i>tag-name</i> —Identify the set of addresses for the SNMP managers authorized to use the community string.  |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration. |

## tag-list

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag-list tag-list;</code>                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure an SNMP tag list used to select target addresses.                                                                                                                                    |
| <b>Options</b>                  | <i>tag-list</i> —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.                          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Trap Target Address</i></li></ul>                                                                                                   |

---

## target-address

---

**Syntax**    `target-address target-address-name {  
                  address address;  
                  address-mask address-mask;  
                  port port-number;  
                  retry-count number;  
                  tag-list tag-list;  
                  target-parameters target-parameters-name;  
                  timeout seconds;  
                  }`

**Hierarchy Level**    [edit snmp v3]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure the address of an SNMP management application and the parameters to be used in sending notifications.

**Options**    *target-address-name*—String that identifies the target address.

The remaining statements are explained separately.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                  snmp-control—To add this statement to the configuration.

## target-parameters

---

**Syntax** At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {
 profile-name;
 parameters {
 message-processing-model (v1 | v2c | V3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 security-name security-name;
 }
}
```

At the `[edit snmp v3 target-address target-address-name]` hierarchy level:

```
target-parameters target-parameters-name;
```

**Hierarchy Level** `[edit snmp v3]`  
`[edit snmp v3 target-address target-address-name]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the `[edit snmp v3]` hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the `[edit snmp v3 target-parameters target-parameters-name]` hierarchy level to the target address configuration at the `[edit snmp v3]` hierarchy level.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- *Defining and Configuring the Trap Target Parameters*
- *Applying Target Parameters*



## targets

---

|                                 |                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>targets {<br/>    <i>address</i>;<br/>}</code>                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-group <i>group-name</i>]</code>                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                        |
| <b>Description</b>              | Configure one or more systems to receive SNMP traps.                                                                     |
| <b>Options</b>                  | <b><i>address</i></b> —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Trap Groups on page 6195</a></li> </ul>            |

## timeout

---

|                                 |                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>timeout <i>seconds</i>;</code>                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-address <i>target-address-name</i>]</code>                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                       |
| <b>Description</b>              | Configure the timeout period (in seconds) for SNMP informs.                                                                                                             |
| <b>Default</b>                  | 15 seconds                                                                                                                                                              |
| <b>Options</b>                  | <b><i>seconds</i></b> —Number of seconds to wait for an inform acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                           |

## traceoptions (SNMP)

---

|                            |                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable  <br/>        no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>    no-remote-trace;<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit snmp]                                                                                                                                                                                                                                                          |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                    |
| <b>Description</b>         | Track the activities of SNMP agents on the switch and record the information in log files.                                                                                                                                                                           |



**NOTE:** The **traceoptions** statement is not supported on the QFabric system.

---

The output of the tracing operations is placed into log files in the **/var/log** directory. Each log file is named after the SNMP agent that generates it. The following logs are created in the **/var/log** directory when the **traceoptions** statement is used:

- chassisd
- craftd
- ilmid
- mib2d
- rmopd
- serviced
- snmpd

**Options** **file *filename***—By default, the name of the log file that records trace output is the name of the process being traced (for example, **mib2d** or **snmpd**). Use this option to specify another name.

**files *number***—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, **snmpd**) reaches its maximum size, it is archived by being renamed to **snmpd.0**. The previous **snmpd.1** is renamed to **snmpd.2**, and so on. The oldest archived file is deleted.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Log all SNMP events.

- **configuration**—Log reading of configuration at the **[edit snmp]** hierarchy level.
- **database**—Log events involving storage and retrieval in the events database.
- **events**—Log important events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **policy**—Log policy processing.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **server**—Log communication with processes that are generating events.
- **subagent**—Log subagent restarts.
- **timer-events**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**size *size***—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

**Range:** 10 KB through 1 GB

**Default:** 1000 KB

**world-readable | no-world-readable**—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

**Required Privilege Level**

|                                                          |
|----------------------------------------------------------|
| snmp—To view this statement in the configuration.        |
| snmp-control—To add this statement to the configuration. |

**Related Documentation**

- [Understanding Tracing and Logging Operations on page 6081](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 6393](#)

## trap-group

---

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>trap-group group-name {<br/>    categories {<br/>        category;<br/>    }<br/>    destination-port port-number;<br/>    targets {<br/>        address;<br/>    }<br/>}</pre>                                                                                                     |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for QFX Series switches.                                                                                                                                                                                                                   |
| <b>Description</b>              | Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent. |
| <b>Options</b>                  | <p><b>group-name</b>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>                                                                                                             |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6195</a></li></ul>                                                                                                                                                                              |

---

## trap-options

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>trap-options {<br/>    agent-address outgoing-interface;<br/>    source-address address;<br/>}</pre>                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Options</i></li></ul>                                                                                                                                                                                                                                                                                            |

## type (RMON Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>type type;</code>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit snmp rmon event <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the type of notification generated when a rising or falling threshold is crossed.                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | <code>log-and-trap</code>                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>type</b>—Type of notification. It can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>log</b>—Add an entry to the <b>logTable</b> object.</li><li>• <b>log-and-trap</b>—Send an SNMP trap and add a log entry.</li><li>• <b>none</b>—No notifications are sent.</li><li>• <b>snmptrap</b>—Send an SNMP trap.</li></ul>                                                                  |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## type (SNMPv3)

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>type (inform   trap);</code>                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 notify <i>name</i>]</code>                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>inform</b> option added in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the type of SNMP notification.                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>inform</b> —Defines the type of notification as an inform. SNMP informs are confirmed notifications.<br><br><b>trap</b> —Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.                                     |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Informs on page 6206</a></li> <li>• <a href="#">Configuring the SNMPv3 Trap Notification</a></li> </ul>                                                                         |

## user

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>user <i>username</i>;</code>                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 usm local-engine],</code><br><code>[edit snmp v3 usm remote-engine <i>engine-id</i>]</code>                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Specify a user associated with an SNMPv3 group on a local or remote SNMP engine.                                                                                                               |
| <b>Options</b>                  | <b><i>username</i></b> —SNMPv3 user-based security model (USM) username.                                                                                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Creating SNMPv3 Users on page 6201</a></li> </ul>                                                                                         |

## usm

```

Syntax usm {
 local-engine {
 user username {
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 authentication-sha {
 authentication-password authentication-password;
 }
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
 }
 remote-engine engine-id {
 user username {
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-none;
 authentication-sha {
 authentication-password authentication-password;
 }
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
 }
 }
}

```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 7.4.



Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure user-based security model (USM) information.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.

**Level** snmp-control—To add this statement to the configuration.

**Related** • [Creating SNMPv3 Users on page 6201](#)

**Documentation** • *Configuring the Remote Engine and Remote User*

## v3

---

```
Syntax v3 {
 notify name {
 tag tag-name;
 type trap;
 }
 notify-filter profile-name {
 oid object-identifier (include | exclude);
 }
 snmp-community community-index {
 community-name community-name;
 security-name security-name;
 tag tag-name;
 }
 target-address target-address-name {
 address address;
 address-mask address-mask;
 port port-number;
 retry-count number;
 tag-list tag-list;
 target-parameters target-parameters-name;
 timeout seconds;
 }
 target-parameters target-parameters-name {
 notify-filter profile-name;
 parameters {
 message-processing-model (v1 | v2c | v3);
 security-level (authentication | none | privacy);
 security-model (usm | v1 | v2c);
 security-name security-name;
 }
 }
}

usm {
 local-engine {
 user username {
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none;
 }
 }
}
```

```

}
remote-engine engine-id {
 user username {
 authentication-md5 {
 authentication-password authentication-password;
 }
 authentication-sha {
 authentication-password authentication-password;
 }
 authentication-none;
 privacy-aes128 {
 privacy-password privacy-password;
 }
 privacy-des {
 privacy-password privacy-password;
 }
 privacy-3des {
 privacy-password privacy-password;
 }
 privacy-none {
 privacy-password privacy-password;
 }
 }
}
}
}
vacm {
 access {
 group group-name {
 (default-context-prefix | context-prefix context-prefix) {
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }
 }
}
security-to-group {
 security-model (usm | v1 | v2c) {
 security-name security-name {
 group group-name;
 }
 }
}
}
}
}

```

**Hierarchy Level** [edit snmp]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure SNMPv3.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.  
**Level** snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6118](#)

---

## vacm

---

**Syntax**

```
vacm {
 access {
 group group-name {
 (default-context-prefix | context-prefix context-prefix){
 security-model (any | usm | v1 | v2c) {
 security-level (authentication | none | privacy) {
 notify-view view-name;
 read-view view-name;
 write-view view-name;
 }
 }
 }
 }
 }
 security-to-group {
 security-model (usm | v1 | v2c);
 security-name security-name {
 group group-name;
 }
 }
}
```

**Hierarchy Level** [edit snmp v3]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure view-based access control model (VACM) information, including access privileges such as security model and security level for a group of users.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.  
**Level** snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Defining Access Privileges for an SNMP Group](#)

## variable

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>variable <i>oid-variable</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Set the object identifier (OID) of the MIB object (also called variable) to be monitored when you configure an SNMP RMON alarm. If the value of the monitored variable exceeds the configured rising threshold or falling threshold, an alarm is triggered and a corresponding event may be generated.                                                                                                                      |
| <b>Options</b>                  | <i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1) or the name of the MIB object—for example, <code>ifInOctets.1</code> .                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul> |

## version

---

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | version (all   v1   v2);                                                                                                         |
| <b>Hierarchy Level</b>          | [edit snmp trap-group <i>group-name</i> ]                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                |
| <b>Description</b>              | Specify the version number of SNMP traps.                                                                                        |
| <b>Default</b>                  | all—Send an SNMPv1 and SNMPv2 trap for every trap condition.                                                                     |
| <b>Options</b>                  | all—Send an SNMPv1 and SNMPv2 trap for every trap condition.<br><br>v1—Send SNMPv1 traps only.<br><br>v2—Send SNMPv2 traps only. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6195</a></li></ul>                      |

## view (Configuring a MIB View)

|                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                     | <code>view <i>view-name</i> {<br/>    <b>oid</b> <i>object-identifier</i> (include   exclude);<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                            | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                        | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                | Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The <b>view</b> statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the <b>view</b> statement at the <b>[edit snmp community <i>community-name</i>]</b> hierarchy level. |
| <div>  <b>NOTE:</b> To remove an OID completely, use the <code>delete view all oid oid-number</code> command but omit the <code>include</code> parameter. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                                                                                    | <p><b><i>view-name</i></b>—Name of the view.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                   | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6197</a></li> <li>• <i>Associating MIB Views with an SNMP User Group</i></li> <li>• <a href="#">community on page 1766</a></li> </ul>                                                                                                                                                                                                                                                         |

## view (Associating MIB View with a Community)

---

|                                 |                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>view view-name;</code>                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit snmp community community-name]</code>                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                       |
| <b>Description</b>              | Associate a view with a community. A view represents a group of MIB objects.                                                                            |
| <b>Options</b>                  | <b>view-name</b> —Name of the view. You must use a view name already configured in the <b>view</b> statement at the <b>[edit snmp]</b> hierarchy level. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the SNMP Community String</i></li></ul>                                                          |

## write-view

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>write-view view-name;</code>                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm access group group-name (default-context-prefix   context-prefix context-prefix) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series switches.  |
| <b>Description</b>              | Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                          |
| <b>Options</b>                  | <b>view-name</b> —Name of the view for which the SNMP user group has write permission.                                                                                                                   |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring MIB Views on page 6197</a></li><li>• <i>Configuring the Write View</i></li></ul>                                                         |

## Configuration Statements for System Log Messages

---

- [archive \(All System Log Files\) on page 6366](#)
- [archive \(Individual System Log File\) on page 6368](#)
- [archive \(QFabric System\) on page 6369](#)



- [console \(System Logging\) on page 6370](#)
- [explicit-priority on page 6371](#)
- [facility-override on page 6371](#)
- [file \(QFabric System\) on page 6372](#)
- [file \(System Logging\) on page 6373](#)
- [files on page 6374](#)
- [host \(System\) on page 6375](#)
- [log-prefix \(System\) on page 6377](#)
- [match on page 6377](#)
- [size \(System\) on page 6378](#)
- [structured-data on page 6379](#)
- [syslog \(System\) on page 6380](#)
- [syslog \(QFabric System\) on page 6382](#)
- [time-format on page 6383](#)
- [user \(System Logging\) on page 6384](#)

## archive (All System Log Files)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>archive &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;start-time <i>time</i>&gt; &lt;transfer-interval <i>interval</i>&gt;<br/>&lt;binary-data   no-binary-data&gt;;<br/>&lt;world-readable   no-world-readable&gt; ;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit system <a href="#">syslog</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | Configure archiving properties for all system log files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>             | <p><b>files <i>number</i></b>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i>, it closes the file, compresses it, and renames it <i>logfile.0.gz</i> (the amount of data is determined by the <b>size</b> statement at this hierarchy level). The utility then opens and writes to a new file called <i>logfile</i>. When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i>, and the new file is closed, compressed, and renamed <i>logfile.0.gz</i>. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p><b>Range:</b> 1 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>size <i>size</i></b>—Maximum amount of data that the Junos OS logging utility writes to a log file <i>logfile</i> before archiving it (closing it, compressing it, and changing its name to <i>logfile.0.gz</i>). The utility then opens and writes to a new file called <i>logfile</i>.</p> <p><b>Syntax:</b> <i>x k</i> to specify the number of kilobytes, <i>x m</i> for the number of megabytes, or <i>x g</i> for the number of gigabytes</p> <p><b>Range:</b> 64 KB through 1 GB</p> <p><b>Default:</b></p> <ul style="list-style-type: none"><li>• 128 KB for EX Series switches and J Series routers</li><li>• 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch</li><li>• 10 MB for TX Matrix and TX Matrix Plus routers</li></ul> <p><b>binary-data   no-binary-data</b>—Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems)..</p> <p><b>Default:</b> no-binary-data</p> <p><b>world-readable   no-world-readable</b>—Grant all users permission to read archived log files, or restrict the permission only to the <b>root</b> user and users who have the Junos OS <b>maintenance</b> permission.</p> <p><b>Default:</b> no-world-readable</p> |

**Required Privilege** system—To view this statement in the configuration.

**Level** system-control—To add this statement to the configuration.

**Related Documentation** • [Specifying Log File Size, Number, and Archiving Properties on page 6220](#)

## archive (Individual System Log File)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | archive <archive-sites ( <i>ftp-url</i> <password <i>password</i> >)> <files <i>number</i> > <size <i>size</i> > <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval <i>minutes</i> > <world-readable   no-world-readable>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Hierarchy Level     | [edit system <b>syslog file</b> <i>filename</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information | Statement introduced before Junos OS Release 7.4.<br><b>start-time</b> and <b>transfer-interval</b> statements introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description         | Configure archiving properties for a specific system log file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Options             | <p><b>archive-sites</b> <i>site-name</i>—FTP URL representing the destination for the archived log file (for information about how to specify valid FTP URLs, see <a href="#">“Format for Specifying Filenames and URLs in Junos OS CLI Commands” on page 73</a>). If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the filename specified at the <b>[edit system syslog]</b> hierarchy level.</p> <p><b>files</b> <i>number</i>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i>, it closes the file, compresses it, and renames it <i>logfile.0.gz</i> (the amount of data is determined by the <b>size</b> statement at this hierarchy level). The utility then opens and writes to a new file called <i>logfile</i>. When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i>, and the new file is closed, compressed, and renamed <i>logfile.0.gz</i>. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p><b>Range:</b> 1 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>password</b> <i>password</i>—Password for authenticating with the site specified by the <b>archive-sites</b> statement.</p> <p><b>size</b> <i>size</i>—Maximum amount of data that the Junos OS logging utility writes to a log file <i>logfile</i> before archiving it (closing it, compressing it, and changing its name to <i>logfile.0.gz</i>). The utility then opens and writes to a new file called <i>logfile</i>.</p> <p><b>Syntax:</b> <b>xk</b> to specify the number of kilobytes, <b>xm</b> for the number of megabytes, or <b>xg</b> for the number of gigabytes</p> <p><b>Range:</b> 64 KB through 1 GB</p> <p><b>Default:</b> 128 KB for J Series routers; 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch; 10 MB for TX Matrix and TX Matrix Plus routers</p> |

**start-time "YYYY-MM-DD.hh:mm"**—Date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

**transfer-interval *interval***—Interval at which to transfer the log file to an archive site.

**Range:** 5 through 2880 minutes

**world-readable | no-world-readable**—Grant all users permission to read archived log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

**Default:** no-world-readable

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Specifying Log File Size, Number, and Archiving Properties on page 6220](#)

## archive (QFabric System)

**Syntax** archive {  
size *size*;  
}

**Hierarchy Level** [edit system [syslog](#) file *filename*]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the archiving properties for the system message log file.

**Options** **size *size***—Maximum amount of system log message data that the QFabric system stores in the log file.

**Syntax:** *xk* to specify the number of kilobytes, *xm* for the number of megabytes, or *xg* for the number of gigabytes

**Range:** 65 KB through 1 GB

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [syslog on page 1403](#)

## console (System Logging)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>console {<br/>    <i>facility severity</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit system <a href="#">syslog</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the logging of system messages to the system console.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b><i>facility</i></b>—Class of messages to log. To specify multiple classes, include multiple <b><i>facility severity</i></b> statements. For a list of the facilities, see <a href="#">Table 607 on page 6222</a>.</p> <p><b><i>severity</i></b>—Severity of the messages that belong to the facility specified by the paired <b><i>facility</i></b> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see <a href="#">Table 608 on page 6223</a>.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Directing System Log Messages to the Console on page 6211</a></li><li>• <i>Junos OS System Log Messages Reference</i></li></ul>                                                                                                                                                                                                                                                                                                                                  |

## explicit-priority

|                                 |                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>explicit-priority;</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ],<br>[edit logical-systems <i>logical-system-name</i> system syslog host],<br>[edit system syslog file <i>filename</i> ],<br>[edit system syslog host]                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                      |
| <b>Description</b>              | Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.<br><br>When the <b>structured-data</b> statement is also included at the [edit system syslog file <i>filename</i> ] hierarchy level, this statement is ignored for the file. |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Including Priority Information in System Log Messages on page 6213</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> <li>• <a href="#">structured-data on page 6379</a></li> </ul>                                                    |

## facility-override

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>facility-override <i>facility</i>;</code>                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit system syslog host]                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                           |
| <b>Description</b>              | Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.                                                                                                     |
| <b>Options</b>                  | <i>facility</i> —Alternate facility to substitute for the default facilities. For a list of the possible facilities, see <a href="#">Table 610 on page 6224</a> .                                                        |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Changing the Alternative Facility Name for Remote System Log Messages on page 6225</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> </ul> |

## file (QFabric System)

---

**Syntax**    file *filename* {  
              archive {  
                  size *maximum-file-size*;  
              }  
              explicit-priority;  
              facility *severity*;  
              match "*regular-expression*";  
              structured-data {  
                  brief;  
              }  
          }

**Hierarchy Level**    [edit system [syslog](#)]

**Release Information**    Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Configure the logging of system messages to a file.



**NOTE:** On the QFabric system, a syslog file named **messages** is configured implicitly with facility and severity levels of **any any** and a size of 100 MBs.

---

**Options**    *facility*—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements.

*filename*—Filename that you specify with the **show log** command.

**Default:** Filename **messages**

*severity*—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities at the specified level and higher are logged.

The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**    • [syslog on page 1403](#)



## file (System Logging)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> file <i>filename</i> {     <i>facility severity</i>;     archive {         <i>files number</i>;         <i>size size</i>;         (no-world-readable   world-readable);     }     explicit-priority;     match "<i>regular-expression</i>";     structured-data {         brief;     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit system <a href="#">syslog</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the logging of system messages to a file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>facility</i></b>—Class of messages to log. To specify multiple classes, include multiple <b><i>facility severity</i></b> statements. For a list of the facilities, see <a href="#">Table 607 on page 6222</a>.</p> <p><b><i>file filename</i></b>—File in the <b><i>severity</i></b> directory in which to log messages from the specified facility. To log messages to more than one file, include more than one <b><i>file</i></b> statement.</p> <p><b><i>severity</i></b>—Severity of the messages that belong to the facility specified by the paired <b><i>facility</i></b> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see <a href="#">Table 608 on page 6223</a>.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Directing System Log Messages to a Log File on page 6209</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## files

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>files <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit system syslog archive],<br>[edit system syslog file <i>filename</i> <a href="#">archive</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i> , it closes the file, compresses it, and renames it to <i>logfile.0.gz</i> (for information about the maximum file size, see <a href="#">size</a> ). The utility then opens and writes to a new file called <i>logfile</i> . When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i> , and the new file is closed, compressed, and renamed <i>logfile.0.gz</i> . By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file). |
| <b>Options</b>                  | <i>number</i> —Maximum number of archived files.<br><b>Range:</b> 1 through 1000<br><b>Default:</b> 10 files                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying Log File Size, Number, and Archiving Properties on page 6220</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li><li>• <a href="#">size on page 6378</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## host (System)

|                                                |                                                                                                                                                                                                                                                                               |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                  | <pre> host (hostname   other-routing-engine) {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     source-address source-address;     structured-data {         brief;     } } </pre>    |
| <b>QFX Series</b>                              | <pre> host (hostname {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>                                                             |
| <b>TX Matrix Router and EX Series Switches</b> | <pre> host (hostname   other-routing-engine   scc-master) {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>                        |
| <b>TX Matrix Plus Router</b>                   | <pre> host (hostname   other-routing-engine   sfc0-master) {     facility severity;     allow-duplicates;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre> |
| <b>Hierarchy Level</b>                         | <pre> [edit logical-systems logical-system-name system syslog], [edit system syslog] </pre>                                                                                                                                                                                   |
| <b>Release Information</b>                     | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                 |
| <b>Description</b>                             | Configure the logging of system messages to a remote destination.                                                                                                                                                                                                             |

- Options**
- facility**—Class of messages to log. To specify multiple classes, include multiple **facility severity** statements. For a list of the facilities, see [Table 607 on page 6222](#).
  - hostname**—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a **host** statement for each one.
  - other-routing-engine**—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.



**NOTE:** The **other-routing-engine** option is not applicable to the QFX Series.

---

- port**—Port number of the remote syslog server that can be modified.
- scc-master**—(TX Matrix routers only) On a T640 router that is part of a routing matrix, direct messages to the TX Matrix router.
- severity**—Severity of the messages that belong to the facility specified by the paired **facility** name. Messages with severities of the specified level and higher are logged. For a list of the severities, see [Table 608 on page 6223](#).
- sfc0-master**—(TX Matrix Plus routers only) On a T1600 or T4000 router that is part of a routing matrix, direct messages to the TX Matrix Plus router.

The remaining statements are explained separately.

- Required Privilege Level**
- system—To view this statement in the configuration.
  - system-control—To add this statement to the configuration.

- Related Documentation**
- *Directing System Log Messages to a Remote Machine or the Other Routing Engine*
  - *Directing Messages to a Remote Destination from the Routing Matrix Based on the TX Matrix Router*
  - *Directing Messages to a Remote Destination from the Routing Matrix Based on a TX Matrix Plus Router*
  - *Junos OS System Log Messages Reference*

## log-prefix (System)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>log-prefix <i>string</i>;</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit system syslog host]                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Include a text string in each message directed to a remote destination.                                                                                                                        |
| <b>Options</b>                  | <i>string</i> —Text string to include in each message.                                                                                                                                         |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Adding a Text String to System Log Messages on page 6209</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> </ul> |

## match


|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>match "regular-expression";</code>                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ],<br>[edit logical-systems <i>logical-system-name</i> system syslog user ( <i>username</i>   *)],<br>[edit system syslog file <i>filename</i> ],<br>[edit system syslog host <i>hostname</i>   other-routing-engine  scc-master)],<br>[edit system syslog user ( <i>username</i>   *)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                              |
| <b>Description</b>              | Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Using Regular Expressions to Refine the Set of Logged Messages on page 6227</a></li> </ul>                                                                                                                                                                                                                             |

## size (System)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>size size;</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit system syslog archive],<br>[edit system syslog file <i>filename</i> archive]                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the maximum amount of data that the Junos OS logging utility writes to a log file <b>logfile</b> before archiving it (closing it, compressing it, and changing its name to <b>logfile.0.gz</b> ). The utility then opens and writes to a new file called <b>logfile</b> . For information about the number of archive files that the utility creates in this way, see <a href="#">files</a> . |
| <b>Options</b>                  | <b>size</b> —Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).<br><b>Syntax:</b> <b>xk</b> to specify the number of kilobytes, <b>xm</b> for the number of megabytes, or <b>xg</b> for the number of gigabytes<br><b>Range:</b> 64 KB through 1 GB<br><b>Default:</b> 1 MB for MX Series routers and the QFX Series                                           |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying Log File Size, Number, and Archiving Properties on page 6220</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li><li>• <a href="#">files on page 6374</a></li></ul>                                                                                                                                                 |

## structured-data

|                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                     | structured-data {<br>brief;<br>}                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                            | [edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ],<br>[edit system syslog file <i>filename</i> ]                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                        | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                             |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                | Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, <i>The syslog Protocol</i> ( <a href="http://tools.ietf.org/html/draft-ietf-syslog-protocol-23">http://tools.ietf.org/html/draft-ietf-syslog-protocol-23</a> ). |
| <div>  <p><b>NOTE:</b> When this statement is included, other statements that specify the format for messages written to the file are ignored (the <code>explicit-priority</code> statement at the [edit system syslog file <i>filename</i>] hierarchy level and the <code>time-format</code> statement at the [edit system syslog] hierarchy level).</p> </div> |                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                   | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li>• <i>Logging Messages in Structured-Data Format</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">explicit-priority on page 6371</a></li> <li>• <a href="#">time-format on page 6383</a></li> </ul>                           |

## syslog (System)

```
Syntax syslog {
 archive {
 (binary-data| no-binary-data);
 files number;
 size maximum-file-size;
 start-time "YYYY-MM-DD.hh:mm";
 transfer-interval minutes;
 (world-readable | no-world-readable);
 }
 console {
 facility severity;
 }
 file filename {
 facility severity;
 explicit-priority;
 match "regular-expression";
 archive {
 (binary-data| no-binary-data);
 files number;
 size maximum-file-size;
 start-time "YYYY-MM-DD.hh:mm";
 transfer-interval minutes;
 (world-readable | no-world-readable);
 }
 structured-data {
 brief;
 }
 }
 host (hostname | other-routing-engine | scc-master) {
 facility severity;
 explicit-priority;
 facility-override facility;
 log-prefix string;
 match "regular-expression";
 source-address source-address;
 structured-data {
 brief;
 }
 port port number;
 }
 log-rotate-frequency frequency;
 server server name;
 source-address source-address;
 time-format (millisecond | year | year millisecond);
 user (username | *) {
 facility severity;
 match "regular-expression";
 }
 }
```

Hierarchy Level [edit logical-systems *logical-system-name* system],  
[edit system]



|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Support at the <b>[edit logical-systems logical-system-name system]</b> hierarchy level introduced in Junos OS Release 11.4.</p>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>archive</b>—Define parameters for archiving log messages.</p> <p><b>console</b>—Send log messages of a specified class and severity to the console.</p> <p><b>file</b>—Send log messages to a named file.</p> <p><b>host</b> —Remote location to be notified of specific log messages.</p> <p><b>log-rotate-frequency</b>—Configure the interval for checking logfile size and archiving messages.</p> <p><b>server</b>—Name of the system log server in the inet.0 routing instance.</p> <p><b>source-address</b>—Include a specified address as the source address for log messages.</p> <p><b>time-format</b>—Additional information to include in the system log time stamp.</p> <p><b>user</b>—Notify a specific user of the log event.</p> |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Junos OS System Log Configuration Overview</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6155</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## syslog (QFabric System)

---

**Syntax**    `syslog {  
          file filename {  
            archive {  
              size maximum-file-size;  
            }  
            explicit-priority;  
            facility severity;  
            match "regular-expression";  
            structured-data;  
          }  
          filter all {  
            facility severity;  
            match "regular-expression";  
          }  
          host hostname {  
            explicit-priority;  
            facility severity;  
            facility-override facility;  
            log-prefix string;  
            match "regular-expression";  
            structured-data;  
          }  
          }  
          }`

**Hierarchy Level**    [edit system]

**Release Information**    Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Configure system log messages for the QFabric system.


The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding the Implementation of System Log Messages on the QFabric System on page 1227](#)

## time-format

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>time-format (year   millisecond   year millisecond);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit system syslog]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a <b>file</b>, <b>console</b>, or <b>user</b> statement at the <code>[edit system syslog]</code> hierarchy level. As of Junos OS Release 11.4, the additional time information is also sent to destinations configured by a <b>host</b> statement.</p> <p>By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, <b>Aug 21 12:36:30</b>. However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the <code>[edit system syslog time-format]</code> statement.</p> |
|                                 | <p> <b>NOTE:</b> When the <b>structured-data</b> statement is included at the <code>[edit system syslog file <i>filename</i>]</code> hierarchy level, this statement is ignored for the file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>millisecond</b>—Include the millisecond in the timestamp.</p> <p><b>year</b>—Include the year in the timestamp.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Including the Year or Millisecond in Timestamps on page 167</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> <li>• <a href="#">structured-data on page 6379</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## user (System Logging)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>user (username   *) {<br/>    facility severity;<br/>    match "regular-expression";<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system syslog]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the logging of system messages to user terminals.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>*</b> (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.</p> <p><b>facility</b>—Class of messages to log. To specify multiple classes, include multiple <b>facility severity</b> statements. For a list of the facilities, see <a href="#">Table 607 on page 6222</a>.</p> <p><b>severity</b>—Severity of the messages that belong to the facility specified by the paired <b>facility</b> name. Messages with severities the specified level and higher are logged. For a list of the severities, see <a href="#">Table 608 on page 6223</a>.</p> <p><b>username</b>—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one <b>user</b> statement.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Directing System Log Messages to a User Terminal on page 6211</a></li><li>• <a href="#">Junos OS System Logging Facilities and Message Severity Levels on page 6222</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

# Administration

- [Routine Monitoring Using the CLI on page 6385](#)
- [Monitoring Commands on page 6400](#)

## Routine Monitoring Using the CLI

---

- [Displaying a Log File from a Single-Chassis System on page 6385](#)
- [Monitoring Traffic Through the Router or Switch on page 6386](#)
- [Monitoring RMON MIB Tables on page 6389](#)
- [Monitoring SNMP on page 6389](#)
- [Monitoring System Log Messages on page 6391](#)
- [Pinging Hosts on page 6392](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 6393](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage on page 6396](#)
- [Displaying Commit Script Output on page 6398](#)

## Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system such as the QFX3500 switch, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
```

```

process): new instance detected (variable: sysAppElemRunMemory.5.6.2293)
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysAppElemRunMemory.5.8.2292)
...
Nov 4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov 4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov 4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages

```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```

user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...

```

#### Related Documentation

- [Interpreting Messages Generated in Standard Format on page 6219](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6216](#)

## Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 6386](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 6387](#)

### Displaying Real-Time Statistics About All Interfaces on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through all interfaces on the router or switch.

**Action** To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

## Sample Output

```

user@host> monitor interface traffic
host name Seconds: 15 Time: 12:31:09
Interface Link Input packets (pps) Output packets (pps)
so-1/0/0 Down 0 (0) 0 (0)
so-1/1/0 Down 0 (0) 0 (0)
so-1/1/1 Down 0 (0) 0 (0)
so-1/1/2 Down 0 (0) 0 (0)
so-1/1/3 Down 0 (0) 0 (0)
t3-1/2/0 Down 0 (0) 0 (0)
t3-1/2/1 Down 0 (0) 0 (0)
t3-1/2/2 Down 0 (0) 0 (0)
t3-1/2/3 Down 0 (0) 0 (0)
so-2/0/0 Up 211035 (1) 36778 (0)
so-2/0/1 Up 192753 (1) 36782 (0)
so-2/0/2 Up 211020 (1) 36779 (0)
so-2/0/3 Up 211029 (1) 36776 (0)
so-2/1/0 Up 189378 (1) 36349 (0)
so-2/1/1 Down 0 (0) 18747 (0)
so-2/1/2 Down 0 (0) 16078 (0)
so-2/1/3 Up 0 (0) 80338 (0)
at-2/3/0 Up 0 (0) 0 (0)
at-2/3/1 Down 0 (0) 0 (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

```

**Meaning** The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

## Displaying Real-Time Statistics About an Interface on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through an interface on the router or switch.

**Action** To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

## Sample Output

```

user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
 Input bytes: 5856541 (88 bps)
 Output bytes: 6271468 (96 bps)
 Input packets: 157629 (0 pps)
 Output packets: 157024 (0 pps)
Encapsulation statistics:
 Input keepalives: 42353
 Output keepalives: 42320
 LCP state: Opened
Error statistics:

```

```

Input errors: 0
Input drops: 0
Input framing errors: 0
Input runs: 0
Input giants: 0
Policed discards: 0
L3 incompletes: 0
L2 channel errors: 0
L2 mismatch timeouts: 0
Carrier transitions: 1
Output errors: 0
Output drops: 0
Aged packets: 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
 LOS count 1
 LOF count 1
 SEF count 1
 ES-S 77
 SES-S 77
SONET statistics:
 BIP-B1 0
 BIP-B2 0
 REI-L 0
 BIP-B3 0
 REI-P 0
Received SONET overhead: F1 : 0x00 J0 : 0xZ

```

**Meaning** The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 612 on page 6388](#).

**Table 612: Output Control Keys for the monitor interface Command**

| Action                                                                                                                                                                                                                       | Key      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Display information about the next interface. The <b>monitor interface</b> command scrolls through the physical or logical interfaces in the same order that they are displayed by the <b>show interfaces terse</b> command. | <b>N</b> |
| Display information about a different interface. The command prompts you for the name of a specific interface.                                                                                                               | <b>I</b> |
| Freeze the display, halting the display of updated statistics.                                                                                                                                                               | <b>F</b> |
| Thaw the display, resuming the display of updated statistics.                                                                                                                                                                | <b>T</b> |
| Clear (zero) the current delta counters since <b>monitor interface</b> was started. It does not clear the accumulative counter.                                                                                              | <b>C</b> |
| Stop the <b>monitor interface</b> command.                                                                                                                                                                                   | <b>Q</b> |



See the [CLI Explorer](#) for details on using match conditions with the **monitor traffic** command.

## Monitoring RMON MIB Tables

**Purpose** Monitor remote monitoring (RMON) alarm, event, and log tables.

**Action** To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index Variable description Value State

 5 monitor
 jnxOperatingCPU.9.1.0.0 5 falling threshold

Event
Index Type Last Event
 1 log and trap 2010-07-10 11:34:17 PDT
Event Index: 1
 Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
 Time: 2010-07-10 11:34:07 PDT
 Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
 Time: 2010-07-10 11:34:17 PDT
```

**Meaning** The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6198](#)
- [show snmp rmon on page 6455](#)
- [show snmp rmon history on page 6459](#)
- [clear snmp statistics on page 6403](#)
- [clear snmp history on page 6402](#)

## Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.

- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
Alarm
Index Variable description Value State

32768 Health Monitor: root file system utilization
 jnxHrStoragePercentUsed.1 58 active

32769 Health Monitor: /config file system utilization
 jnxHrStoragePercentUsed.2 0 active

32770 Health Monitor: RE 0 CPU utilization
 jnxOperatingCPU.9.1.0.0 0 active

32773 Health Monitor: RE 0 Memory utilization
 jnxOperatingBuffer.9.1.0.0 35 active

32775 Health Monitor: jkernel daemon CPU utilization
 Init daemon 0 active
 Chassis daemon 50 active
 Firewall daemon 0 active
 Interface daemon 5 active
 SNMP daemon 11 active
 MIB2 daemon 42 active
 ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system

sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx

Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 24444184
sysContact.0 = J Smith
sysName.0 = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics

SNMP statistics:
Input:
 Packets: 0, Bad versions: 0, Bad community names: 0,
 Bad community uses: 0, ASN parse errors: 0,
 Too big: 0, No such names: 0, Bad values: 0,
 Read only: 0, General errors: 0,
```

```

Total request varbinds: 0, Total set varbinds: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0,
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
Throttle drops: 0, Duplicate request drops: 0
Output:
Packets: 0, Too bigs: 0, No such names: 0,
Bad values: 0, General errors: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0

```

- Related Documentation**
- [health-monitor on page 1773](#)
  - [show snmp mib on page 6452](#)
  - [show snmp statistics on page 1859](#)

## Monitoring System Log Messages

**Purpose** Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

**Action** To view system log messages:

```
user@switch1> show log messages
```

## Sample Output

```

Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApp1ElmtRunMemory.5.6.2293)
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApp1ElmtRunMemory.5.8.2292)
...
Nov 4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages

```

**Meaning** The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of 128 kilobytes (KB).
- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user **jsmith**.

## Pinging Hosts

**Purpose** Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

**Action** To use the **ping** command to send four requests (ping count) to **host3**:  
**ping host count number**

## Sample Output

```
ping host3 count 4
user@switch> ping host3 count 4
PING host3.site.net (176.26.232.111): 56 data bytes
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

- Meaning**
- The **ping** results show the following information:
    - Size of the ping response packet (in bytes).
    - IP address of the host from which the response was sent.
    - Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
    - Time-to-live (ttl) hop-count value of the ping response packet.
    - Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
    - Number of ping requests (probes) sent to the host.
    - Number of ping responses received from the host.

- Packet loss percentage.
- Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

## Tracing SNMP Activity on a Device Running Junos OS

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
  - chassisd
  - craftd
  - ilmid
  - mib2d
  - rmopd
  - serviced
  - snmpd
- When a trace file named **filename** reaches its maximum size, it is renamed **filename.0**, then **filename.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
 file <files number> <match regular-expression> <size size> <world-readable |
 no-world-readable>;
 flag flag;
 no-remote-trace;
}
```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 6394](#)
- [Configuring Access to the Log File on page 6394](#)

- [Configuring a Regular Expression for Lines to Be Logged on page 6394](#)
- [Configuring the Trace Operations on page 6395](#)

---

### Configuring the Number and Size of SNMP Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed *filename.0*, then *filename.1*, and so on, until there are three trace files. Then the oldest trace file (*filename.2*) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

---

### Configuring Access to the Log File

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

---

### Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

## Configuring the Trace Operations

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
 all;
 configuration;
 database;
 events;
 general;
 interface-stats;
 nonvolatile-sets;
 pdu;
 policy;
 protocol-timeouts;
 routing-socket;
 server;
 subagent;
 timer;
 varbind-error;
}
```

Table 613 on page 6395 describes the meaning of the SNMP tracing flags.

**Table 613: SNMP Tracing Flags**

| Flag                     | Description                                                                 | Default Setting |
|--------------------------|-----------------------------------------------------------------------------|-----------------|
| <b>all</b>               | Log all operations.                                                         | Off             |
| <b>configuration</b>     | Log reading of the configuration at the <b>[edit snmp]</b> hierarchy level. | Off             |
| <b>database</b>          | Log events involving storage and retrieval in the events database.          | Off             |
| <b>events</b>            | Log important events.                                                       | Off             |
| <b>general</b>           | Log general events.                                                         | Off             |
| <b>interface-stats</b>   | Log physical and logical interface statistics.                              | Off             |
| <b>nonvolatile-set</b>   | Log nonvolatile SNMP set request handling.                                  | Off             |
| <b>pdu</b>               | Log SNMP request and response packets.                                      | Off             |
| <b>policy</b>            | Log policy processing.                                                      | Off             |
| <b>protocol-timeouts</b> | Log SNMP response timeouts.                                                 | Off             |
| <b>routing-socket</b>    | Log routing socket calls.                                                   | Off             |

Table 613: SNMP Tracing Flags (*continued*)

| Flag                 | Description                                                  | Default Setting |
|----------------------|--------------------------------------------------------------|-----------------|
| <b>server</b>        | Log communication with processes that are generating events. | Off             |
| <b>subagent</b>      | Log subagent restarts.                                       | Off             |
| <b>timer</b>         | Log internal timer events.                                   | Off             |
| <b>varbind-error</b> | Log variable binding errors.                                 | Off             |

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

#### Related Documentation

- *Configuring SNMP on a Device Running Junos OS*
- *Configuration Statements at the [edit snmp] Hierarchy Level*
- *Example: Tracing SNMP Activity*
- [Configuring SNMP on page 1703](#)

## Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage

Even though the Junos OS has built-in performance metrics and monitoring options, you might need to have customized performance metrics. To make it easier for you to monitor such customized data through a standard monitoring system, the Junos OS provides you with an enterprise-specific Utility MIB that can store such data and thus extend SNMP support for managing and monitoring the data of your choice.

The enterprise-specific Utility MIB provides you with container objects of the following types: **32-bit counters**, **64-bit counters**, **signed integers**, **unsigned integers**, and **octet strings**. You can use these container MIB objects to store the data that are otherwise not supported for SNMP operations. You can populate data for these objects either by using CLI commands or with the help of Op scripts and an RPC API that can invoke the CLI commands.

The following CLI commands enable you to set and clear Utility MIB object values:

- **request snmp utility-mib set instance *name* object-type <counter | counter 64 | integer | string | unsigned integer> object-value *value***
- **request snmp utility-mib clear instance *name* object-type <counter | counter 64 | integer | string | unsigned integer>**

The **instance *name*** option of the **request snmp utility-mib <set | clear>** command specifies the name of the data instance and is the main identifier of the data. The **object-type**



<counter | counter 64 | integer | string | unsigned integer> option enables you specify the object type, and the **object-value value** option enables you to set the value of the object.

To automate the process of populating Utility MIB data, you can use a combination of an event policy and event script. The following examples show the configuration for an event policy to run **show system buffers** every hour and to store the **show system buffers** data in Utility MIB objects by running an event script (**check-mbufs.slax**).

**Event Policy Configuration** To configure an event policy that runs the **show system buffers** command every hour and invokes **check-mbufs.slax** to store the **show system buffers** data into Utility MIB objects, include the following statements at the [edit] hierarchy level:

```
event-options {
 generate-event {
 1-HOUR time-interval 600;
 }
 policy Mbufs {
 events 1-HOUR;
 then {
 event-script check-mbufs.slax; # script stored at /var/db/scripts/event/
 }
 }
 event-script {
 file check-mbufs.slax;
 }
}
```

**check-mbufs.slax Script** The following example shows the **check-mbufs.slax** script that is stored under **/var/db/scripts/event/**:

```
----- script START -----
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

match / {
 <op-script-results>{
 var $cmd = <command> "show system buffers";
 var $out = jcs:invoke($cmd);

 var $lines = jcs:break_lines($out);
 for-each ($lines) {
 if (contains(., "current/peak/max")) {
 var $pattern = "([0-9]+)/([0-9]+)/([0-9]+) mbufs";
 var $split = jcs:regex($pattern, .);
 var $result = $split[2];

 var $rpc = <request-snmp-utility-mib-set> {
 <object-type> "integer";
 <instance> "current-mbufs";
 <object-value> $result;
 }
 var $res = jcs:invoke($rpc);
 }
 }
 }
}
```

```

 }
 }
}
----- script END -----

```

You can run the following command to check the data stored in the Utility MIB as a result of the event policy and script shown in the preceding examples:

```

user@host> show snmp mib walk jnxUtilData ascii jnxUtilIntegerValue."current-mbufs"
= 0 jnxUtilIntegerTime."current-mbufs" = 07 da 05 0c 03 14 2c 00 2d 07 00
regress@caramels>

```



**NOTE:** The `show snmp mib walk` command is not available on the QFabric system, but you can use external SNMP client applications to perform this operation.

#### Related Documentation

- *Understanding SNMP Implementation in Junos OS*
- *Configuring SNMP on Devices Running Junos OS*
- *Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS*
- *Optimizing the Network Management System Configuration for the Best Results*
- *Configuring Options on Managed Devices for Better SNMP Response Time*
- *Managing Traps and Informs*
- [Understanding the Implementation of SNMP on the QFabric System on page 1225](#)

## Displaying Commit Script Output

[Table 614 on page 6398](#) summarizes the Junos OS command-line interface (CLI) commands you can use to monitor and troubleshoot commit scripts. For more information about the `cscript.log` file, see *Tracing Commit Script Processing*.



**NOTE:** Tracing commit script processing, including the `cscript.log` file, is not supported on the QFX3000-G QFabric system.

**Table 614: Commit Script Configuration and Operational Mode Commands**

| Task                                                          | Command                                          |
|---------------------------------------------------------------|--------------------------------------------------|
| <b>Configuration Mode Commands</b>                            |                                                  |
| Display errors and warnings generated by commit scripts.      | <code>commit</code> or <code>commit check</code> |
| Display detailed information.                                 | <code>commit   display detail</code>             |
| Display the underlying Extensible Markup Language (XML) data. | <code>commit   display xml</code>                |

**Table 614: Commit Script Configuration and Operational Mode Commands** (*continued*)

| Task                                                                                                                                                                                                              | Command                                                                          |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Display the postinheritance contents of the configuration database. This view includes transient changes, but does not include changes made in configuration groups.                                              | <b>show   display commit-scripts</b>                                             |
| Display the postinheritance contents of the configuration database. This view excludes transient changes.                                                                                                         | <b>show   display commit-scripts no-transients</b>                               |
| Display the postinheritance configuration in XML format.<br><br>Viewing the configuration in XML format can be helpful when you are writing XML Path Language (XPath) expressions and configuration element tags. | <b>show   display commit-scripts view</b>                                        |
| Display the postinheritance configuration in XML format, but exclude transient changes.                                                                                                                           | <b>show   display commit-scripts view   display commit-scripts no-transients</b> |
| Display all configuration groups data, including script-generated changes to the groups.                                                                                                                          | <b>show groups   display commit-scripts</b>                                      |
| Display a particular configuration group, including script-generated changes to the group.                                                                                                                        | <b>show groups <i>group-name</i>   display commit-scripts</b>                    |
| <b>Operational Mode Commands</b>                                                                                                                                                                                  |                                                                                  |
| Display logging data associated with all commit script processing.                                                                                                                                                | <b>show log cscript.log</b>                                                      |
| Display processing for only the most recent commit operation.                                                                                                                                                     | <b>show log cscript.log   last</b>                                               |
| Display processing for script errors.                                                                                                                                                                             | <b>show log cscript.log   match error</b>                                        |
| Display processing for a particular script.                                                                                                                                                                       | <b>show log cscript.log   match <i>filename</i></b>                              |

**Related Documentation**

- *Tracing Commit Script Processing*

## Monitoring Commands

---

- clear sflow collector statistics
- clear snmp history
- clear snmp statistics
- monitor traffic
- ping
- ping fabric multicast-flow
- ping fabric unicast-flow
- request snmp spoof-trap
- request snmp utility-mib clear instance
- request snmp utility-mib set instance
- show oam fabric flow-specification
- show oam fabric interfaces
- show log
- show sflow
- show sflow collector
- show sflow interface
- show snmp health-monitor
- show snmp inform-statistics
- show snmp mib
- show snmp rmon
- show snmp rmon history
- show snmp statistics
- show snmp v3
- traceroute fabric unicast-flow

## clear sflow collector statistics

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear sflow collector statistics                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                               |
| <b>Description</b>              | Clear the sample counters for all sFlow collectors.                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> <li>• <a href="#">show sflow collector on page 6442</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">clear sflow collector statistics on page 6401</a>                                                                                                                                                                                                                 |

### Sample Output

#### clear sflow collector statistics

The following example shows two output examples for the **show sflow collector** command, one before and one after the **clear sflow collector statistics** command was issued.

```
user@host> show sflow collector
Collector Udp-port No. of samples
address
10.1.1.1 6343 3174
10.1.2.1 6343 3562
```

```
user@host> clear sflow collector statistics
```

```
user@host> show sflow collector
Collector Udp-port No. of samples
address
10.1.1.1 6343 0
10.1.2.1 6343 0
```

## clear snmp history

---

**Syntax**    `clear snmp history (index | all)`

**Release Information**    Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Delete the samples of Ethernet statistics collected for a history group.

**Options**    *all*—Clear all the entries in the history index.

*index*—Clear the contents of the specified entry in the history index.

**Required Privilege Level**    clear

**Related Documentation**    • [clear snmp statistics on page 6403](#)

## clear snmp statistics

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear snmp statistics                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Clear Simple Network Management Protocol (SNMP) statistics.                                                                                                                              |
| <b>Options</b>                  | This command has no options.                                                                                                                                                             |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show snmp statistics on page 1859</a></li> </ul>                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">clear snmp statistics on page 6403</a>                                                                                                                                       |
| <b>Output Fields</b>            | See <a href="#">show snmp statistics</a> for an explanation of output fields.                                                                                                            |

## Sample Output

### clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
 Input:
 Packets: 8, Bad versions: 0, Bad community names: 0,
 Bad community uses: 0, ASN parse errors: 0,
 Too bigs: 0, No such names: 0, Bad values: 0,
 Read onlys: 0, General errors: 0,
 Total request varbinds: 8, Total set varbinds: 0,
 Get requests: 0, Get nexts: 8, Set requests: 0,
 Get responses: 0, Traps: 0,
 Silent drops: 0, Proxy drops 0
 Output:
 Packets: 2298, Too bigs: 0, No such names: 0,
 Bad values: 0, General errors: 0,
 Get requests: 0, Get nexts: 0, Set requests: 0,
 Get responses: 8, Traps: 2290
```

```
user@host> clear snmp statistics
```

```
user@host> show snmp statistics
SNMP statistics:
 Input:
 Packets: 0, Bad versions: 0, Bad community names: 0,
 Bad community uses: 0, ASN parse errors: 0,
 Too bigs: 0, No such names: 0, Bad values: 0,
 Read onlys: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0,
Silent drops: 0, Proxy drops 0
Output:
Packets: 0, Too bigs: 0, No such names: 0,
Bad values: 0, General errors: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0
```



## monitor traffic

**Syntax**    monitor traffic  
               <brief | detail | extensive>  
               <absolute-sequence>  
               <count *count*>  
               <interface *interface-name*>  
               <layer2-headers>  
               <matching *matching*>  
               <no-domain-names>  
               <no-promiscuous>  
               <no-resolve>  
               <no-timestamp>  
               <print-ascii>  
               <print-hex>  
               <resolve-timeout>  
               <size *size*>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                               Command introduced in Junos OS Release 9.0 for EX Series switches.  
                               Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display packet headers or packets received and sent from the Routing Engine.



### NOTE:

- Using the **monitor-traffic** command can degrade router or switch performance.
- Delays from DNS resolution can be eliminated by using the **no-resolve** option.



**NOTE:** This command is not supported on the QFabric system.

**Options**    **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**absolute-sequence**—(Optional) Display absolute TCP sequence numbers.

**count *count***—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The **monitor traffic** command quits automatically after displaying the number of packets specified.

**interface *interface-name***—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

**layer2-headers**—(Optional) Display the link-level header on each line.

**matching *matching***—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

**no-domain-names**—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

**no-promiscuous**—(Optional) Do not put the interface into promiscuous mode.

**no-resolve**—(Optional) Suppress reverse lookup of the IP addresses.

**no-timestamp**—(Optional) Suppress timestamps on displayed packets.

**print-ascii**—(Optional) Display each packet in ASCII format.

**print-hex**—(Optional) Display each packet, except the link-level header, in hexadecimal format.

**resolve-timeout *timeout***—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

**size *size***—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

**Additional Information** In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

**monitor traffic matching "*expression*"**

Replace ***expression*** with one or more of the match conditions listed in [Table 615 on page 6407](#).

Table 615: Match Conditions for the monitor traffic Command

| Match Type    | Condition                                             | Description                                                                                                                                                                                                                                                                            |
|---------------|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Entity        | <b>host</b> [ <i>address</i>   <i>hostname</i> ]      | Matches packets that contain the specified address or hostname.<br><br>The protocol match conditions <b>arp</b> , <b>ip</b> , or <b>rarp</b> , or any of the directional match conditions can be prepended to the <b>host</b> match condition.                                         |
|               | <b>net</b> <i>address</i>                             | Matches packets with source or destination addresses containing the specified network address.                                                                                                                                                                                         |
|               | <b>net</b> <i>address mask mask</i>                   | Matches packets containing the specified network address and subnet mask.                                                                                                                                                                                                              |
|               | <b>port</b> ( <i>port-number</i>   <i>port-name</i> ) | Matches packets containing the specified source or destination TCP or UDP port number or port name.<br><br>In place of the numeric port address, you can specify a text synonym, such as <b>bgp</b> (179), <b>dhcp</b> (67), or <b>domain</b> (53) (the port numbers are also listed). |
| Directional   | <b>dst</b>                                            | Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                  |
|               | <b>src</b>                                            | Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                             |
|               | <b>src and dst</b>                                    | Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                         |
|               | <b>src or dst</b>                                     | Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                        |
| Packet Length | <b>less</b> <i>value</i>                              | Matches packets shorter than or equal to the specified value, in bytes.                                                                                                                                                                                                                |
|               | <b>greater</b> <i>value</i>                           | Matches packets longer than or equal to the specified value, in bytes.                                                                                                                                                                                                                 |

Table 615: Match Conditions for the monitor traffic Command (*continued*)

| Match Type | Condition                                                | Description                                                                                                                                                                                                                                                                                                                     |
|------------|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Protocol   | <b>amt</b>                                               | Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.                                                                                                                                                                                                |
|            | <b>arp</b>                                               | Matches all ARP packets.                                                                                                                                                                                                                                                                                                        |
|            | <b>ether</b>                                             | Matches all Ethernet packets.                                                                                                                                                                                                                                                                                                   |
|            | <b>ether (broadcast   multicast)</b>                     | Matches broadcast or multicast Ethernet frames. This match condition can be prepended with <b>src</b> and <b>dst</b> .                                                                                                                                                                                                          |
|            | <b>ether protocol (address   (arp   ip   rarp))</b>      | Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The <b>ether protocol</b> arguments <b>arp</b> , <b>ip</b> , and <b>rarp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ether protocol</b> match condition. |
|            | <b>icmp</b>                                              | Matches all ICMP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>ip</b>                                                | Matches all IP packets.                                                                                                                                                                                                                                                                                                         |
|            | <b>ip (broadcast   multicast)</b>                        | Matches broadcast or multicast IP packets.                                                                                                                                                                                                                                                                                      |
|            | <b>ip protocol (address   (icmp   igmp   tcp   udp))</b> | Matches packets with the specified address or protocol type. The <b>ip protocol</b> arguments <b>icmp</b> , <b>tcp</b> , and <b>udp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ip protocol</b> match condition.                                                 |
|            | <b>isis</b>                                              | Matches all IS-IS routing messages.                                                                                                                                                                                                                                                                                             |
|            | <b>rarp</b>                                              | Matches all RARP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>tcp</b>                                               | Matches all TCP datagrams.                                                                                                                                                                                                                                                                                                      |
|            | <b>udp</b>                                               | Matches all UDP datagrams.                                                                                                                                                                                                                                                                                                      |

To combine expressions, use the logical operators listed in [Table 616 on page 6408](#).

Table 616: Logical Operators for the monitor traffic Command

| Logical Operator (Highest to Lowest Precedence) | Description                                                                          |
|-------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>!</b>                                        | Logical NOT. If the first condition does not match, the next condition is evaluated. |

Table 616: Logical Operators for the monitor traffic Command (*continued*)

| Logical Operator (Highest to Lowest Precedence) | Description                                                                                                                                         |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| &&                                              | Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped. |
|                                                 | Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.  |
| ( )                                             | Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).        |

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 617 on page 6410](#).



**NOTE:** Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 615 on page 6407](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

**Table 617: Arithmetic and Relational Operators for the monitor traffic Command**

| Arithmetic or Relational Operator                         | Description                                                                         |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------|
| <b>Arithmetic Operator</b>                                |                                                                                     |
| +                                                         | Addition operator.                                                                  |
| -                                                         | Subtraction operator.                                                               |
| /                                                         | Division operator.                                                                  |
| &                                                         | Bitwise AND.                                                                        |
| *                                                         | Bitwise exclusive OR.                                                               |
|                                                           | Bitwise inclusive OR.                                                               |
| <b>Relational Operator (Highest to Lowest Precedence)</b> |                                                                                     |
| <=                                                        | If the first expression is less than or equal to the second, the packet matches.    |
| >=                                                        | If the first expression is greater than or equal to the second, the packet matches. |
| <                                                         | If the first expression is less than the second, the packet matches.                |
| >                                                         | If the first expression is greater than the second, the packet matches.             |
| =                                                         | If the compared expressions are equal, the packet matches.                          |
| !=                                                        | If the compared expressions are unequal, the packet matches.                        |

**Required Privilege Level** trace  
maintenance

**List of Sample Output** [monitor traffic count on page 6411](#)  
[monitor traffic detail count on page 6411](#)  
[monitor traffic extensive \(Absolute Sequence\) on page 6411](#)  
[monitor traffic extensive \(Relative Sequence\) on page 6411](#)  
[monitor traffic extensive count on page 6411](#)  
[monitor traffic interface on page 6412](#)  
[monitor traffic matching on page 6412](#)  
[monitor traffic \(TX Matrix Plus Router\) on page 6412](#)  
[monitor traffic \(QFX3500 Switch\) on page 6413](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

### monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

### monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive count

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
```

```
reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)
```

### monitor traffic interface

```
user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...
```

### monitor traffic matching

```
user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...
```

### monitor traffic (TX Matrix Plus Router)

```
user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >
```



```

summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell11.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

### monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.

```

```
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```


## ping

**List of Syntax**   [Syntax on page 6415](#)  
                               [Syntax \(QFX Series\) on page 6415](#)

**Syntax**   `ping host`  
                   `<bypass-routing>`  
                   `<count requests>`  
                   `<detail>`  
                   `<do-not-fragment>`  
                   `<inet | inet6>`  
                   `<interface source-interface>`  
                   `<interval seconds>`  
                   `<logical-system logical-system-name>`  
                   `<loose-source value>`  
                   `<mac-address mac-address>`  
                   `<no-resolve>`  
                   `<pattern string>`  
                   `<rapid>`  
                   `<record-route>`  
                   `<routing-instance routing-instance-name>`  
                   `<size bytes>`  
                   `<source source-address>`  
                   `<strict >`  
                   `<strict-source value.>`  
                   `<tos type-of-service>`  
                   `<ttl value>`  
                   `<verbose>`  
                   `<vpls instance-name>`  
                   `<wait seconds>`

**Syntax (QFX Series)**   `ping host`  
                               `<bypass-routing>`  
                               `<count requests>`  
                               `<detail>`  
                               `<do-not-fragment>`  
                               `<inet>`  
                               `<interface source-interface>`  
                               `<interval seconds>`  
                               `<logical-system logical-system-name>`  
                               `<loose-source value>`  
                               `<mac-address mac-address>`  
                               `<no-resolve>`  
                               `<pattern string>`  
                               `<rapid>`  
                               `<record-route>`  
                               `<routing-instance routing-instance-name>`  
                               `<size bytes>`  
                               `<source source-address>`  
                               `<strict>`  
                               `< strict-source value>`  
                               `<tos type-of-service>`  
                               `<ttl value>`  
                               `<verbose>`

<wait *seconds*>

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b> | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>         | Check host reachability and network connectivity. The <b>ping</b> command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>host</b>—IP address or hostname of the remote system to ping.</p> <p><b>bypass-routing</b>—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p><b>count requests</b>—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p><b>detail</b>—(Optional) Include in the output the interface on which the ping reply was received.</p> <p><b>do-not-fragment</b>—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div><p><b>NOTE:</b> In Junos OS Release 11.1 and later, when issuing the <b>ping</b> command for an IPv6 route with the <b>do-not-fragment</b> option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p></div> <p><b>inet</b>—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p><b>inet6</b>—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p><b>interface source-interface</b>—(Optional) Interface to use to send the ping requests.</p> <p><b>interval seconds</b>—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p><b>logical-system logical-system-name</b>—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the <b>set cli logical-system logical-system-name</b> command and then run the <b>ping</b> command. To return to the main router, enter the <b>clear cli logical-system</b> command.</p> |

**loose-source *value***—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

**mac-address *mac-address***—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**pattern *string***—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

**rapid**—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

**record-route**—(Optional) Record and report the packet's path (IPv4).

**routing-instance *routing-instance-name***—(Optional) Name of the routing instance for the ping attempt.

**size *bytes***—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**strict**—(Optional) Use the strict source route option (IPv4).

**strict-source *value***—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

**tos *type-of-service***—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

**ttl *value***—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

**verbose**—(Optional) Display detailed output.

**vpls *instance-name***—(Optional) Ping the instance to which this VPLS belongs.

**wait *seconds***—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

**Required Privilege Level** network

**Related Documentation**

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

**List of Sample Output**   [ping hostname on page 6418](#)  
                              [ping hostname rapid on page 6418](#)  
                              [ping hostname size count on page 6418](#)

**Output Fields**   When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

### ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

### ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

## ping fabric multicast-flow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ping fabric multicast-flow source-fmep-id <i>source-fmep-id</i> fma-name <i>fma-name</i><br/>flow-spec-name <i>flow-specification-name</i><br/>&lt;verbose&gt;</code>                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Send a fabric multicast ping command to the multicast group destination fabric maintenance endpoints (FMEPs) configured in the flow specification. Send a single PDU with a timeout after 1 second.                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>flow-spec-name <i>flow-specification-name</i></b>—Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.</p> <p><b>fma-name <i>fma-name</i></b>—Name of the FMA.</p> <p><b>source-fmep-id <i>source-fmep-id</i></b>—FMEP ID that is the source of the fabric ping operation.</p> <p><b>verbose</b>—(Optional) Detailed version of the output display.</p>                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> <li>• <a href="#">ping fabric unicast-flow on page 6421</a></li> <li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">ping fabric multicast-flow on page 6420</a>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | <a href="#">Table 618 on page 6419</a> lists the output fields for the <b>ping fabric-multicast-flow</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                         |

**Table 618: ping fabric multicast-flow Output Fields**

| Field Name                      | Field Description                              |
|---------------------------------|------------------------------------------------|
| Using fabric flow-specification | Name of the flow specification.                |
| Ethernet frame-size             | Ethernet frame size in bytes.                  |
| Source-IP                       | MAC address of the source FMEP interface.      |
| Destination-IP                  | MAC address of the destination FMEP interface. |

Table 618: ping fabric multicast-flow Output Fields (*continued*)

| Field Name                     | Field Description                          |
|--------------------------------|--------------------------------------------|
| Protocol                       | IP protocol.                               |
| received response from fmep-id | FMEP ID from which responses are received. |

## Sample Output

### ping fabric multicast-flow

```
user@host> ping fabric multicast-flow source-fmep-id 3 fma-name fma1 flow-spec-name fspec2
verbose
```

```
Using fabric flow-specification: fspec2
```

```
Ethernet frame-size: 256
```

```
Source-IP: 87.238.205.21
```

```
Destination-IP: 225.0.0.1 IP
```

```
Protocol: 17
```

```
received response from fmep-id 32775...
```

```
received response from fmep-id 32794...
```

```
received response from fmep-id 32795...
```

```
sent 1 requests, received 3 responses
```



## ping fabric unicast-flow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ping fabric unicast-flow <i>fma-name</i> <i>fma-name</i> source-fmep-id <i>source-fmep-id</i> dest-fmep-id <i>dest-fmep-id</i> flow-spec-name <i>flow-specification-name</i><br><count <i>count</i> ><br><forced-fte-interface <i>fte-interface</i> ><br><verbose>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Send a unicast fabric ping command to the destination fabric maintenance endpoints (FMEPs) specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>count</b> <i>count</i>—(Optional) Number of fabric ping PDUs to send.<br/>The maximum PDU count is 5. The default PDU count is 1.</p> <p><b>dest-fmep-id</b> <i>destination-fmep-id</i>—ID of the destination FMEP of the fabric ping operation.</p> <p><b>flow-spec-name</b> <i>flow-specification-name</i>—Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.</p> <p><b>fma-name</b> <i>fma-name</i>—Name of the FMA.</p> <p><b>forced-fte-interface</b> <i>fte-interface</i>—(Optional) Forces the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.</p> <p><b>source-fmep-id</b> <i>source-fmep-id</i>—ID of the FMEP that is the source of the fabric ping operation.</p> <p><b>verbose</b>—(Optional) Detailed version of the output display.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> <li>• <a href="#">ping fabric multicast-flow on page 6419</a></li> <li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <p><a href="#">ping fabric unicast-flow on page 6422</a></p> <p><a href="#">ping fabric unicast-flow (With Masking Enabled on the Flow Specification) on page 6422</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 619 on page 6422</a> lists the output fields for the <b>ping fabric-unicast-flow</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 619: ping fabric unicast-flow Output Fields

| Field Name                                                                                       | Field Description                                                       |
|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Fabric flow ping between source <i>source-fmep-name</i> destination <i>destination-fmep-name</i> | Path of the unicast-flow ping between the source and destination FMEPs. |
| Using fabric flow-specification                                                                  | Name of the flow specification.                                         |
| Ethernet frame-size                                                                              | Ethernet frame size in bytes.                                           |
| Source-MAC                                                                                       | MAC address of the source FMEP interface.                               |
| Destination-MAC                                                                                  | MAC address of the destination FMEP interface.                          |
| Ethertype                                                                                        | EtherType protocol.                                                     |
| received response from fmep-id                                                                   | ID of the FMEP sending the response.                                    |

## Sample Output

### ping fabric unicast-flow

```
user@host> ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1 verbose
```

```
Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:2:B1:61:3B:96
Destination-MAC: 0:26:78:90:70:21
EtherType: 8295
received response from fmep-id 2...
sent 1 requests, received 1 responses
```

### ping fabric unicast-flow (With Masking Enabled on the Flow Specification)

```
user@host> ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1 verbose
```

```
Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:1:20:A0:0:1
Destination-MAC: 0:8A:3A:B3:7D:67
Source-MAC Mask: FF:F8:FF:FF:FF:FF
EtherType: 26785
Current Source-MAC value: 0:1:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:2:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:3:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:4:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:5:20:a0:0:1
received response from fmep-id 2...
```

```
Current Source-MAC value: 0:6:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:7:20:a0:0:1
received response from fmep-id 2...
sent 7 requests, received 7 responses
```

## request snmp spoof-trap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request snmp spoof-trap</b><br><b>&lt;trap&gt; variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.2.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>&lt;trap&gt;</b> —Name of the trap to spoof.<br><br><b>variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b> —(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, <b>ifIndex[14] = 14</b> ). Enclose the list of variable bindings in quotation marks ( " ") and use a comma to separate each object name, instance, and value definition (for example, <b>variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"</b> ). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.<br><br><b>&lt;dummy name&gt;</b> —A dummy trap name to display the list of available traps.<br><br><b>Question mark (?)</b> —Question mark? to display possible completions. |
| <b>Required Privilege Level</b> | request                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">request snmp spoof-trap (with Variable Bindings) on page 6424</a><br><a href="#">request snmp spoof-trap (Illegal Trap Name) on page 6424</a><br><a href="#">request snmp spoof-trap (Question Mark ?) on page 6428</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmpspoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

### request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
```

ads1AturPerfLossThreshTrap  
ads1AturPerfLossThreshTrap  
ads1AturPerfLprsThreshTrap  
ads1AturRateChangeTrap  
apsEventChannelMismatch  
apsEventFEPLF  
apsEventModeMismatch  
apsEventPSBF  
apsEventSwitchover  
authenticationFailure  
bfdSessDown  
bfdSessUp  
bgpBackwardTransition  
bgpEstablished  
coldStart  
dlswTrapCircuitDown  
dlswTrapCircuitUp  
dlswTrapTConnDown  
dlswTrapTConnPartnerReject  
dlswTrapTConnProtViolation  
dlswTrapTConnUp  
dsx1LineStatusChange  
dsx3LineStatusChange  
entConfigChange  
fallingAlarm  
frDLCIStatusChange  
ggsnTrapChanged  
ggsnTrapCleared  
ggsnTrapNew  
gmplsTunnelDown  
ifMauJabberTrap  
ipv6IfStateChange  
isisAreaMismatch  
isisAttemptToExceedMaxSequence  
isisAuthenticationFailure  
isisAuthenticationTypeFailure  
isisCorruptedLSPDetected  
isisDatabaseOverload  
isisIDLenMismatch  
isisLSPTooLargeToPropagate  
isisManualAddressDrops  
isisMaxAreaAddressesMismatch  
isisOriginatingLSPBufferSizeMismatch  
isisOwnLSPPurge  
isisProtocolsSupportedMismatch  
isisRejectedAdjacency  
isisSequenceNumberSkip  
isisVersionSkew  
jnxAccessAuthServerDisabled  
jnxAccessAuthServerEnabled  
jnxAccessAuthServiceDown  
jnxAccessAuthServiceUp  
jnxBfdSessDetectionTimeHigh  
jnxBfdSessTxIntervalHigh  
jnxBgpM2BackwardTransition  
jnxBgpM2Established  
jnxCmCfgChange  
jnxCmRescueChange  
jnxCollFlowOverload  
jnxCollFlowOverloadCleared  
jnxCollFtpSwitchover

jnxCollMemoryAvailable  
jnxCollMemoryUnavailable  
jnxCollUnavailableDest  
jnxCollUnavailableDestCleared  
jnxCollUnsuccessfulTransfer  
jnxDfcHardMemThresholdExceeded  
jnxDfcHardMemUnderThreshold  
jnxDfcHardPpsThresholdExceeded  
jnxDfcHardPpsUnderThreshold  
jnxDfcSoftMemThresholdExceeded  
jnxDfcSoftMemUnderThreshold  
jnxDfcSoftPpsThresholdExceeded  
jnxDfcSoftPpsUnderThreshold  
jnxEventTrap  
jnxExampleStartup  
jnxFEBSwitchover  
jnxFanFailure  
jnxFanOK  
jnxFruCheck  
jnxFruFailed  
jnxFruInsertion  
jnxFruOK  
jnxFruOffline  
jnxFruOnline  
jnxFruPowerOff  
jnxFruPowerOn  
jnxFruRemoval  
jnxHardDiskFailed  
jnxHardDiskMissing  
jnxJsAvPatternUpdateTrap  
jnxJsChassisClusterSwitchover  
jnxJsFwAuthCapacityExceeded  
jnxJsFwAuthFailure  
jnxJsFwAuthServiceDown  
jnxJsFwAuthServiceUp  
jnxJsNatAddrPoolThresholdStatus  
jnxJsScreenAttack  
jnxJsScreenCfgChange  
jnxLdpLspDown  
jnxLdpLspUp  
jnxLdpSesDown  
jnxLdpSesUp  
jnxMIMstCistPortLoopProtectStateChangeTrap  
jnxMIMstCistPortRootProtectStateChangeTrap  
jnxMIMstErrTrap  
jnxMIMstGenTrap  
jnxMIMstInvalidBpduRxdTrap  
jnxMIMstMstiPortLoopProtectStateChangeTrap  
jnxMIMstMstiPortRootProtectStateChangeTrap  
jnxMIMstNewRootTrap  
jnxMIMstProtocolMigrationTrap  
jnxMIMstRegionConfigChangeTrap  
jnxMIMstTopologyChgTrap  
jnxMacChangedNotification  
jnxMplsLdpInitSesThresholdExceeded  
jnxMplsLdpPathVectorLimitMismatch  
jnxMplsLdpSessionDown  
jnxMplsLdpSessionUp  
jnxOspfV3IfConfigError  
jnxOspfV3IfRxBadPacket  
jnxOspfV3IfStateChange

jnxOspfV3LsdbApproachingOverflow  
jnxOspfV3LsdbOverflow  
jnxOspfV3NbrRestartHelperStatusChange  
jnxOspfV3NbrStateChange  
jnxOspfV3NssaTranslatorStatusChange  
jnxOspfV3RestartStatusChange  
jnxOspfV3VirtIfConfigError  
jnxOspfV3VirtIfRxBadPacket  
jnxOspfV3VirtIfStateChange  
jnxOspfV3VirtNbrRestartHelperStatusChange  
jnxOspfV3VirtNbrStateChange  
jnxOtnAlarmCleared  
jnxOtnAlarmSet  
jnxOverTemperature  
jnxPMonOverloadCleared  
jnxPMonOverloadSet  
jnxPingEgressJitterThresholdExceeded  
jnxPingEgressStdDevThresholdExceeded  
jnxPingEgressThresholdExceeded  
jnxPingIngressJitterThresholdExceeded  
jnxPingIngressStdDevThresholdExceeded  
jnxPingIngressThresholdExceeded  
jnxPingRttJitterThresholdExceeded  
jnxPingRttStdDevThresholdExceeded  
jnxPingRttThresholdExceeded  
jnxPortBpduErrorStatusChangeTrap  
jnxPortLoopProtectStateChangeTrap  
jnxPortRootProtectStateChangeTrap  
jnxPowerSupplyFailure  
jnxPowerSupplyOK  
jnxRedundancySwitchover  
jnxRmonAlarmGetFailure  
jnxRmonGetOk  
jnxSecAccessIfMacLimitExceeded  
jnxSecAccessSdsRateLimitCrossed  
jnxSonetAlarmCleared  
jnxSonetAlarmSet  
jnxSpSvcSetCpuExceeded  
jnxSpSvcSetCpuOk  
jnxSpSvcSetZoneEntered  
jnxSpSvcSetZoneExited  
jnxStormEventNotification  
jnxSyslogTrap  
jnxTemperatureOK  
jnxVccpPortDown  
jnxVccpPortUp  
jnxVpnIfDown  
jnxVpnIfUp  
jnxVpnPwDown  
jnxVpnPwUp  
jnxl2aldGlobalMacLimit  
jnxl2aldInterfaceMacLimit  
jnxl2aldRoutingInstMacLimit  
linkDown  
linkUp  
lldpRemTablesChange  
mfrMibTrapBundleLinkMismatch  
mplsLspChange  
mplsLspDown  
mplsLspInfoChange  
mplsLspInfoDown

mplsLspInfoPathDown  
mplsLspInfoPathUp  
mplsLspInfoUp  
mplsLspPathDown  
mplsLspPathUp  
mplsLspUp  
mplsNumVrfRouteMaxThreshExceeded  
mplsNumVrfRouteMidThreshExceeded  
mplsNumVrfSecIllglLb1ThrshExcd  
mplsTunnelDown  
mplsTunnelReoptimized  
mplsTunnelRerouted  
mplsTunnelUp  
mplsVrfIfDown  
mplsVrfIfUp  
mplsXCDown  
mplsXCUp  
msdpBackwardTransition  
msdpEstablished  
newRoot  
ospfIfAuthFailure  
ospfIfConfigError  
ospfIfRxBadPacket  
ospfIfStateChange  
ospfLsdbApproachingOverflow  
ospfLsdbOverflow  
ospfMaxAgeLsa  
ospfNbrStateChange  
ospfOriginateLsa  
ospfTxRetransmit  
ospfVirtIfAuthFailure  
ospfVirtIfConfigError  
ospfVirtIfRxBadPacket  
ospfVirtIfStateChange  
ospfVirtIfTxRetransmit  
ospfVirtNbrStateChange  
pethMainPowerUsageOffNotification  
pethMainPowerUsageOnNotification  
pethPsePortOnOffNotification  
pingProbeFailed  
pingTestCompleted  
pingTestFailed  
ptopoConfigChange  
risingAlarm  
rpMauJabberTrap  
sd1cLSStatusChange  
sd1cPortStatusChange  
topologyChange  
traceRoutePathChange  
traceRouteTestCompleted  
traceRouteTestFailed  
vrrpTrapAuthFailure  
vrrpTrapNewMaster  
warmStart

#### request snmp spoof-trap (Question Mark ?)

```
user@host> request snmp spoof-trap ?
Possible completions:
<trap> The name of the trap to spoof
ads1AtucInitFailureTrap
```



```
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
---(more 10%)---
```

## request snmp utility-mib clear instance

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request snmp utility-mib clear instance <i>name</i></code><br><code>object-type <i>type</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Clear the data stored in the specified container object in the SNMP Utility MIB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>name</i></b>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><b><i>object-type type</i></b>—Type of container object in which the data is stored. The following container object types are supported:</p> <ul style="list-style-type: none"><li>• <b>counter</b>—Stores a 32-bit counter value.</li><li>• <b>counter64</b>—Stores a 64-bit counter value.</li><li>• <b>integer</b>—Stores a 32-bit signed integer value.</li><li>• <b>unsigned-integer</b>—Stores a 32-bit unsigned integer value.</li></ul> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">SNMP MIBs Support on page 6124</a></li><li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 1225</a></li><li>• <a href="#">request snmp utility-mib set instance on page 6431</a></li></ul>                                                                                                                                                                                                                                                                                      |

## request snmp utility-mib set instance

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | request snmp utility-mib set instance <i>name</i><br>object-type <i>type</i><br>object-value <i>value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Store data in the specified container object in the SNMP Utility MIB. The data may be retrieved by SNMP operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b><i>name</i></b>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><b>object-type <i>type</i></b>—Type of container object in which to store data. The following container object types are supported:</p> <ul style="list-style-type: none"> <li>• <b>counter</b>—Stores a 32-bit counter value.</li> <li>• <b>counter64</b>—Stores a 64-bit counter value.</li> <li>• <b>integer</b>—Stores a 32-bit signed integer value.</li> <li>• <b>unsigned-integer</b>—Stores a 32-bit unsigned integer value.</li> <li>• <b>string</b>—Stores an octet string value.</li> </ul> <p><b>object-value <i>value</i></b>—Data that is stored in the container object.</p> |
| <b>Required Privilege Level</b> | request                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6124</a></li> <li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 1225</a></li> <li>• <a href="#">request snmp utility-mib clear instance on page 6430</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                            |

## show oam fabric flow-specification

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show oam fabric flow-specification</code><br><code>&lt;flow-specification-name&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the fabric flow specifications that are configured in a QFabric system Node group.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>flow-specification-name</i> —(Optional) Name of a flow specification.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> <li>• <a href="#">ping fabric multicast-flow on page 6419</a></li> <li>• <a href="#">ping fabric unicast-flow on page 6421</a></li> <li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam fabric flow-specification (All) on page 6433</a><br><a href="#">show oam fabric flow-specification (Flow Specification Specified) on page 6434</a>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 620 on page 6432</a> lists the output fields for the <b>show oam fabric flow-specification</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                    |

Table 620: show oam fabric flow-specification Output Fields

| Field Name              | Field Description                                                                                                                                                                |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Flow specification name | Name of the flow specification.                                                                                                                                                  |
| Type                    | Flow specification type: <ul style="list-style-type: none"> <li>• Ethernet</li> <li>• Ethernet Unicast IPV4</li> <li>• Multicast IPV4</li> <li>• Multicast VLAN flood</li> </ul> |
| Ethernet frame size     | Ethernet frame size in bytes.                                                                                                                                                    |
| Ether type              | EtherType protocol.                                                                                                                                                              |
| Source-MAC              | MAC address of the source FMEP interface.                                                                                                                                        |

Table 620: show oam fabric flow-specification Output Fields (*continued*)

| Field Name                     | Field Description                                   |
|--------------------------------|-----------------------------------------------------|
| Source-MAC Mask                | MAC address mask of the source FMEP interface.      |
| Destination-MAC                | MAC address of the destination FMEP interface.      |
| Destination-MAC Mask           | MAC address mask of the destination FMEP interface. |
| Source-IP                      | IP address of the source FMEP interface.            |
| Source-IP Mask                 | IP address mask of the source FMEP interface.       |
| Destination-IP                 | IP address of the destination FMEP interface.       |
| Destination-IP mask            | IP address mask of the destination FMEP interface.  |
| IP-Protocol                    | IPv4 protocol configured.                           |
| Source-L4-Port                 | Source TCP or UDP port.                             |
| Destination-L4-Port            | Destination TCP or UDP port.                        |
| Destination multicast group IP | IP address of the destination multicast group.      |

## Sample Output

### show oam fabric flow-specification (All)

```

user@host> show oam fabric flow-specification

Flow specification name : fspec1 Type : Ethernet
 Ethernet frame size : Unspecified
 Ether type : Unspecified
 Source-MAC : Unspecified
 Source-MAC Mask : Unspecified
 Destination-MAC : Unspecified
 Destination-MAC Mask : Unspecified
Flow specification name : fspec2 Type : Ethernet
 Ethernet frame size : Unspecified
 Ether type : 1792
 Source-MAC : 0:0:a0:f0:cc:22
 Source-MAC Mask : Unspecified
 Destination-MAC : 0:11:22:33:a0:e2
 Destination-MAC Mask : Unspecified
Flow specification name : fspec3 Type : Ethernet Unicast IPV4
 Ethernet frame size : Unspecified
 Source-IP : 121.0.0.1
 Source-IP Mask : Unspecified
 Destination-IP : 132.1.0.1
 Destination-IP Mask : Unspecified
 Destination-mac : Unspecified
 IP-protocol : TCP
 Source-L4-port : 270

```

```
Destination-L4-port : Unspecified
Flow specification name : fspec4 Type : Multicast IPV4
Ethernet frame size : Unspecified
Source IP : 100.0.0.23
Destination multicast group IP: 225.0.0.1
Flow specification name : fspec5 Type : Multicast VLAN flood
Ethernet frame size : Unspecified
```

#### show oam fabric flow-specification (Flow Specification Specified)

```
user@host> show oam fabric flow-specification fspec3
```

```
Flow specification name : fspec3 Type : Ethernet Unicast IPV4
Ethernet frame size : Unspecified
Source-IP : 121.0.0.1
Source-IP Mask : Unspecified
Destination-IP : 132.1.0.1
Destination-IP Mask : Unspecified
Destination-mac : Unspecified
IP-protocol : TCP
Source-L4-port : 270
Destination-L4-port : Unspecified
```

## show oam fabric interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam fabric interfaces<br><interface-name><br><fabric-maintenance-association <i>fma-name</i> ><br><brief>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the fabric maintenance associations (FMAs) and fabric maintenance endpoints (FMEPs) that are configured in a QFabric system Node group.<br><br>If you do not specify an interface, display information for an FMA (if one is specified) or all FMAs configured in the Node group.                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>brief</b> —(Optional) Summarized version of the output.<br><br><b>fabric-maintenance-association <i>fma-name</i></b> —(Optional) Name of a specific FMA.<br><br><b>interface-name</b> —(Optional) Name of a specific interface.                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Internal Fabric Monitoring on page 6099</a></li> <li>• <a href="#">Configuring a Fabric Maintenance Association on page 6184</a></li> <li>• <a href="#">Configuring Flow Specifications on page 6185</a></li> <li>• <a href="#">fabric (OAM) on page 6256</a></li> <li>• <a href="#">ping fabric multicast-flow on page 6419</a></li> <li>• <a href="#">ping fabric unicast-flow on page 6421</a></li> <li>• <a href="#">traceroute fabric unicast-flow on page 6467</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam fabric interfaces on page 6436</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 621 on page 6435</a> lists the output fields for the <b>show oam fabric interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                            |

**Table 621: show oam fabric interfaces Output Fields**

| Field Name                     | Field Description                   |
|--------------------------------|-------------------------------------|
| Interface-name                 | Name of the interface.              |
| Fabric-Maintenance Association | FMA name.                           |
| VLAN                           | VLAN name.                          |
| Interface state                | State of the interface: Up or down. |

Table 621: show oam fabric interfaces Output Fields (*continued*)

| Field Name     | Field Description |
|----------------|-------------------|
| MEP Identifier | FMEP identifier.  |
| MEP Name       | FMEP name.        |

## Sample Output

### show oam fabric interfaces

```
user@host> show oam fabric interfaces
```

| Interface-name        | Fabric-Maintenance Association | VLAN | Interface state | MEP Identifier | MEP Name |
|-----------------------|--------------------------------|------|-----------------|----------------|----------|
| ED1479:NULL           | fma-default                    | *    | up              | 32774          |          |
| fmeop-default-ED1479  |                                |      |                 |                |          |
| ED1494:NULL           | fma-default                    | *    | up              | 32773          |          |
| fmeop-default-ED1494  |                                |      |                 |                |          |
| ED1494:xe-0/0/7.0     | fma1                           | v100 | up              | 1              |          |
| ED1497:NULL           | fma-default                    | *    | up              | 32775          |          |
| fmeop-default-ED1497  |                                |      |                 |                |          |
| ED1497:xe-0/0/8.0     | fma1                           | v100 | up              | 2              |          |
| P3613-C:NULL          | fma-default                    | *    | up              | 32777          |          |
| fmeop-default-P3613-C |                                |      |                 |                |          |



## show log

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>             | <a href="#">Syntax on page 6437</a><br><a href="#">Syntax (QFabric System) on page 6437</a><br><a href="#">Syntax (TX Matrix Routers) on page 6437</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                     | <pre>show log &lt;filename   user &lt;username&gt;&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (QFabric System)</b>    | <pre>show log filename &lt;device-type (device-id   device-alias)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (TX Matrix Routers)</b> | <pre>show log &lt;all-lcc   lcc number   scc&gt; &lt;filename   user &lt;username&gt;&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                | List log files, display log file contents, or display information about users who have logged in to the router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                    | <p><b>none</b>—List all log files.</p> <p><b>&lt;all-lcc   lcc number   scc&gt;</b>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p><b>device-type</b>—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p> <ul style="list-style-type: none"> <li>• <b>director-device</b>—Display logs for Director devices.</li> <li>• <b>infrastructure-device</b>—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).</li> <li>• <b>interconnect-device</b>—Display logs for Interconnect devices.</li> <li>• <b>node-device</b>—Display logs for Node devices.</li> </ul> |



**NOTE:** If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



**NOTE:** The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

**Required Privilege Level** trace

**List of Sample Output** [show log on page 6438](#)  
[show log filename on page 6438](#)  
[show log filename \(QFabric System\) on page 6439](#)  
[show log user on page 6439](#)

## Sample Output

### show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin 211663 Oct 1 19:44 dcd
-rw-r--r-- 1 root bin 999947 Oct 1 19:41 dcd.0
-rw-r--r-- 1 root bin 999994 Oct 1 17:48 dcd.1
-rw-r--r-- 1 root bin 238815 Oct 1 19:44 rpd
-rw-r--r-- 1 root bin 1049098 Oct 1 18:00 rpd.0
-rw-r--r-- 1 root bin 1061095 Oct 1 12:13 rpd.1
-rw-r--r-- 1 root bin 1052026 Oct 1 06:08 rpd.2
-rw-r--r-- 1 root bin 1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin 1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin 1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin 1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin 1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin 19656 Oct 1 19:37 wtmp
```

### show log filename

```
user@host> show log rpd
Oct 1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct 1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct 1 18:00:18
Oct 1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct 1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct 1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct 1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct 1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct 1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct 1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

### show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0-RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

### show log user

```

user@host> show log user
darius mg2546 Thu Oct 1 19:37 still logged in
darius mg2529 Thu Oct 1 19:08 - 19:36 (00:28)
darius mg2518 Thu Oct 1 18:53 - 18:58 (00:04)
root mg1575 Wed Sep 30 18:39 - 18:41 (00:02)
root ttyp2 jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex ttyp1 192.168.1.2 Wed Sep 30 01:03 - 01:22 (00:19)

```

## show sflow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow<br><collector><br><interface>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display sFlow configuration information.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Display all sFlow configuration information.</p> <p><b>collector</b>—(Optional) Display a list of configured sFlow collectors and their properties.</p> <p><b>interface</b>—(Optional) Display the interfaces on which sFlow technology is enabled and the sampling parameters.</p>                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show sflow interface on page 6443</a></li> <li>• <a href="#">show sflow collector on page 6442</a></li> <li>• <a href="#">clear sflow collector statistics on page 6401</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow on page 6441</a>                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 622 on page 6440</a> lists the output fields for the <b>show sflow</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                           |

**Table 622: show sflow Output Fields**

| Field Name          | Field Description                                                                                                  | Level of Output |
|---------------------|--------------------------------------------------------------------------------------------------------------------|-----------------|
| sFlow               | Status of the feature: <b>Enabled</b> or <b>Disabled</b> .                                                         | All levels      |
| Sample limit        | Number of packets sampled per second. This sample limit cannot be configured and is set to 300 packets per second. | All levels      |
| Polling interval    | Interval at which the sFlow agent polls the interface.                                                             | All levels      |
| Sample rate egress  | Rate at which egress packets are sampled.                                                                          | All levels      |
| Sample rate ingress | Rate at which ingress packets are sampled.                                                                         | All levels      |
| Agent ID            | IP address assigned to the sFlow agent.                                                                            | All levels      |
| Source IP address   | Source IP address for the sFlow packets.                                                                           | All levels      |

## Sample Output

show sflow

```
user@host> show sflow
```

```
sFlow : Enabled
Sample limit : 300 packets/second
Polling interval : 20 second
Sample rate egress : 1:2048: Disabled
Sample rate ingress : 1:1000: Enabled
Agent ID : 10.93.54.7
Source IP address : 10.93.54.7
```

## show sflow collector

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow collector                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display a list of configured sFlow collectors and their properties.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear sflow collector statistics on page 6401</a></li> <li>• <a href="#">show sflow on page 6440</a></li> <li>• <a href="#">show sflow interface on page 6443</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow collector on page 6442</a>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | <a href="#">Table 623 on page 6442</a> lists the output fields for the <b>show sflow collector</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                       |

**Table 623: show sflow collector Output Fields**

| Field Name        | Field Description                 | Level of Output |
|-------------------|-----------------------------------|-----------------|
| Collector address | IP address of the collector.      | All levels      |
| UDP-Port          | UDP port number of the collector. | All levels      |
| No. of samples    | Number of samples collected.      | All levels      |

## Sample Output

### show sflow collector

```
user@host> show sflow collector
```

```

Collector Udp-port No. of samples
address
10.204.32.46 6343 1000
100.204.32.76 3400 1000
```

## show sflow interface

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow interface                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display the interfaces on which sFlow is enabled and the sampling parameters for the interface.                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show sflow on page 6440</a></li> <li>• <a href="#">show sflow collector on page 6442</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6165</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6189</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow interface (QFX3500 Switch in Standalone Mode) on page 6443</a><br><a href="#">show sflow interface (QFabric System) on page 6444</a>                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 624 on page 6443</a> lists the output fields for the <b>show sflow interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                              |

**Table 624: show sflow interface Output Fields**

| Field Name                  | Field Description                                      | Level of Output |
|-----------------------------|--------------------------------------------------------|-----------------|
| Interface                   | Interface on which sFlow technology is enabled.        | All levels      |
| Status Egress               | Indicates whether an egress sample rate is enabled.    | All levels      |
| Status Ingress              | Indicates whether an ingress sample rate is enabled.   | All levels      |
| Sample rate Egress          | Rate at which egress packets are sampled.              | All levels      |
| Sample rate Ingress         | Rate at which ingress packets are sampled.             | All levels      |
| Adapted sample rate Egress  | Adapted rate at which egress packets are sampled.      | All levels      |
| Adapted sample rate Ingress | Adapted rate at which ingress packets are sampled.     | All levels      |
| Polling-interval            | Interval at which the sFlow agent polls the interface. | All levels      |

## Sample Output

### show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
```

| Interface  | Status  | Sample rate |         | Adapted sample rate |         | Polling-interval |    |
|------------|---------|-------------|---------|---------------------|---------|------------------|----|
|            |         | Egress      | Ingress | Egress              | Ingress |                  |    |
| xe-0/0/0.0 | Enabled | Disabled    | 1000    | 2048                | 1000    | 2048             | 20 |
| xe-1/0/1.0 | Enabled | Disabled    | 1000    | 2048                | 1000    | 2048             | 20 |

## Sample Output

### show sflow interface (QFabric System)

```
user@host> show sflow interface
```

| Interface        | Status  | Sample rate |         | Adapted sample rate |         | Polling-interval |
|------------------|---------|-------------|---------|---------------------|---------|------------------|
|                  |         | Egress      | Ingress | Egress              | Ingress |                  |
| node1:xe-0/0/0.0 | Enabled | Disabled    | 1000    | 2048                | 1000    | 2048             |
| node2:xe-1/0/1.0 | Enabled | Disabled    | 1000    | 2048                | 1000    | 2048             |
| node4:xe-1/0/0.0 | Enabled | Disabled    | 1000    | 2048                | 1000    | 2048             |



## show snmp health-monitor

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp health-monitor<br><alarms (brief   detail)   logs>                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display information about all health monitor alarms and logs.</p> <p><b>alarms (brief   detail)</b>—(Optional) Display information about health monitor alarms. Optionally, specify brief or detailed information about the alarms.</p> <p><b>logs</b>—(Optional) Display information about health monitor logs.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show snmp health-monitor on page 6447</a><br><a href="#">show snmp health-monitor alarms detail on page 6447</a>                                                                                                                                                                                                        |
| <b>Output Fields</b>            | Table 625 on page 6445 describes the output fields for the <b>show snmp health-monitor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                         |

**Table 625: show snmp health-monitor Output Fields**

| Field Name                  | Field Description                                                           | Level of Output |
|-----------------------------|-----------------------------------------------------------------------------|-----------------|
| <b>Alarm Index</b>          | Alarm identifier.                                                           | All levels      |
| <b>Variable description</b> | Description of the health monitor object instance being monitored.          | All levels      |
| <b>Variable name</b>        | Name of the health monitor object instance being monitored.                 | All levels      |
| <b>Value</b>                | Current value of the monitored variable in the most recent sample interval. | All levels      |

Table 625: show snmp health-monitor Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>            | <p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> <li>Alarms: <ul style="list-style-type: none"> <li><b>active</b>—Entry is fully configured and activated.</li> <li><b>falling threshold crossed</b>—Value of the variable has crossed the lower threshold limit.</li> <li><b>rising threshold crossed</b>—Value of the variable has crossed the upper threshold limit.</li> <li><b>under creation</b>—Entry is being configured and is not yet activated.</li> <li><b>startup</b>—Alarm is waiting for the first sample of the monitored variable.</li> <li><b>object not available</b>—Monitored variable of that type is not available to the health monitor agent.</li> <li><b>instance not available</b>—Monitored variable's instance is not available to the health monitor agent.</li> <li><b>object type invalid</b>—Monitored variable is not a numeric value.</li> <li><b>object processing errored</b>—An error occurred when the monitored variable was processed.</li> <li><b>unknown</b>—State is not one of the above.</li> </ul> </li> </ul> | All levels      |
| <b>Variable OID</b>     | Object ID to which the variable name is resolved. The format is x.x.x.x.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail          |
| <b>Sample type</b>      | Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value <b>absolute value</b> or <b>delta value</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail          |
| <b>Startup alarm</b>    | <p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either <b>rising alarm</b> or <b>rising or falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <b>falling alarm</b> or <b>rising or falling alarm</b>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <b>falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <b>rising alarm</b>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul>                                                                         | detail          |
| <b>Owner</b>            | Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail          |
| <b>Creator</b>          | Mechanism by which the entry was configured ( <b>Health Monitor</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail          |
| <b>Sample interval</b>  | Time period between samples (in seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail          |
| <b>Rising threshold</b> | Upper limit threshold value as a percentage of the maximum possible value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail          |

Table 625: show snmp health-monitor Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                               | Level of Output |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Falling threshold   | Lower limit threshold value as a percentage of the maximum possible value.                                                                                                      | detail          |
| Rising event index  | Index number of the event triggered when the rising threshold is crossed.                                                                                                       | detail          |
| Falling event index | Index number of the event triggered when the falling threshold is crossed. Details include the value of the falling event instance and the state of the falling event instance. | detail          |

## Sample Output

### show snmp health-monitor

```

user@switch> show snmp health-monitor

Alarm
Index Variable description Value State

32768 Health Monitor: root file system utilization
 jnxHrStoragePercentUsed.1 59 active

32769 Health Monitor: /config file system utilization
 jnxHrStoragePercentUsed.2 0 active

32770 Health Monitor: RE 0 CPU utilization
 jnxOperatingCPU.9.1.0.0 9 falling threshold

32772 Health Monitor: RE 0 memory utilization
 jnxOperatingBuffer.9.1.0.0 23 active

32774 Health Monitor: Max Kernel Memory Used (%)
 jnxBoxKernelMemoryUsedPercent.0 3 active
Event Index: 32768
Description: Health Monitor: RE 0 CPU utilization crossed falling threshold
70 (value: 5), (variable: jnxOperatingCPU.9.1.0.0)
Time: 2011-01-09 19:18:35 PST

```

### show snmp health-monitor alarms detail

```

user@switch> show snmp health-monitor alarms detail

Alarm Index 32768:
Variable name jnxHrStoragePercentUsed.1
Variable OID 1.3.6.1.4.1.2636.3.31.1.1.1.1.1
Sample type absolute value
Startup alarm rising alarm
Owner Health Monitor: root file system
 utilization
Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 80

```

Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 59  
Instance State: active

Alarm Index 32769:

Variable name jnxHrStoragePercentUsed.2  
Variable OID 1.3.6.1.4.1.2636.3.31.1.1.1.2  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: /config file system  
utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 0  
Instance State: active

Alarm Index 32770:

Variable name jnxOperatingCPU.9.1.0.0  
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: RE 0 CPU utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 9  
Instance State: falling threshold

Alarm Index 32772:

Variable name jnxOperatingBuffer.9.1.0.0  
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: RE 0 memory utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 23  
Instance State: active

Alarm Index 32774:

Variable name jnxBoxKernelMemoryUsedPercent.0  
Variable OID 1.3.6.1.4.1.2636.3.1.16.0  
Sample type absolute value

|                     |                                            |
|---------------------|--------------------------------------------|
| Startup alarm       | rising alarm                               |
| Owner               | Health Monitor: Max Kernel Memory Used (%) |
| Creator             | Health Monitor                             |
| State               | active                                     |
| Sample interval     | 300 seconds                                |
| Rising threshold    | 80                                         |
| Falling threshold   | 70                                         |
| Rising event index  | 32768                                      |
| Falling event index | 32768                                      |
| Instance Value:     | 3                                          |
| Instance State:     | active                                     |

## show snmp inform-statistics

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp inform-statistics                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.           |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) inform requests.                                                                                                           |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show snmp inform-statistics on page 6450</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 626 on page 6450</a> describes the output fields for the <b>show snmp inform-statistics</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 626: show snmp inform-statistics Output Fields**

| Field Name            | Field Description                                                                                                                                  |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Target Name</b>    | Name of the device configured to receive and respond to SNMP informs.                                                                              |
| <b>Address</b>        | IP address of the target device.                                                                                                                   |
| <b>Sent</b>           | Number of informs sent to the target device and acknowledged by the target device.                                                                 |
| <b>Pending</b>        | Number of informs held in memory pending a response from the target device.                                                                        |
| <b>Discarded</b>      | Number of informs discarded after the specified number of retransmissions to the target device were attempted.                                     |
| <b>Timeouts</b>       | Number of informs that did not receive an acknowledgement from the target device within the timeout specified.                                     |
| <b>Probe Failures</b> | Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address). |

## Sample Output

### show snmp inform-statistics

```

user@host> show snmp inform-statistics
Inform Request Statistics:
 Target Name: TA1_v3_md5_none Address: 172.17.20.184
 Sent: 176, Pending: 0
 Discarded: 0, Timeouts: 0, Probe Failures: 0
 Target Name: TA2_v3_sha_none Address: 192.168.110.59

```

Sent: 0, Pending: 4  
Discarded: 84, Timeouts: 0, Probe Failures: 258  
Target Name: TA5\_v2\_none Address: 172.17.20.184  
Sent: 0, Pending: 0  
Discarded: 2, Timeouts: 10, Probe Failures: 0

## show snmp mib

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show snmp mib (get   get-next   walk) (ascii   decimal) <i>object-id</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>get</b>—Retrieve and display one or more SNMP object values.</p> <p><b>get-next</b>—Retrieve and display the next SNMP object values.</p> <p><b>walk</b>—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.</p> <p><b>ascii</b>—Display the SNMP object's string indices as an ASCII-key representation.</p> <p><b>decimal</b>—Display the SNMP object values in the decimal (default) format. The <b>decimal</b> option is the default option for this command. Therefore, issuing the <b>show snmp mib (get   get-next   walk) decimal object-id</b> and the <b>show snmp mib (get   get-next   walk) object-id</b> commands display the same output.</p> <p><b>object-id</b>—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as <b>interfaces</b>). When entering multiple objects, enclose the objects in quotation marks.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show snmp mib get on page 6453</a></p> <p><a href="#">show snmp mib get (Multiple Objects) on page 6453</a></p> <p><a href="#">show snmp mib get (Layer 2 Policer) on page 6453</a></p> <p><a href="#">show snmp mib get-next on page 6453</a></p> <p><a href="#">show snmp mib get-next (Specify an OID) on page 6453</a></p> <p><a href="#">show snmp mib walk on page 6453</a></p> <p><a href="#">show snmp mib walk (QFX Series) on page 6453</a></p> <p><a href="#">show snmp mib walk decimal on page 6454</a></p> <p><a href="#">show snmp mib walk (ASCII) on page 6454</a></p> <p><a href="#">show snmp mib walk (Multiple Indices) on page 6454</a></p> <p><a href="#">show snmp mib walk decimal (Multiple Indices) on page 6454</a></p>                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | <p>Table 627 on page 6453 describes the output fields for the <b>show snmp mib</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



Table 627: show snmp mib Output Fields

| Field Name          | Field Description                                                               |
|---------------------|---------------------------------------------------------------------------------|
| <i>name</i>         | Object name and numeric instance value.                                         |
| <i>object value</i> | Object value. The Junos OS translates OIDs into the corresponding object names. |

## Sample Output

### show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

### show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

### show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

### show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

### show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

### show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

### show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 138980301
sysContact.0 = System Contact
```

```
sysName.0 = LabQFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

#### show snmp mib walk decimal

```
user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

#### show snmp mib walk (ASCII)

```
show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

#### show snmp mib walk (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

#### show snmp mib walk decimal (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

## show snmp rmon

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp rmon<br><alarms (brief   detail)><br><events (brief   detail)><br><logs>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms, events, and logs.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display information about all RMON alarms and events.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about RMON alarms or events.</p> <p><b>alarms</b>—(Optional) Display information about RMON alarms.</p> <p><b>events</b>—(Optional) Display information about RMON events.</p> <p><b>logs</b>—(Optional) Display information about RMON monitoring logs.</p>                                                                                                                          |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li> <li>• <a href="#">Understanding RMON on page 6119</a></li> <li>• <a href="#">clear snmp statistics on page 6403</a></li> <li>• <a href="#">clear snmp history on page 6402</a></li> <li>• <a href="#">show snmp rmon history on page 6459</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show snmp rmon on page 6457</a><br><a href="#">show snmp rmon alarms detail on page 6458</a><br><a href="#">show snmp rmon events detail on page 6458</a><br><a href="#">show snmp rmon logs on page 6458</a>                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <p><a href="#">Table 628 on page 6455</a> describes the output fields for the <b>show snmp rmon</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                 |

**Table 628: show snmp rmon Output Fields**

| Field Name  | Field Description | Level of Output |
|-------------|-------------------|-----------------|
| Alarm Index | Alarm identifier. | All levels      |

Table 628: show snmp rmon Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>         | <p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry is fully configured and activated.</li> <li>• <b>falling threshold crossed</b>—Value of the variable has crossed the lower threshold limit.</li> <li>• <b>rising threshold crossed</b>—Value of the variable has crossed the upper threshold limit.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>startup</b>—Alarm is waiting for the first sample of the monitored variable.</li> <li>• <b>object not available</b>—Monitored variable of that type is not available to the SNMP agent.</li> <li>• <b>instance not available</b>—Monitored variable's instance is not available to the SNMP agent.</li> <li>• <b>object type invalid</b>—Monitored variable is not a numeric value.</li> <li>• <b>object processing errored</b>—An error occurred when the monitored variable was processed.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul> <p>Events:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry has been fully configured and activated.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul> | All levels      |
| <b>Variable name</b> | Name of the SNMP object instance being monitored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Event Index</b>   | Event identifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Type</b>          | <p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>log</b>—A system log message is generated and an entry is made to the log table.</li> <li>• <b>snmptrap</b>—An SNMP trap is sent to the configured destination.</li> <li>• <b>log and trap</b>—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination.</li> <li>• <b>none</b>—Neither log nor trap will be sent.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Last Event</b>    | Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief</b>    |
| <b>Community</b>     | Trap group used for sending the SNMP trap.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Variable OID</b>  | Object ID to which the variable name is resolved. The format is x.x.x.x.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |
| <b>Sample type</b>   | Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of <b>absolute value</b> or <b>delta value</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail</b>   |

Table 628: show snmp rmon Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Startup alarm</b>       | Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either <b>rising alarm</b> or <b>rising or falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <b>falling alarm</b> or <b>rising or falling alarm</b>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <b>falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <b>rising alarm</b>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul> | <b>detail</b>   |
| <b>Owner</b>               | Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>Creator</b>             | Mechanism by which the entry was configured ( <b>CLI</b> or <b>SNMP</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>   |
| <b>Sample interval</b>     | Time period between samples (in seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Rising threshold</b>    | Upper limit threshold value configured by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>Falling threshold</b>   | Lower limit threshold value configured by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>Rising event index</b>  | Event triggered when the rising threshold is crossed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Falling event index</b> | Event triggered when the falling threshold is crossed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>   |
| <b>Current value</b>       | Current value of the monitored variable in the most recent sample interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>   |

## Sample Output

### show snmp rmon

```

user@host> show snmp rmon
Alarm
Index Variable description Value State

 5 monitor
 jnxOperatingCPU.9.1.0.0 5 falling threshold

Event
Index Type Last Event
 1 log and trap 2009-07-10 11:34:17 PDT
Event Index: 1
Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
Time: 2009-07-10 11:34:07 PDT

```

Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,  
(variable: jnxOperatingCPU.9.1.0.0, value: 5)  
Time: 2009-07-10 11:34:17 PDT

#### show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail
Alarm Index 5:
 Variable name jnxOperatingCPU.9.1.0.0
 Variable OID 1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
 Sample type absolute value
 Startup alarm rising or falling alarm
 Owner monitor

 Creator CLI
 State active
 Sample interval 5 seconds
 Rising threshold 90
 Falling threshold 75
 Rising event index 1
 Falling event index 1
 Instance Value: 4
 Instance State: falling threshold
```

#### show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
 Description rmon event
 Type log and trap
 Community rmon-trap-group
 Last event 2009-07-10 11:34:17 PDT
 Creator CLI
 State active
```

#### show snmp rmon logs

```
user@host> show snmp rmon logs
Event Index: 1
 Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
 Time: 2009-07-10 11:34:07 PDT
 Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
 Time: 2009-07-10 11:34:17 PDT
```

---

## show snmp rmon history

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp rmon history<br><history-index><br>sample-index <sample-index>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display the contents of the RMON history group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display all the entries in the RMON history group.</p> <p><b>history-index</b>—(Optional) Display the contents of the specified entry in the RMON history group.</p> <p><b>sample-index sample-index</b>—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.</p>                                                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6121</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6389</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6198</a></li><li>• <a href="#">Understanding RMON on page 6119</a></li><li>• <a href="#">clear snmp statistics on page 6403</a></li><li>• <a href="#">clear snmp history on page 6402</a></li><li>• <a href="#">show snmp rmon on page 6455</a></li></ul> |

## show snmp statistics

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp statistics                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.                                                                    |
| <b>Options</b>                  | This command has no options.                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear snmp statistics on page 6403</a></li> </ul>                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show snmp statistics on page 6463</a>                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 178 on page 1859</a> describes the output fields for the <b>show snmp statistics</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 629: show snmp statistics Output Fields

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input</b> | <p>Information about received packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets(snmplnPkts)</b>—Total number of messages delivered to the SNMP entity from the transport service.</li> <li>• <b>Bad versions—(snmplnBadVersions)</b> Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version.</li> <li>• <b>Bad community names—(snmplnBadCommunityNames)</b> Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity.</li> <li>• <b>Bad community uses—(snmplnBadCommunityUses)</b> Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.</li> <li>• <b>ASN parse errors—(snmplnASNParseErrs)</b> Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.</li> <li>• <b>Too big—(snmplnTooBig)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmplnNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmplnBadValues)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>Read only—(snmplnReadOnly)</b> Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of <b>readOnly</b>. Only incorrect implementations of SNMP generate this error.</li> </ul> |



Table 629: show snmp statistics Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input (continued) | <ul style="list-style-type: none"> <li>• <b>General errors—(snmpInGenErrs)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Total requests varbinds—(snmpInTotalReqVars)</b> Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP <b>GetRequest</b> and <b>GetNext</b> PDUs.</li> <li>• <b>Total set varbinds—(snmpInSetVars)</b> Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP <b>SetRequest</b> PDUs.</li> <li>• <b>Get requests—(snmpInGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get nexts—(snmpInGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Set requests—(snmpInSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get responses—(snmpInGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Traps—(snmpInTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> <li>• <b>Silent drops—(snmpSilentDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.</li> <li>• <b>Proxy drops—(snmpProxyDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.</li> <li>• <b>Commit pending drops</b>—Number of SNMP packets for <b>Set</b> requests dropped because of a previous pending SNMP <b>Set</b> request on the committed configuration.</li> <li>• <b>Throttle drops</b>—Number of SNMP packets for any requests dropped reaching the throttle limit.</li> </ul> |

Table 629: show snmp statistics Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V3 Input   | <p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> <li>• <b>Unknown security models—(snmpUnknownSecurityModels)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine.</li> <li>• <b>Invalid messages—(snmpInvalidMsgs)</b> Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message.</li> <li>• <b>Unknown pdu handlers—(snmpUnknownPDUHandlers)</b> Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type.</li> <li>• <b>Unavailable contexts—(snmpUnavailableContexts)</b> Number of requests received for a context that is known to the SNMP engine, but is currently unavailable.</li> <li>• <b>Unknown contexts—(snmpUnknownContexts)</b> Total number of requests received for a context that is unknown to the SNMP engine.</li> <li>• <b>Unsupported security levels—(usmStatsUnsupportedSecLevels)</b> Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable).</li> <li>• <b>Not in time windows—(usmStatsNotInTimeWindows)</b> Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window.</li> <li>• <b>Unknown user names—(usmStatsUnknownUserNames)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine.</li> <li>• <b>Unknown engine ids—(usmStatsUnknownEngineIDs)</b> Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine.</li> <li>• <b>Wrong digests—(usmStatsWrongDigests)</b> Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value.</li> <li>• <b>Decryption errors—(usmStatsDecryptionErrors)</b> Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.</li> </ul> |

Table 629: show snmp statistics Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output</b> | <p>Information about transmitted packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets—(snmpOutPkts)</b> Total number of messages passed from the SNMP entity to the transport service.</li> <li>• <b>Too big—(snmpOutTooBig)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmpOutNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmpOutBadValues)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>General errors—(snmpOutGenErrs)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Get requests—(snmpOutGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get nexts—(snmpOutGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs generated by the SNMP entity.</li> <li>• <b>Set requests—(snmpOutSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get responses—(snmpOutGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs generated by the SNMP entity.</li> <li>• <b>Traps—(snmpOutTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> </ul> |

## Sample Output

### show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
 Input:
 Packets: 246213, Bad versions: 12, Bad community names: 12,
 Bad community uses: 0, ASN parse errors: 96,
 Too big: 0, No such names: 0, Bad values: 0,
 Read onlys: 0, General errors: 0,
 Total request varbinds: 227084, Total set varbinds: 67,
 Get requests: 44942, Get nexts: 190371, Set requests: 10712,
 Get responses: 0, Traps: 0,
 Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
 Throttle drops: 0,
 V3 Input:
 Unknown security models: 0, Invalid messages: 0
 Unknown pdu handlers: 0, Unavailable contexts: 0
 Unknown contexts: 0, Unsupported security levels: 1
 Not in time windows: 0, Unknown user names: 0
 Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
 Output:
 Packets: 246093, Too big: 0, No such names: 31561,
 Bad values: 0, General errors: 2,
 Get requests: 0, Get nexts: 0, Set requests: 0,
 Get responses: 246025, Traps: 0

```

## show snmp v3

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show snmp v3</code><br><code>&lt;access &lt;brief   detail&gt;   community   general   groups   notify &lt;filter&gt;   target &lt;address   parameters&gt;   users&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>none</b>—Display all of the SNMPv3 operating configuration.</p> <p><b>access</b>—(Optional) Display SNMPv3 access information.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about SNMPv3 access information.</p> <p><b>community</b>—(Optional) Display SNMPv3 community information.</p> <p><b>general</b>—(Optional) Display SNMPv3 general information.</p> <p><b>groups</b>—(Optional) Display SNMPv3 security-to-group information.</p> <p><b>notify &lt;filter&gt;</b>—(Optional) Display SNMPv3 notify information and, optionally, notify filter information.</p> <p><b>target &lt;address   parameters&gt;</b>—(Optional) Display SNMPv3 target information and, optionally, either target address or target parameter information.</p> <p><b>users</b>—(Optional) Display SNMPv3 user information.</p> |
| <b>Additional Information</b>   | To edit the default display of the <b>show snmp v3</b> command, specify options in the <b>show</b> statement at the <b>[edit snmp v3]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show snmp v3 on page 6465</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 630 on page 6465</a> describes the output fields for the <b>show snmp v3</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 630: show snmp v3 Output Fields

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local engine              | <p>Information about the local SNMP engine configuration:</p> <ul style="list-style-type: none"> <li>• <b>Local engine ID</b>—Unique Identifier of the local SNMPv3 engine.</li> <li>• <b>Engine boots</b>—Number of times the local SNMPv3 engine has rebooted or reinitialized since this engine ID was configured.</li> <li>• <b>Engine time</b>—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized.</li> <li>• <b>Max msg size</b>—Maximum message size the sender can accommodate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Engine ID (local engine)  | <p>Information about the local SNMP engine ID and the associated users:</p> <ul style="list-style-type: none"> <li>• <b>User</b>—SNMPv3 username.</li> <li>• <b>Auth/Priv</b>—Authentication and encryption algorithm that is configured for the user.</li> <li>• <b>Storage</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of the user as listed in the SNMPv3 user table. Only rows with an active status in the table are used by the SNMPv3 engine.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Engine ID (remote engine) | <p>Information about a remote SNMP engine, associated users, user groups, and user access policies:</p> <ul style="list-style-type: none"> <li>• <b>User</b>—SNMPv3 username.</li> <li>• <b>Auth/Priv</b>—Authentication and encryption algorithm that is configured for the user.</li> <li>• <b>Storage</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of a new user that has been activated. Only users with an active status can use SNMPv3.</li> <li>• <b>Group name</b>—Name of a group of users for which the configured access privileges apply.</li> <li>• <b>Security model</b>—Security model (such as <b>usm</b>, <b>v1</b>, <b>v2c</b>, or <b>any</b>) that is configured for the group. The security model is used with the security name to ensure messaging security.</li> <li>• <b>Security name</b>—Security name that is associated with a user, and which is used with the security model to ensure messaging security.</li> <li>• <b>Storage type</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of a user in a group. Only users with an active status can use SNMPv3.</li> </ul> |
| Access control            | <p>Information about access control:</p> <ul style="list-style-type: none"> <li>• <b>Group name</b>—Name of a group of users for which the configured access privileges apply.</li> <li>• <b>Context prefix</b>—SNMPv3 context for which the configured access privileges apply.</li> <li>• <b>Security model/level</b>—Security model and security level combination that is configured for user access privileges.</li> <li>• <b>Read view</b>—Identifies the MIB view used for SNMPv3 read operations.</li> <li>• <b>Write view</b>—Identifies the MIB view used for SNMPv3 write operations.</li> <li>• <b>Notify view</b>—Identifies the MIB view used for outbound SNMP notifications.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show snmp v3

```
user@host> show snmp v3
```

Local engine ID: 80 00 0a 4c e04 31 32 33 34  
Engine boots: 38  
Engine time: 64583 seconds  
Max msg size: 2048 bytes

Engine ID: local

| User  | Auth/Priv | Storage     | Status |
|-------|-----------|-------------|--------|
| user1 | md5/des   | nonvolatile | active |
| user2 | sha/none  | nonvolatile | active |
| user3 | none/none | nonvolatile | active |

Engine ID: 81 00 0a 4c 04 64 64 64 64

| User | Auth/Priv | Storage     | Status |
|------|-----------|-------------|--------|
| UNEW | md5/none  | nonvolatile | active |

| Group name | Security model | Security name | Storage type | Status |
|------------|----------------|---------------|--------------|--------|
| g1         | usm            | user1         | nonvolatile  | active |
| g2         | usm            | user2         | nonvolatile  | active |
| g3         | usm            | user3         | nonvolatile  | active |

Access control:

| Group | Context prefix | Security model/level | Read view | Write view | Notify view |
|-------|----------------|----------------------|-----------|------------|-------------|
| g1    |                | usm/privacy          | v1        | v1         |             |
| g2    |                | usm/authent          | v1        | v1         |             |
| g3    |                | usm/none             | v1        | v1         |             |

## traceroute fabric unicast-flow

**Syntax** `traceroute fabric unicast-flow flow-spec-name flow-specification-name fma fma-name  
 <source-fmep fmep-name>  
 <dest-fmep fmep-name>  
 <forced-fte-interface fte-interface>  
 <count count>  
 <verbose>`

**Release Information** Command introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Trace the path taken by a specific unicast flow (ping) operation across a VLAN on the QFabric system from a source maintenance endpoint (FMEP) to a destination FMEP. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices.

The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the protocol data unit (PDU) travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.



**NOTE:** This command is not supported on the QFX3000-M QFabric system.

**Options** `dest-fmep-id source-fmep-id`—(Optional) ID of the destination FMEP of the traceroute operation.

`flow-spec-name flow-specification-name`—Name of the flow specification that defines the parameters of the traceroute operation.

`fma-name fma-name`—Name of the FMA that associates a VLAN with the FMEPs.

`forced-fte-interface fte-interface`—(Optional) Option to force the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.

`source-fmep-id source-fmep-id`—(Optional) ID of the FMEP that is the source of the traceroute operation.

`verbose`—(Optional) Detailed version of the output display.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of Internal Fabric Monitoring on page 6099](#)
- [Configuring a Fabric Maintenance Association on page 6184](#)

- [Configuring Flow Specifications on page 6185](#)
- [fabric \(OAM\) on page 6256](#)
- [ping fabric multicast-flow on page 6419](#)
- [ping fabric unicast-flow on page 6421](#)

**List of Sample Output** [traceroute fabric unicast-flow on page 6468](#)  
[traceroute fabric unicast-flow \(Verbose\) on page 6468](#)

**Output Fields** [Table 631 on page 6468](#) lists the output fields for the **traceroute fabric-unicast-flow** command. Output fields are listed in the approximate order in which they appear.

**Table 631: traceroute fabric unicast-flow Output Fields**

| Field Name                                                                                             | Field Description                                                                                      |
|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Fabric flow traceroute between source <i>source-fmep-name</i> destination <i>destination-fmep-name</i> | Source and destination FMEP names.                                                                     |
| Using fabric flow-specification                                                                        | Name of the flow specification.                                                                        |
| Ethernet frame-size                                                                                    | Ethernet frame size.                                                                                   |
| Source-MAC                                                                                             | MAC address of the source FMEP interface.                                                              |
| Destination-MAC                                                                                        | MAC address of the destination FMEP interface.                                                         |
| Ethertype                                                                                              | EtherType protocol.                                                                                    |
| Received response from fmip-id, interface                                                              | FMIP identifier and interface name of the FMIP that received the PDU and responded to the source FMEP. |
| Received total <i>number</i> responses                                                                 | Number of responses that are received after the traceroute operation has finished.                     |

## Sample Output

### traceroute fabric unicast-flow

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```

## Sample Output

### traceroute fabric unicast-flow (Verbose)

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1 verbose
```



```
Fabric flow traceroute between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:EF:7:99:85:10
Destination-MAC: 0:D3:A6:11:39:21
Ethertype: 11564
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```



## CHAPTER 70

# Troubleshooting

- [Troubleshooting Overview on page 6471](#)
- [Troubleshooting Procedures on page 6475](#)

### Troubleshooting Overview

---

- [Understanding Troubleshooting Resources on page 6471](#)
- [Troubleshooting Overview on page 6473](#)

### Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

[Table 632 on page 6471](#) provides a list of some of the troubleshooting resources.

**Table 632: Troubleshooting Resources on the QFX Series**

| Troubleshooting Resource              | Description                                                                                                                                               | Documentation                                                             |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| Chassis alarms                        | Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch. | <a href="#">"Chassis Alarm Messages on a QFX3500 Device" on page 6488</a> |
| Chassis Status LEDs and Fan Tray LEDs | A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.                | <i>Chassis Status LEDs on a QFX3500 Device</i>                            |
| Interface alarms                      | A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.                                           | <a href="#">"Interface Alarm Messages" on page 6491</a>                   |
| System alarms                         | A predefined alarm is triggered by a missing rescue configuration or problem with the software license.                                                   | <a href="#">"Understanding Alarms" on page 6487</a>                       |

Table 632: Troubleshooting Resources on the QFX Series (*continued*)

| Troubleshooting Resource                      | Description                                                                                                                                                                                                                                                                                                                                                                                                | Documentation                                                                                                                                                                                                                                                    |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System log messages                           | The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6155</a></li> <li>• <a href="#">Junos OS System Log Configuration Statements on page 6208</a></li> </ul>                                    |
| Junos OS operational mode commands            | Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the <b>traceroute monitor</b> command to locate points of failure in a network.                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 315</a></li> <li>• <a href="#">Monitoring System Properties on page 316</a></li> <li>• <a href="#">traceroute monitor</a></li> </ul>                          |
| Junos OS automation scripts (event scripts)   | Event scripts can be used to automate network troubleshooting and management tasks.                                                                                                                                                                                                                                                                                                                        | <i>Junos OS Configuration and Operations Automation Guide</i>                                                                                                                                                                                                    |
| Junos OS XML operational tags                 | XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.                                                                                                                                                                                                                                                    | <i>Junos XML API Operational Developer Reference</i>                                                                                                                                                                                                             |
| NETCONF XML management protocol               | The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations. | <i>NETCONF XML Management Protocol Guide</i>                                                                                                                                                                                                                     |
| SNMP MIBs and traps                           | MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6124</a></li> <li>• <a href="#">SNMP Traps Support on page 6139</a></li> <li>• <a href="#">Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS</a></li> </ul> |
| AI-Scripts and Advanced Insight Manager (AIM) | AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.                                                                                                                                                   | <a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>                                                                                                                                                                                              |
| Junos Space Service Now                       | This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.                                                                                                                         | <i>Service Automation</i>                                                                                                                                                                                                                                        |

Table 632: Troubleshooting Resources on the QFX Series (*continued*)

| Troubleshooting Resource        | Description                                                                                                                                                                                                                                                                                                                                 | Documentation                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| Junos Space Service Insight     | This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now. | <i>Service Automation</i>                                 |
| Juniper Networks Knowledge Base | You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.                                                                                                                                                                                                                        | <a href="http://kb.juniper.net">http://kb.juniper.net</a> |

## Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series product.

[Table 633 on page 6473](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 633: Troubleshooting on the QFX Series

| Problem Category           | Symptom or Problem                                                          | Recommended Action                                                                                     |
|----------------------------|-----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Switch hardware components | LCD panel shows a chassis alarm count.                                      | See <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 6488</a> .                        |
|                            | Fan tray LED is blinking amber.                                             | See <i>Fan Tray LED on a QFX3500 Device</i> .                                                          |
|                            | Chassis status LED for the power is blinking amber.                         | See <i>Chassis Status LEDs on a QFX3500 Device</i> .                                                   |
|                            | Chassis status LED for the fan (on the management board) is blinking amber. | Replace the management board as soon as possible. See <i>Chassis Status LEDs on a QFX3500 Device</i> . |

Table 633: Troubleshooting on the QFX Series (*continued*)

| Problem Category               | Symptom or Problem                                                                                          | Recommended Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port configuration             | Cannot configure a port as a Gigabit Ethernet port.                                                         | <p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                | Cannot configure a port as a Fibre Channel port.                                                            | <p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                                | Cannot configure a port as a 10-Gigabit Ethernet port.                                                      | <p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p> |
|                                | Cannot configure a 40-Gbps QSFP+ interface.                                                                 | <p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                         |
| External devices (USB devices) | Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state. | Unplug the USB device and reboot the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Initial device configuration   | Cannot configure management Ethernet ports.                                                                 | <p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p><b>NOTE:</b> The management ports are on the front panel of the QFX3500 switch. They are labeled <b>C0</b> and <b>C1</b> on the front panel. In the CLI they are referred to as <b>me0</b> and <b>me1</b>.</p> <p>See <a href="#">“Configuring a QFX3500 Device as a Standalone Switch” on page 163</a>.</p>                                                                                                                                                                                                                                                  |

Table 633: Troubleshooting on the QFX Series (*continued*)

| Problem Category                   | Symptom or Problem                                                                                                                                                    | Recommended Action                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Software upgrade and configuration | Failed software upgrade.                                                                                                                                              | See <a href="#">“Recovering from a Failed Software Installation” on page 119.</a>                                                                                                                                                                                                                                                                                                                                                                   |
|                                    | Active partition becomes inactive after upgrade.                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                    | Problem with the active configuration file.                                                                                                                           | See the following topics: <ul style="list-style-type: none"> <li>• <a href="#">Loading a Previous Configuration File on page 1600</a></li> <li>• <a href="#">Reverting to the Default Factory Configuration on page 173</a></li> <li>• <a href="#">Reverting to the Rescue Configuration on page 175</a></li> <li>• <a href="#">Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111</a></li> </ul> |
|                                    | Root password is lost or forgotten.                                                                                                                                   | Recover the root password. See <a href="#">“Recovering the Root Password” on page 1159.</a>                                                                                                                                                                                                                                                                                                                                                         |
| Network interfaces                 | An aggregated Ethernet interface is down.                                                                                                                             | See <a href="#">“Troubleshooting an Aggregated Ethernet Interface” on page 1161.</a>                                                                                                                                                                                                                                                                                                                                                                |
|                                    | Interface on built-in network port is down.                                                                                                                           | See <a href="#">“Troubleshooting Network Interfaces” on page 1160.</a>                                                                                                                                                                                                                                                                                                                                                                              |
|                                    | Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Ethernet switching                 | A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. | See <a href="#">“Troubleshooting Ethernet Switching” on page 2261.</a>                                                                                                                                                                                                                                                                                                                                                                              |
| Firewall filter                    | Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.                                                                            | See <a href="#">“Troubleshooting Firewall Filter Configuration” on page 4873.</a>                                                                                                                                                                                                                                                                                                                                                                   |

## Troubleshooting Procedures

- [Recovering from a Failed Software Installation on page 6475](#)
- [Loading a Previous Configuration File on page 6476](#)
- [Reverting to the Default Factory Configuration on page 6477](#)
- [Reverting to the Rescue Configuration on page 6478](#)
- [Recovering the Root Password on page 6478](#)

## Recovering from a Failed Software Installation

**Problem**    **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution** If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **ftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Loading a Previous Configuration File

You can use the **rollback <number>** command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

### Syntax

```
rollback <number>
```

### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.



- **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
- **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related Documentation** • [Configuration File Terms on page 48](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

**Related Documentation** • [Understanding Configuration Files on page 1590](#)  
 • [Loading a Previous Configuration File on page 1600](#)  
 • [Reverting to the Rescue Configuration on page 175](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

### Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1609](#)
- [Reverting to the Default Factory Configuration on page 173](#)
- [Configuration File Terms on page 48](#)

## Recovering the Root Password

If you forget the root password for the QFX3500 switch, you can use the password recovery procedure to reset the root password.



NOTE: The root password cannot be recovered on a QFabric system.



NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the QFX3500 switch power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the switch into the RJ-45-to-DB-9 serial port adapter supplied with the switch.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the switch.
6. Turn on the power to the management device.

7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the switch by (if necessary) plugging the power cords into the QFX3500 switch power supply, or turning on the power to the device or switch by switching on the AC power outlet the device is plugged into.  
  
The terminal emulation screen on your management device displays the switch's boot sequence.
10. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:  
  
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
11. At the following prompt, enter **boot -s** to start up the system in single-user mode.  
  
ok **boot -s**
12. At the following prompt, enter **recovery** to start the root password recovery procedure.  
  
Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: **recovery**
13. Enter configuration mode in the CLI.
14. Set the root password. For example:  
  
user@switch# **set system root-authentication plain-text-password**
15. At the following prompt, enter the new root password. For example:  
  
New password: **juniper1**  
Retype new password:
16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.  
  
root@host# **commit**  
commit complete
18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the switch.  
  
Reboot the system? [y/n] **y**

**Related  
Documentation**

- [Configuring the Root Password on page 1702](#)



## PART 22

# Troubleshooting

- [Overview on page 6483](#)
- [Administration on page 6493](#)
- [Troubleshooting on page 6505](#)



## CHAPTER 71

# Overview

- [General Troubleshooting on page 6483](#)
- [Alarms on page 6487](#)

## General Troubleshooting

---

- [Understanding Troubleshooting Resources on page 6483](#)
- [Troubleshooting Overview on page 6485](#)

## Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

[Table 632 on page 6471](#) provides a list of some of the troubleshooting resources.

**Table 634: Troubleshooting Resources on the QFX Series**

| Troubleshooting Resource              | Description                                                                                                                                               | Documentation                                                             |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| Chassis alarms                        | Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch. | <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 6488</a> |
| Chassis Status LEDs and Fan Tray LEDs | A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.                | <i>Chassis Status LEDs on a QFX3500 Device</i>                            |
| Interface alarms                      | A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.                                           | <a href="#">“Interface Alarm Messages” on page 6491</a>                   |
| System alarms                         | A predefined alarm is triggered by a missing rescue configuration or problem with the software license.                                                   | <a href="#">“Understanding Alarms” on page 6487</a>                       |

Table 634: Troubleshooting Resources on the QFX Series (*continued*)

| Troubleshooting Resource                      | Description                                                                                                                                                                                                                                                                                                                                                                                                | Documentation                                                                                                                                                                                                                                                    |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System log messages                           | The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6155</a></li> <li>• <a href="#">Junos OS System Log Configuration Statements on page 6208</a></li> </ul>                                    |
| Junos OS operational mode commands            | Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the <b>traceroute monitor</b> command to locate points of failure in a network.                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 315</a></li> <li>• <a href="#">Monitoring System Properties on page 316</a></li> <li>• <a href="#">traceroute monitor</a></li> </ul>                          |
| Junos OS automation scripts (event scripts)   | Event scripts can be used to automate network troubleshooting and management tasks.                                                                                                                                                                                                                                                                                                                        | <i>Junos OS Configuration and Operations Automation Guide</i>                                                                                                                                                                                                    |
| Junos OS XML operational tags                 | XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.                                                                                                                                                                                                                                                    | <i>Junos XML API Operational Developer Reference</i>                                                                                                                                                                                                             |
| NETCONF XML management protocol               | The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations. | <i>NETCONF XML Management Protocol Guide</i>                                                                                                                                                                                                                     |
| SNMP MIBs and traps                           | MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6124</a></li> <li>• <a href="#">SNMP Traps Support on page 6139</a></li> <li>• <a href="#">Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS</a></li> </ul> |
| AI-Scripts and Advanced Insight Manager (AIM) | AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.                                                                                                                                                   | <a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>                                                                                                                                                                                              |
| Junos Space Service Now                       | This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.                                                                                                                         | <i>Service Automation</i>                                                                                                                                                                                                                                        |



Table 634: Troubleshooting Resources on the QFX Series (*continued*)

| Troubleshooting Resource        | Description                                                                                                                                                                                                                                                                                                                                 | Documentation                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| Junos Space Service Insight     | This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now. | <i>Service Automation</i>                                 |
| Juniper Networks Knowledge Base | You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.                                                                                                                                                                                                                        | <a href="http://kb.juniper.net">http://kb.juniper.net</a> |

## Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series product.

[Table 633 on page 6473](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 635: Troubleshooting on the QFX Series

| Problem Category           | Symptom or Problem                                                          | Recommended Action                                                                                     |
|----------------------------|-----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Switch hardware components | LCD panel shows a chassis alarm count.                                      | See <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 6488</a> .                        |
|                            | Fan tray LED is blinking amber.                                             | See <i>Fan Tray LED on a QFX3500 Device</i> .                                                          |
|                            | Chassis status LED for the power is blinking amber.                         | See <i>Chassis Status LEDs on a QFX3500 Device</i> .                                                   |
|                            | Chassis status LED for the fan (on the management board) is blinking amber. | Replace the management board as soon as possible. See <i>Chassis Status LEDs on a QFX3500 Device</i> . |

Table 635: Troubleshooting on the QFX Series (*continued*)

| Problem Category               | Symptom or Problem                                                                                          | Recommended Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port configuration             | Cannot configure a port as a Gigabit Ethernet port.                                                         | <p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                | Cannot configure a port as a Fibre Channel port.                                                            | <p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                                | Cannot configure a port as a 10-Gigabit Ethernet port.                                                      | <p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p> |
|                                | Cannot configure a 40-Gbps QSFP+ interface.                                                                 | <p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <a href="#">“QFX3500 Device Overview” on page 3</a>.</p>                                                                                                                                                                                                                                                                                                                         |
| External devices (USB devices) | Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state. | Unplug the USB device and reboot the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Initial device configuration   | Cannot configure management Ethernet ports.                                                                 | <p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p><b>NOTE:</b> The management ports are on the front panel of the QFX3500 switch. They are labeled <b>C0</b> and <b>C1</b> on the front panel. In the CLI they are referred to as <b>me0</b> and <b>me1</b>.</p> <p>See <a href="#">“Configuring a QFX3500 Device as a Standalone Switch” on page 163</a>.</p>                                                                                                                                                                                                                                                  |

Table 635: Troubleshooting on the QFX Series (*continued*)

| Problem Category                   | Symptom or Problem                                                                                                                                                    | Recommended Action                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Software upgrade and configuration | Failed software upgrade.                                                                                                                                              | See <a href="#">“Recovering from a Failed Software Installation” on page 119.</a>                                                                                                                                                                                                                                                                                                                                                                   |
|                                    | Active partition becomes inactive after upgrade.                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                    | Problem with the active configuration file.                                                                                                                           | See the following topics: <ul style="list-style-type: none"> <li>• <a href="#">Loading a Previous Configuration File on page 1600</a></li> <li>• <a href="#">Reverting to the Default Factory Configuration on page 173</a></li> <li>• <a href="#">Reverting to the Rescue Configuration on page 175</a></li> <li>• <a href="#">Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111</a></li> </ul> |
|                                    | Root password is lost or forgotten.                                                                                                                                   | Recover the root password. See <a href="#">“Recovering the Root Password” on page 1159.</a>                                                                                                                                                                                                                                                                                                                                                         |
| Network interfaces                 | An aggregated Ethernet interface is down.                                                                                                                             | See <a href="#">“Troubleshooting an Aggregated Ethernet Interface” on page 1161.</a>                                                                                                                                                                                                                                                                                                                                                                |
|                                    | Interface on built-in network port is down.                                                                                                                           | See <a href="#">“Troubleshooting Network Interfaces” on page 1160.</a>                                                                                                                                                                                                                                                                                                                                                                              |
|                                    | Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Ethernet switching                 | A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. | See <a href="#">“Troubleshooting Ethernet Switching” on page 2261.</a>                                                                                                                                                                                                                                                                                                                                                                              |
| Firewall filter                    | Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.                                                                            | See <a href="#">“Troubleshooting Firewall Filter Configuration” on page 4873.</a>                                                                                                                                                                                                                                                                                                                                                                   |

## Alarms

- [Understanding Alarms on page 6487](#)
- [Chassis Alarm Messages on a QFX3500 Device on page 6488](#)
- [Interface Alarm Messages on page 6491](#)

## Understanding Alarms

The QFX Series support different alarm types and severity levels. [Table 636 on page 6488](#) provides a list of alarm terms and definitions that may help you in monitoring the device.

Table 636: Alarm Terms and Definitions

| Term                  | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm                 | Signal alerting you to conditions that might prevent normal operation. On the device, alarm indicators include the LCD panel and LEDs on the front. The LCD panel displays the chassis alarm message count. Blinking amber LEDs indicate yellow alarm conditions for chassis components.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Alarm condition       | Failure event that triggers an alarm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Alarm severity levels | <p>Seriousness of the alarm. The level of severity can be either major (red) or minor (yellow).</p> <ul style="list-style-type: none"> <li>Major (red)—Indicates a critical situation on the device that has resulted from one of the following conditions. A red alarm condition requires immediate action. <ul style="list-style-type: none"> <li>One or more hardware components have failed.</li> <li>One or more hardware components have exceeded temperature thresholds.</li> <li>An alarm condition configured on an interface has triggered a critical warning.</li> </ul> </li> <li>Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance. For example, a missing rescue configuration generates a yellow system alarm.</li> </ul> |
| Alarm types           | <p>Alarms include the following types:</p> <ul style="list-style-type: none"> <li>Chassis alarm—Predefined alarm triggered by a physical condition on the device such as a power supply failure or excessive component temperature.</li> <li>Interface alarm—Alarm you configure to alert you when an interface link is down. Applies to <b>ethernet</b>, <b>fibre-channel</b>, and <b>management-ethernet</b> interfaces. You can configure a red (major) or yellow (minor) alarm for the link-down condition, or have the condition ignored.</li> <li>System alarm—Predefined alarm triggered by a missing rescue configuration or failure to install a license for a licensed software feature.</li> </ul>                                                                                                                                                                                                             |

#### Related Documentation

- [Chassis Alarm Messages on a QFX3008-I Interconnect Device](#)
- [Chassis Alarm Messages on a QFX3500 Device on page 6488](#)
- [Interface Alarm Messages on page 6491](#)

## Chassis Alarm Messages on a QFX3500 Device

Chassis alarms indicate a failure on the device or one of its components. Chassis alarms are preset and cannot be modified.

The chassis alarm message count is displayed on the LCD panel on the front of the device. To view the chassis alarm message text remotely, use the **show chassis lcd** CLI command.

Chassis alarms on QFX3500 devices have two severity levels:

- Major (red)—Indicates a critical situation on the device that has resulted from one of the conditions described in [Table 637 on page 6489](#). A red alarm condition requires immediate action.

- Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

Table 637 on page 6489 describes the chassis alarm messages on QFX3500 devices.

**Table 637: QFX3500 Chassis Alarm Messages**

| Component | Alarm Type  | CLI Message                               | Recommended Action                                                                                                                                                                                                                                                                                                        |
|-----------|-------------|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fans      | Major (red) | <b>Fan/Blower Absent</b>                  | The fan is missing. Install a fan.                                                                                                                                                                                                                                                                                        |
|           |             | <b>Fan Failure</b>                        | Replace the fan and report the failure to customer support.                                                                                                                                                                                                                                                               |
|           |             | <b>Fan I2C Failure</b>                    | Check the system log for one of the following messages and report the error message to customer support: <ul style="list-style-type: none"> <li>• <b>CM ENV Monitor: Get fan speed failed.</b></li> <li>• <b><i>Fan-number</i> is NOT spinning @ correct speed</b>, where <i>fan-number</i> may be 1, 2, or 3.</li> </ul> |
|           |             | <b>Fan <i>fan-number</i> Not Spinning</b> | Remove and check the fan for obstructions, and then reinsert the fan. If the problem persists, replace the fan.                                                                                                                                                                                                           |

Table 637: QFX3500 Chassis Alarm Messages (*continued*)

| Component      | Alarm Type     | CLI Message                                                       | Recommended Action                                                                                                                                                                                                                                                                                                                                                                   |
|----------------|----------------|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Power supplies | Major (red)    | <b>PEM <i>pem-number</i> Airflow not matching Chassis Airflow</b> | The power supply airflow direction is the opposite of the chassis airflow direction. Replace the power supply with a power supply that supports the same airflow direction as the chassis.                                                                                                                                                                                           |
|                |                | <b>PEM <i>pem-number</i> I2C Failure</b>                          | Check the system log for one of the following messages and report the error message to customer support: <ul style="list-style-type: none"> <li><b>I2C Read failed for device <i>number</i>,</b> where <i>number</i> may be from 123 to 125.</li> <li><b>PS <i>number</i>: Transitioning from online to offline,</b> where power supply (PS) <i>number</i> may be 1 or 2.</li> </ul> |
|                |                | <b>PEM <i>pem-number</i> is not powered</b>                       | For information only. Check the power cord connection and reconnect it if necessary.                                                                                                                                                                                                                                                                                                 |
|                |                | <b>PEM <i>pem-number</i> is not supported</b>                     | Indicates a power supply problem, or the power supply is not supported on the device. Report the problem to customer support.                                                                                                                                                                                                                                                        |
|                |                | <b>PEM <i>pem-number</i> Not OK</b>                               | Indicates a problem with the incoming AC or outgoing DC power. Replace the power supply.                                                                                                                                                                                                                                                                                             |
|                | Minor (yellow) | <b>PEM <i>pem-number</i> Absent</b>                               | For information only. Indicates the device was powered on with two power supplies installed, but now one is missing. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.                                                                                                               |
|                |                | <b>PEM <i>pem-number</i> Power Supply Type Mismatch</b>           | For information only. Indicates that an AC power supply and DC power supply have been installed in the same chassis. If you wish to remove this alarm message, reboot the device with two AC power supplies or two DC power supplies.                                                                                                                                                |
|                |                | <b>PEM <i>pem-number</i> Removed</b>                              | For information only. Indicates the device was powered on with two power supplies installed, but one has been removed. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.                                                                                                             |

Table 637: QFX3500 Chassis Alarm Messages (*continued*)

| Component           | Alarm Type     | CLI Message                                 | Recommended Action                                                                                                                                                                                                                          |
|---------------------|----------------|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Temperature sensors | Major (red)    | <i>sensor-location</i> Temp Sensor Fail     | Check the system log for the following message and report it to customer support:<br><br><b>Temp sensor <i>sensor-number</i> failed</b> , where <i>sensor-number</i> may range from 1 through 10.                                           |
|                     |                | <i>sensor-location</i> Temp Sensor Too Hot  | Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor. If the condition persists, the device may shut down. |
|                     | Minor (yellow) | <i>sensor-location</i> Temp Sensor Too Warm | For information only. Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor.                                |

**Related Documentation**

- [Front Panel of a QFX3500 Device](#)
- [Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types on page 147](#)
- [alarm on page 2554](#)

## Interface Alarm Messages

Interface alarms are alarms that you configure to alert you when an interface is down.

To configure an interface link-down condition to trigger a red or yellow alarm, or to configure the link-down condition to be ignored, use the **alarm** statement at the **[edit chassis]** hierarchy level. You can specify the **ethernet**, **fibre-channel**, or **management-ethernet** interface type.



**NOTE:** Fibre Channel alarms are only valid on QFX3500 devices.

By default, major alarms are configured for interface link-down conditions on the control plane and management network interfaces in a QFabric system. The link-down alarms indicate that connectivity to the control plane network is down. You can configure these alarms to be ignored using the **alarm** statement at the **[edit chassis]** hierarchy level.



NOTE: If you configure a yellow alarm on the QFX3008-I Interconnect device, it will be handled as a red alarm.

**Related  
Documentation**

- [Understanding Alarms on page 6487](#)



## CHAPTER 72

# Administration

- [Routine Monitoring Using the CLI on page 6493](#)

### Routine Monitoring Using the CLI

---

- [Monitoring SNMP on page 6493](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 6495](#)
- [Monitoring RMON MIB Tables on page 6498](#)
- [Displaying a Log File from a Single-Chassis System on page 6499](#)
- [Monitoring System Log Messages on page 6500](#)
- [Monitoring Traffic Through the Router or Switch on page 6501](#)
- [Pinging Hosts on page 6503](#)

### Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
```

```
Alarm
```

| Index | Variable description                                                         | Value | State  |
|-------|------------------------------------------------------------------------------|-------|--------|
| 32768 | Health Monitor: root file system utilization<br>jnxHrStoragePercentUsed.1    | 58    | active |
| 32769 | Health Monitor: /config file system utilization<br>jnxHrStoragePercentUsed.2 | 0     | active |

```
32770 Health Monitor: RE 0 CPU utilization
 jnxOperatingCPU.9.1.0.0 0 active

32773 Health Monitor: RE 0 Memory utilization
 jnxOperatingBuffer.9.1.0.0 35 active

32775 Health Monitor: jkernel daemon CPU utilization
 Init daemon 0 active
 Chassis daemon 50 active
 Firewall daemon 0 active
 Interface daemon 5 active
 SNMP daemon 11 active
 MIB2 daemon 42 active
 ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system
```

```
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx
```

```
Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 24444184
sysContact.0 = J Smith
sysName.0 = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics
```

```
SNMP statistics:
 Input:
 Packets: 0, Bad versions: 0, Bad community names: 0,
 Bad community uses: 0, ASN parse errors: 0,
 Too bigs: 0, No such names: 0, Bad values: 0,
 Read onlys: 0, General errors: 0,
 Total request varbinds: 0, Total set varbinds: 0,
 Get requests: 0, Get nexts: 0, Set requests: 0,
 Get responses: 0, Traps: 0,
 Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
 Throttle drops: 0, Duplicate request drops: 0
 Output:
 Packets: 0, Too bigs: 0, No such names: 0,
 Bad values: 0, General errors: 0,
 Get requests: 0, Get nexts: 0, Set requests: 0,
 Get responses: 0, Traps: 0
```

- Related Documentation
- [health-monitor on page 1773](#)
  - [show snmp mib on page 6452](#)
  - [show snmp statistics on page 1859](#)

## Tracing SNMP Activity on a Device Running Junos OS

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
  - chassisd
  - craftd
  - ilmid
  - mib2d
  - rmopd
  - serviced
  - snmpd
- When a trace file named **filename** reaches its maximum size, it is renamed **filename.0**, then **filename.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
 file <files number> <match regular-expression> <size size> <world-readable |
 no-world-readable>;
 flag flag;
 no-remote-trace;
}
```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 6496](#)
- [Configuring Access to the Log File on page 6496](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 6496](#)
- [Configuring the Trace Operations on page 6496](#)

### Configuring the Number and Size of SNMP Log Files

---

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed *filename.0*, then *filename.1*, and so on, until there are three trace files. Then the oldest trace file (*filename.2*) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

### Configuring Access to the Log File

---

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

### Configuring a Regular Expression for Lines to Be Logged

---

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

### Configuring the Trace Operations

---

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
 all;
```

```

configuration;
database;
events;
general;
interface-stats;
nonvolatile-sets;
pdu;
policy;
protocol-timeouts;
routing-socket;
server;
subagent;
timer;
varbind-error;
}

```

Table 613 on page 6395 describes the meaning of the SNMP tracing flags.

**Table 638: SNMP Tracing Flags**

| Flag                     | Description                                                                 | Default Setting |
|--------------------------|-----------------------------------------------------------------------------|-----------------|
| <b>all</b>               | Log all operations.                                                         | Off             |
| <b>configuration</b>     | Log reading of the configuration at the <b>[edit snmp]</b> hierarchy level. | Off             |
| <b>database</b>          | Log events involving storage and retrieval in the events database.          | Off             |
| <b>events</b>            | Log important events.                                                       | Off             |
| <b>general</b>           | Log general events.                                                         | Off             |
| <b>interface-stats</b>   | Log physical and logical interface statistics.                              | Off             |
| <b>nonvolatile-set</b>   | Log nonvolatile SNMP set request handling.                                  | Off             |
| <b>pdu</b>               | Log SNMP request and response packets.                                      | Off             |
| <b>policy</b>            | Log policy processing.                                                      | Off             |
| <b>protocol-timeouts</b> | Log SNMP response timeouts.                                                 | Off             |
| <b>routing-socket</b>    | Log routing socket calls.                                                   | Off             |
| <b>server</b>            | Log communication with processes that are generating events.                | Off             |
| <b>subagent</b>          | Log subagent restarts.                                                      | Off             |
| <b>timer</b>             | Log internal timer events.                                                  | Off             |

Table 638: SNMP Tracing Flags (*continued*)

| Flag                 | Description                  | Default Setting |
|----------------------|------------------------------|-----------------|
| <b>varbind-error</b> | Log variable binding errors. | Off             |

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
  - [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
  - [Example: Tracing SNMP Activity](#)
  - [Configuring SNMP on page 1703](#)

## Monitoring RMON MIB Tables

**Purpose** Monitor remote monitoring (RMON) alarm, event, and log tables.

**Action** To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index Variable description Value State

5 monitor
 jnxOperatingCPU.9.1.0.0 5 falling threshold

Event
Index Type Last Event

1 log and trap 2010-07-10 11:34:17 PDT
Event Index: 1
 Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
 Time: 2010-07-10 11:34:07 PDT
 Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
 Time: 2010-07-10 11:34:17 PDT
```

**Meaning** The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

- Related Documentation**
- [Configuring RMON Alarms and Events on page 6198](#)
  - [show snmp rmon on page 6455](#)

- [show snmp rmon history on page 6459](#)
- [clear snmp statistics on page 6403](#)
- [clear snmp history on page 6402](#)

## Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system such as the QFX3500 switch, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysAppElmtRunMemory.5.6.2293)
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysAppElmtRunMemory.5.8.2292)
...
Nov 4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov 4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov 4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
```

```
SNMP trap: cold start
...
```

- Related Documentation**
- [Interpreting Messages Generated in Standard Format on page 6219](#)
  - [Interpreting Messages Generated in Structured-Data Format on page 6216](#)

## Monitoring System Log Messages

**Purpose** Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

**Action** To view system log messages:

```
user@switch1> show log messages
```

## Sample Output

```
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov 4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApp1ElmtRunMemory.5.6.2293)
Nov 4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApp1ElmtRunMemory.5.8.2292)
...
Nov 4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

**Meaning** The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of 128 kilobytes (KB).
- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user **jsmith**.



## Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 6501](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 6501](#)

### Displaying Real-Time Statistics About All Interfaces on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through all interfaces on the router or switch.

**Action** To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

### Sample Output

```
user@host> monitor interface traffic
host name Seconds: 15 Time: 12:31:09
Interface Link Input packets (pps) Output packets (pps)
so-1/0/0 Down 0 (0) 0 (0)
so-1/1/0 Down 0 (0) 0 (0)
so-1/1/1 Down 0 (0) 0 (0)
so-1/1/2 Down 0 (0) 0 (0)
so-1/1/3 Down 0 (0) 0 (0)
t3-1/2/0 Down 0 (0) 0 (0)
t3-1/2/1 Down 0 (0) 0 (0)
t3-1/2/2 Down 0 (0) 0 (0)
t3-1/2/3 Down 0 (0) 0 (0)
so-2/0/0 Up 211035 (1) 36778 (0)
so-2/0/1 Up 192753 (1) 36782 (0)
so-2/0/2 Up 211020 (1) 36779 (0)
so-2/0/3 Up 211029 (1) 36776 (0)
so-2/1/0 Up 189378 (1) 36349 (0)
so-2/1/1 Down 0 (0) 18747 (0)
so-2/1/2 Down 0 (0) 16078 (0)
so-2/1/3 Up 0 (0) 80338 (0)
at-2/3/0 Up 0 (0) 0 (0)
at-2/3/1 Down 0 (0) 0 (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

**Meaning** The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

### Displaying Real-Time Statistics About an Interface on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through an interface on the router or switch.

**Action** To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

## Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
 Input bytes: 5856541 (88 bps)
 Output bytes: 6271468 (96 bps)
 Input packets: 157629 (0 pps)
 Output packets: 157024 (0 pps)
Encapsulation statistics:
 Input keepalives: 42353
 Output keepalives: 42320
 LCP state: Opened
Error statistics:
 Input errors: 0
 Input drops: 0
 Input framing errors: 0
 Input runs: 0
 Input giants: 0
 Policed discards: 0
 L3 incompletes: 0
 L2 channel errors: 0
 L2 mismatch timeouts: 0
 Carrier transitions: 1
 Output errors: 0
 Output drops: 0
 Aged packets: 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
 LOS count 1
 LOF count 1
 SEF count 1
 ES-S 77
 SES-S 77
SONET statistics:
 BIP-B1 0
 BIP-B2 0
 REI-L 0
 BIP-B3 0
 REI-P 0
Received SONET overhead: F1 : 0x00 J0 : 0xZ
```

**Meaning** The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 612 on page 6388](#).

Table 639: Output Control Keys for the monitor interface Command

| Action                                                                                                                                                                                                                       | Key      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Display information about the next interface. The <b>monitor interface</b> command scrolls through the physical or logical interfaces in the same order that they are displayed by the <b>show interfaces terse</b> command. | <b>N</b> |
| Display information about a different interface. The command prompts you for the name of a specific interface.                                                                                                               | <b>I</b> |
| Freeze the display, halting the display of updated statistics.                                                                                                                                                               | <b>F</b> |
| Thaw the display, resuming the display of updated statistics.                                                                                                                                                                | <b>T</b> |
| Clear (zero) the current delta counters since <b>monitor interface</b> was started. It does not clear the accumulative counter.                                                                                              | <b>C</b> |
| Stop the <b>monitor interface</b> command.                                                                                                                                                                                   | <b>Q</b> |

See the [CLI Explorer](#) for details on using match conditions with the **monitor traffic** command.

## Pinging Hosts

**Purpose** Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

**Action** To use the **ping** command to send four requests (ping count) to **host3**:  
**ping host count number**

## Sample Output

```
ping host3 count 4
user@switch> ping host3 count 4
PING host3.site.net (176.26.232.111): 56 data bytes
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

**Meaning**

- The **ping** results show the following information:
  - Size of the ping response packet (in bytes).
  - IP address of the host from which the response was sent.

- Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
- Time-to-live (ttl) hop-count value of the ping response packet.
- Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
- Number of ping requests (probes) sent to the host.
- Number of ping responses received from the host.
- Packet loss percentage.
- Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

## CHAPTER 73

# Troubleshooting

- [Configuration and File Management on page 6505](#)
- [Ethernet Switching on page 6508](#)
- [Hardware on page 6513](#)
- [High Availability on page 6515](#)
- [Interfaces on page 6516](#)
- [Junos OS Basics on page 6522](#)
- [Layer 3 Protocols on page 6536](#)
- [Security on page 6537](#)
- [Services on page 6547](#)
- [Storage on page 6550](#)
- [Traffic Management on page 6555](#)

### Configuration and File Management

---

- [Loading a Previous Configuration File on page 6505](#)
- [Reverting to the Default Factory Configuration on page 6506](#)
- [Reverting to the Rescue Configuration on page 6507](#)
- [Cleaning Up the System File Storage Space on page 6507](#)

### Loading a Previous Configuration File

You can use the **rollback <number>** command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

#### Syntax

**rollback <number>**

#### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.

- **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
- **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related Documentation**

- [Configuration File Terms on page 48](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

**Related Documentation**

- [Understanding Configuration Files on page 1590](#)
- [Loading a Previous Configuration File on page 1600](#)
- [Reverting to the Rescue Configuration on page 175](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a QFX Series product and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

### Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1609](#)
- [Reverting to the Default Factory Configuration on page 173](#)
- [Configuration File Terms on page 48](#)

## Cleaning Up the System File Storage Space

**Problem**    **Description:** The system file storage space on the switch is full. Rebooting the switch does not solve the problem.

The following error message is displayed during a typical operation on the switch after the file storage space is full.

```
user@switch% cli
user@switch> configure
/var: write failed, filesystem is full
```

**Solution**    Clean up the file storage on the switch by deleting system files.

1. Request to delete system files on the switch.

```
user@switch> request system storage cleanup
```

The list of files to be deleted is displayed.

List of files to delete:

|  | Size  | Date         | Name                               |
|--|-------|--------------|------------------------------------|
|  | 11B   | Jul 26 20:55 | /var/jail/tmp/alarmd.ts            |
|  | 124B  | Aug 4 18:05  | /var/log/default-log-messages.0.gz |
|  | 1301B | Jul 26 20:42 | /var/log/install.0.gz              |
|  | 387B  | Jun 3 14:37  | /var/log/install.1.gz              |
|  | 4920B | Aug 4 18:05  | /var/log/messages.0.gz             |
|  | 20.0K | Jul 26 21:00 | /var/log/messages.1.gz             |
|  | 16.3K | Jun 25 13:45 | /var/log/messages.2.gz             |
|  | 804B  | Aug 4 18:05  | /var/log/security.0.gz             |
|  | 16.8K | Aug 3 11:15  | /var/log/security.1.gz             |

```

487B Aug 4 18:04 /var/log/wtmp.0.gz
855B Jul 29 22:54 /var/log/wtmp.1.gz
920B Jun 30 16:32 /var/log/wtmp.2.gz
94B Jun 3 14:36 /var/log/wtmp.3.gz
353.2K Jun 3 14:37 /var/sw/pkg/jloader-qfx-11.2I20110303_1117_dc-builder.tgz

124.0K Jun 3 14:30 /var/tmp/gres-tp/env.dat
0B Apr 14 16:20 /var/tmp/gres-tp/lock
0B Apr 14 17:37 /var/tmp/if-rtbdb/env.lock
12.0K Jul 26 20:55 /var/tmp/if-rtbdb/env.mem
2688.0K Jul 26 20:55 /var/tmp/if-rtbdb/shm_usr1.mem
132.0K Jul 26 20:55 /var/tmp/if-rtbdb/shm_usr2.mem
2048.0K Jul 26 20:55 /var/tmp/if-rtbdb/trace.mem
155B Jul 26 20:55 /var/tmp/krt_gencfg_filter.txt
0B Jul 26 20:55 /var/tmp/rtbdb/if-rtbdb
1400.6K Aug 3 10:13 /var/tmp/sfid.core.0.gz
1398.9K Aug 3 17:01 /var/tmp/sfid.core.1.gz
Delete these files ? [yes,no] (no)

```

2. Enter **yes** to delete the files.

3. Reboot the switch.



**BEST PRACTICE:** We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch.

Related Documentation • [request system storage cleanup on page 423](#)

## Ethernet Switching

- [Troubleshooting Ethernet Switching on page 6508](#)
- [Troubleshooting Layer 2 Protocol Tunneling on page 6509](#)
- [Troubleshooting Private VLANs on page 6510](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 6513](#)

## Troubleshooting Ethernet Switching

**Problem Description:** Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by



issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

**Solution** Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

**Related Documentation**

- [arp on page 2072](#)
- [mac-table-aging-time on page 2101](#)

## Troubleshooting Layer 2 Protocol Tunneling

- [Drop Threshold Statistics Might Be Incorrect on page 6509](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 6509](#)

### Drop Threshold Statistics Might Be Incorrect

**Problem** **Description:** L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets will not be reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit are not reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

**Solution** This is expected behavior.

### Egress Filtering of L2PT Traffic Not Supported

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a

firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Layer 2 Protocol Tunneling on page 1910](#)
- [Configuring Layer 2 Protocol Tunneling on page 2063](#)

## Troubleshooting Private VLANs

Use the following information to troubleshoot a private VLAN configuration.

- [Limitations of Private VLANs on page 6510](#)
- [Forwarding with Private VLANs on page 6510](#)
- [Egress Firewall Filters with Private VLANs on page 6511](#)
- [Egress Port Mirroring with Private VLANs on page 6512](#)

### Limitations of Private VLANs

---

The following constraints apply to private VLAN configurations:

- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.

### Forwarding with Private VLANs

---

**Problem Description:**

- When isolated VLAN or community VLAN tagged traffic is received on a PVLAN trunk port, MAC addresses are learned from the primary VLAN. This means that output from the [show ethernet-switching table](#) command shows that MAC addresses are learned from the primary VLAN and replicated to secondary VLANs. This behavior has no effect on forwarding decisions.
- If a packet with a secondary VLAN tag is received on a promiscuous port, it is accepted and forwarded.
- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.
  - The packet has a community VLAN tag.
  - The packet is destined to a unicast MAC address or multicast group MAC address that was learned on an isolated VLAN.
- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.

- The packet has an isolated VLAN tag.
- The packet is destined to a unicast MAC address or multicast group MAC address that was learned on a community VLAN.
- If a packet with a primary VLAN tag is received by a secondary (isolated or community) VLAN port, the secondary port forwards the packet.
- If you configure a community VLAN on one device and configure another community VLAN on a second device and both community VLANs use the same VLAN ID, traffic for one of the VLANs can be forwarded to the other VLAN. For example, assume the following configuration:
  - Community VLAN comm1 on switch 1 has VLAN ID 50 and is a member of primary VLAN pvlan100.
  - Community VLAN comm2 on switch 2 also has VLAN ID 50 and is a member of primary VLAN pvlan200.
  - Primary VLAN pvlan100 exists on both switches.

If traffic for comm1 is sent from switch 1 to switch 2, it will be sent to the ports participating in comm2. (The traffic will also be forwarded to the ports in comm1, as you would expect.)

**Solution** These are expected behaviors.

### Egress Firewall Filters with Private VLANs

**Problem Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port

- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Port Mirroring with Private VLANs

---

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

- Related Documentation**
- [Understanding Private VLANs on page 1878](#)
  - [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS on page 1887](#)
  - [Creating a Private VLAN on a Single Switch on page 2057](#)
  - [Creating a Private VLAN Spanning Multiple Switches on page 2058](#)
  - [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports on page 2024](#)

## Troubleshooting Q-in-Q and VLAN Translation Configuration

- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 6513](#)
- [Egress Port Mirroring with VLAN Translation on page 6513](#)

### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

**Problem** **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same EtherType.

### Egress Port Mirroring with VLAN Translation

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

- Related Documentation**
- [Understanding Q-in-Q Tunneling and VLAN Translation on page 1906](#)
  - [Example: Setting Up Q-in-Q Tunneling on page 2036](#)

## Hardware

- [Troubleshooting QFX3100 Director Device Isolation on page 6514](#)

## Troubleshooting QFX3100 Director Device Isolation

**Problem**    **Description:** Both connections between the QFX3100 Director devices are broken so that one of the Director devices in a Director group becomes isolated from the group.

The redundant patch cables interconnecting the Director devices are critical links required for the operation of the Director group. The two inter-Director device links must remain connected when the Director devices are online. After the Director devices are installed and the Director group is active, if a single inter-Director device link loses and regains its connection, the operation of the Director group remains intact. However, the loss of both inter-Director device links causes one Director device to isolate itself from the Director group.



**WARNING:** Do not reconnect the inter-Director patch cables before properly restarting the isolated Director device. Restarting the active Director device instead of the isolated Director device can result in both Director devices rebooting, with a subsequent data loss.

**Environment:** This problem occurs between the two QFX3100 Director devices found in QFabric systems.

**Symptoms:** Symptoms of this problem include an unscheduled rebooting of one of the Director devices.

---

### Resolution    *Determine Which Director Device Is Isolated*

Before restoring the inter-Director device links, determine which one of the Director devices is in isolation.

To locate an isolated Director device, use one of the following methods:

- Review logs or management tools for standard SNMP traps issued from the Director group before the Director device became isolated.
  - If eth-2/6 links are down, the Director group cannot communicate. Normally, one of the devices reboots.
  - If both eth-2/6 and eth-7/8/9 links are down, the Director device is isolated from the control plane and is not providing fabric services.
  - Issue **show fabric session-host**.
- Use the CLI to determine the serial numbers of the active Director device.
  - Issue the **show fabric session-host** command.

```
root@qfabric>show fabric session-host
Identifier: 0281042010000013
```

- Issue the **show fabric administration inventory director-group status | grep “dg0|dg1”** command.

```
root@qfabrid> show fabric administration inventory director-group status | grep
“dg0|dg1”
```

```
dg0 online master 10.94.214.80 0% 13597976k 4 4 days, 22:36 hrs
dg1 online master 10.94.214.81 0% 18677380k 3 4 days, 22:25 hrs
dg0 0281042010000013 online master
dg1 0281042010000018 online backup
```

When the Director devices cannot communicate, the **show fabric administration inventory director-group** command only displays the Director device that is online.

### *Power Off the Isolated Director Device and Restore the Inter-Director Device Links*



**CAUTION:** Be sure you know which Director device is active and which is isolated. If you power off the active Director device, both Director devices reboot and cause potential data loss on the system.

To restore communication within the Director group:

1. Power off the isolated Director device.
2. Restore the inter-Director device links (port 3 to port 3) by firmly inserting the redundant patch cables.
3. Power on the previously isolated Director device. The Director device reboots.

#### **Related Documentation**

- *Connecting QFX3100 Director Devices in a Director Group*

## High Availability

- [Troubleshooting VRRP on page 6515](#)

## Troubleshooting VRRP

**Problem** **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

**Solution** Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

**Related Documentation**

- [failover-delay on page 2326](#)

## Interfaces

---

- [Troubleshooting an Aggregated Ethernet Interface on page 6516](#)
- [Troubleshooting Network Interfaces on page 6516](#)
- [Troubleshooting Multichassis Link Aggregation on page 6517](#)

### Troubleshooting an Aggregated Ethernet Interface

**Problem**    **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution**    Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related Documentation**

- [Verifying the Status of a LAG Interface on page 2630](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2454](#)

### Troubleshooting Network Interfaces

**The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down**

---

**Problem**    **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command `show interfaces interface-name`, the disabled port is not listed.

**Cause**    By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution**    Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.



## Troubleshooting Multichassis Link Aggregation

Use the following information to troubleshoot multichassis link aggregation configuration.

- [MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table on page 6517](#)
- [MC-LAG Peer Does Not Go into Standby Mode on page 6518](#)
- [Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive on page 6518](#)
- [Redirect Filters Take Priority over User-Defined Filters on page 6518](#)
- [Operational Command Output Is Wrong on page 6519](#)
- [ICCP Connection Might Take Up to 60 Seconds to Become Active on page 6519](#)
- [MAC Address Age Learned on an MC-AE Interface Is Reset to Zero on page 6519](#)
- [MAC Address Is Not Learned Remotely in a Default VLAN on page 6520](#)
- [Snooping Entries Learned on MC-AE Interfaces Are Not Removed on page 6520](#)
- [ICCP Does Not Come Up After You Add or Delete an Authentication Key on page 6520](#)
- [Local Status Is Standby When It Should Be Active on page 6520](#)
- [Packets Loop on the Server When ICCP Fails on page 6520](#)
- [Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change on page 6521](#)
- [No Commit Checks Are Done for ICL-PL Interfaces on page 6521](#)
- [Double Failover Scenario on page 6521](#)
- [Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up on page 6521](#)
- [Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer on page 6522](#)
- [AE Interfaces Go Down on page 6522](#)
- [Flooding of Upstream Traffic on page 6522](#)

### MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table

**Problem Description:** When both of the multichassis aggregated Ethernet (MC-AE) interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the MC-AE interfaces are not removed from the MAC address table. For example, if you disable the MC-AE interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the MC-AE interfaces of both MC-LAG peers:

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
 VLAN MAC address Type Age Interfaces
 --- -
 v10 * Flood - All-members
 v10 00:10:94:00:00:01 Learn(L) 3:55 ae0.0 (MCAE)
```

|     |                   |          |               |
|-----|-------------------|----------|---------------|
| v10 | 00:10:94:00:00:02 | Learn(R) | 0 xe-0/0/9.0  |
| v20 | *                 | Flood    | - All-members |
| v30 | *                 | Flood    | - All-members |
| v30 | 84:18:88:de:b1:2e | Static   | - Router      |

```
user@switchB> show ethernet-switching table
```

```
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
```

| VLAN | MAC address       | Type     | Age | Interfaces    |
|------|-------------------|----------|-----|---------------|
| v10  | *                 | Flood    |     | - All-members |
| v10  | 00:10:94:00:00:01 | Learn(R) | 0   | ae0.0 (MCAE)  |
| v10  | 00:10:94:00:00:02 | Learn    | 40  | xe-0/0/10.0   |
| v20  | *                 | Flood    |     | - All-members |
| v30  | *                 | Flood    |     | - All-members |
| v30  | 84:18:88:df:83:0a | Static   |     | - Router      |

**Solution** This is expected behavior.

#### MC-LAG Peer Does Not Go into Standby Mode

**Problem** **Description:** A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Interchassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

**Solution** To prevent failure to enter standby mode, make sure the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

#### Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet (MC-AE) interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's MC-AE interfaces with status control set to standby become inactive instead of active.

**Solution** This is expected behavior.

#### Redirect Filters Take Priority over User-Defined Filters

**Problem** **Description:** Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters. This is expected behavior.

**Solution** This is expected behavior.

### Operational Command Output Is Wrong

**Problem Description:** After you deactivate the Interchassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
Redundancy Group IDs Joined: None
```

```
Client Application: lacpd
Redundancy Group IDs Joined: 1
```

```
Client Application: eswd
Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

**Solution** This is expected behavior.

### ICCP Connection Might Take Up to 60 Seconds to Become Active

**Problem Description:** When the Interchassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

**Solution** This is expected behavior.

### MAC Address Age Learned on an MC-AE Interface Is Reset to Zero

**Problem Description:** When you activate and then deactivate an interchassis control link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet (MC-AE) interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine (PFE). The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
VLAN MAC address Type Age Interfaces
v100 * Flood - All-members
v100 00:10:00:00:00:01 Learn(L) 0 ae0.0 (MCAE)
v100 00:10:00:00:00:02 Learn(L) 0 ae0.0 (MCAE)
```

**Solution** This is expected behavior.

### MAC Address Is Not Learned Remotely in a Default VLAN

**Problem**    **Description:** If a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, the Interchassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

**Solution**    This is expected behavior.

### Snooping Entries Learned on MC-AE Interfaces Are Not Removed

**Problem**    **Description:** When multichassis aggregated Ethernet (MC-AE) interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the MC-AE interfaces on the VLAN are not cleared when the MC-AE interfaces go down. This is done to speed up convergence time when the interfaces come up, or come up and go down.

**Solution**    This is expected behavior.

### ICCP Does Not Come Up After You Add or Delete an Authentication Key

**Problem**    **Description:** The Interchassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

**Solution**    Delete the ICCP configuration , and then add the ICCP configuration.

### Local Status Is Standby When It Should Be Active

**Problem**    **Description:** If the multichassis aggregated Ethernet (MC-AE) interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the MC-AE interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

**Solution**    This is expected behavior.

### Packets Loop on the Server When ICCP Fails

**Problem**    **Description:** When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

**Solution**    This is expected behavior.

### Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

---

**Problem**    **Description:** After a reboot or after a new Interchassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis aggregated Ethernet (MC-AE) interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

**Solution**    This is expected behavior.

### No Commit Checks Are Done for ICL-PL Interfaces

---

**Problem**    **Description:** There are no commit checks on the interface being configured as an interchassis control link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

**Solution**    This is expected behavior.

### Double Failover Scenario

---

**Problem**    **Description:** If the following events happen in this exact order—the Interchassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet (MC-AE) interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the MC-AE interface on the MC-LAG in active mode were up and blocks the interchassis control protocol-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

**Solution**    This is expected behavior.

### Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

---

**Problem**    **Description:** When the interchassis control link-protection link (ICL-PL) goes down and up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine (PFE) flag `Ip4McastFloodMode` for the VLAN is changed to `MCAST_FLOOD_ALL`. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

**Solution**    This is expected behavior.

### Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

---

**Problem**    **Description:** When the Interchassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

**Solution**    This is expected behavior.

### AE Interfaces Go Down

---

**Problem**    **Description:** When a multichassis aggregated Ethernet (MC-AE) interface is converted to an aggregated Ethernet (AE) interface, it retains some MC-AE properties. For example, the AE interface might retain the administrative key of the MC-AE. When this happens, the AE interface goes down.

**Solution**    Restart the Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the AE interface to bring up the AE interface. Restarting LACP removes the MC-AE properties of the AE interface.

### Flooding of Upstream Traffic

---

**Problem**    **Description:** When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

**Solution**    Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2435](#)
- [Example: Configuring Multichassis Link Aggregation on page 2462](#)
- [Configuring Multichassis Link Aggregation on page 2540](#)

## Junos OS Basics

---

- [Rebooting and Halting a QFX Series Product on page 6523](#)
- [Recovering from a Failed Software Installation on page 6524](#)
- [Recovering the Root Password on page 6524](#)
- [Creating an Emergency Boot Device for a QFX Series Device on page 6526](#)

- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 6528](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 6529](#)

## Rebooting and Halting a QFX Series Product

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

```
user@switch> request system reboot ?
```

Possible completions:

|           |                                             |
|-----------|---------------------------------------------|
| <[Enter]> | Execute this command                        |
| at        | Time at which to perform the operation      |
| in        | Number of minutes to delay before operation |
| media     | Boot media for next boot                    |
| message   | Message to display to all users             |
|           | Pipe through a command                      |

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

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



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

```
user@switch> request system halt ?
```

Possible completions:

|           |                                             |
|-----------|---------------------------------------------|
| <[Enter]> | Execute this command                        |
| at        | Time at which to perform the operation      |
| in        | Number of minutes to delay before operation |
| media     | Boot media for next boot                    |
| message   | Message to display to all users             |
|           | Pipe through a command                      |



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

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

### Related Documentation

- [clear system reboot on page 332](#)
- [request system reboot on page 390](#)
- [request system halt on page 375](#)
- [request system power-off on page 385](#)
- [Connecting a QFX Series Device to a Management Console](#)

## Recovering from a Failed Software Installation

**Problem**    **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution**    If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Recovering the Root Password

If you forget the root password for the QFX3500 switch, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.

---





**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the QFX3500 switch power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the switch into the RJ-45-to-DB-9 serial port adapter supplied with the switch.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the switch.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the switch by (if necessary) plugging the power cords into the QFX3500 switch power supply, or turning on the power to the device or switch by switching on the AC power outlet the device is plugged into.
 

The terminal emulation screen on your management device displays the switch's boot sequence.
10. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:
 

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```
11. At the following prompt, enter **boot -s** to start up the system in single-user mode.
 

```
ok boot -s
```
12. At the following prompt, enter **recovery** to start the root password recovery procedure.
 

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN
for /bin/sh: recovery
```
13. Enter configuration mode in the CLI.
14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the switch.

```
Reboot the system? [y/n] y
```

#### Related Documentation

- [Configuring the Root Password on page 1702](#)

## Creating an Emergency Boot Device for a QFX Series Device

If Junos OS on the QFX Series is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



**NOTE:** In the following procedure, we assume that you are creating the emergency boot device on a QFX3500 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device from a QFX3500 device:

1. Use FTP to copy the installation media image into the **/var/tmp** directory on the QFX3500 device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the **su** command:

```
% su
Password: password
```



**NOTE:** The password is the root password for the QFX3500 device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

#### Related Documentation

- [USB Port Specifications for the QFX Series](#)
- [Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device on page 111](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 112](#)

## Performing a Recovery Installation on a QFX3008-I, QFX3600-I, QFX3600, or QFX3500 Device

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device for a QFX Series Device” on page 165](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



**WARNING:** The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the QFX Series device.
2. Reboot the QFX Series device.



**NOTE:** Do not power off the device if it is already on.

```
[edit system]
user@device> request system reboot
```

If you do not have access to the CLI, power cycle the QFX Series device.

The emergency boot device (external USB install media) is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following options:

```
External USB install media detected.
You can load Junos from this media onto an internal drive.
Press 'y' to proceed, 'f' to format and install, or 'n' to abort.
Do you wish to continue ([y]/f/n)? f
```

4. Type **f** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.

If you do not want to format the internal flash storage, type **y**.

The following messages are displayed:

```
Installing packages from external USB drive da1
Packages will be installed to da0, media size: 8G
```

```
Processing format options
Fri September 4 01:18:44 UTC 2012
```

```
-- IMPORTANT INFORMATION --
Installer has detected settings to format system boot media.
This operation will erase all data from your system.
```

```
Formatting installation disk .. this will take a while, please wait
Disabling platform watchdog - threshold 12 mins
```

```
Determining installation slice
Fri September 4 01:27:07 UTC 2012
```

5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the device is finished copying the software, you are presented with the following prompt:

```
*** Fri September 4 01:19:00 UTC 2012***
Installation successful..
Please select one of the following options:
Reboot to installed Junos after removing install media (default) ... 1
Reboot to installed Junos by disabling install media 2
Exit to installer debug shell 3
Install Junos to alternate slice 4
Your choice: 4
NOTE: System installer will now install Junos to alternate slice
Do not power off or remove the external installer media or
interrupt the installation mechanism.
```

6. Select **4** to install Junos OS to the alternate slice of the partition, and then press Enter.
7. Remove the emergency boot device when prompted and then press Enter. The device then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the login prompt.
8. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.

#### Related Documentation

- [Creating an Emergency Boot Device for a QFX Series Device on page 165](#)

## Performing a QFabric System Recovery Installation on the Director Group

If the software on your QFabric system is damaged in some way that prevents the software from loading correctly, or you need to upgrade the software on your QFabric system, you may need to perform a recovery installation on the Director group.

If possible, perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device (for example, an external USB flash drive) for each of your Director devices to use during the recovery installation.

You can either use the external USB flash drive containing the software supplied by Juniper Networks, or you can use an external USB flash drive supplied by Juniper Networks on which you install the QFabric system install media.

2. Because the recovery installation process completely overwrites the entire contents of the Director device, make sure you back up any configuration files and initial setup information on a different external USB flash drive before you begin a recovery installation. You will need to restore this information as part of recovery process.

Use the **request system software configuration-backup** command to back up your configuration files and initial setup information:

```
user@switch> request system software configuration-backup path
```



**NOTE:** To recover the Director group, you must upgrade both Director devices in parallel. If you are recovering only one Director device in a Director group, and the software version will remain the same between the two Director devices, make sure that the other Director device is powered on and operational. If the software version of the Director device you are recovering will be different, make sure that the other Director device is powered off and is not operational.

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive on page 6530
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software on page 6532

#### (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive

---

If you do not have an external USB flash drive preloaded with the software from Juniper Networks to use as an emergency boot device, you can create your own, using a blank external USB flash drive provided by Juniper Networks. Download the install media from the Juniper Networks Support website onto your UNIX workstation, uncompress and untar the software, and then burn the software image onto your Juniper Networks external USB (4-gigabyte) flash drive. Make sure you create two emergency boot devices, one for each Director device, so you can perform a recovery installation in parallel.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.
4. In the **QFX Series** section, click the name of the platform for which you want to download software.

5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. Select the complete install media you want to download in the **QFabric System Install Media** section.  
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Log in and save the install media file to your UNIX workstation.
10. Use FTP to access the UNIX workstation where the install media resides.  
**ftp ftp://hostname/pathname install-media-qfabric-<version>.img.tgz**
11. When prompted, enter your username and password.
12. Make sure you are in binary mode by entering **binary** at the prompt.  
**binary**
13. Use the **get** command to transfer the installation package from the FTP host to your UNIX workstation.  
**get install-media-qfabric-<version>.img.tgz**
14. Close the FTP session:  
**bye**
15. Untar the *install-media-qfabric-<version>.img.tgz* file on your UNIX workstation.  
**tar -xvzf install-media-qfabric-11.3X30.6.img.tgz**
16. Insert a blank external USB (4-gigabyte) flash drive supplied by Juniper Networks into your UNIX workstation.
17. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive using the **dd** command:  
**dd if=install-media-qfabric-11.3X30.6.img of=/dev/sdb bs=16k**  
250880+0 records in  
250880+0 records out  
4110417920 bytes (4.1 GB) copied, 5.10768 seconds, 805 MB/s
18. Perform the steps in [“Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software” on page 115](#) to continue with the recovery installation.

## Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software

---

This procedure describes how to perform a recovery installation using an external USB flash drive that contains Junos OS software.



**NOTE:** Since the recovery installation process completely overwrites the entire contents of the Director device, you will need to restore the required configuration files and initial setup information. The following procedure assumes you previously saved these backup files with the **request system software configuration-backup** command. Ensure that you have these backup files available on an external USB flash drive before you perform the following steps.

1. Insert the external USB flash drive into the Director device.
2. Perform one of the following tasks:
  - If you have access to the default partition, reboot the Director device by issuing the **request system reboot director-group** command.
  - If you do not have access to the default partition, power cycle the Director device.

The following menu appears on the Director device console when the Director device boots up:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

3. Type **install** and then press **Enter** to install the software on the Director device.

Once the installation process is complete, the Director device reboots, and the following menu appears on the Director device console:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

4. Press **Enter**.

The Director device reboots from the local disk on which the software was just installed.

5. Log in as root on the Director device.

The following menu appears on the Director device console:

```
Before you can access the QFabric system, you must complete the initial setup
of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of the
Director device and then log back in to restart
this setup menu.
```

```
Continue?[y/n]
```



6. Enter **n** to bypass the initial setup script and enter the Director device root directory, where you can mount the external USB flash drive containing the configuration files and initial setup information.

7. Issue the **ls /mnt** command to list the **mount** directory.

```
root@dg0 ~]# ls /mnt
```

8. Issue the **mkdir** command to create a directory within the mount directory.

```
root@dg0 ~]# mkdir /mnt/myusb
```

9. Issue the **mount /dev/sdb2 /mnt/myusb/** command to mount the external USB flash drive to the local drive of the Director device.

```
root@dg0 ~]# mount /dev/sdb2 /mnt/myusb/
```

10. Issue the **ls -la /mnt/myusb/** command to verify the contents of your mounted external USB flashdrive.

```
root@dg0 ~]# ls -la /mnt/myusb/
total 1770884
drwxr-xr-x 2 root root 4096 Sep 7 05:16 .
drwxr-xr-x 3 root root 4096 Sep 7 10:15 ..
-rw-r--r-- 1 root root 4249 Sep 7 03:52 mybackup-20110907
```

11. Exit the Director device and log back in as root on the Director device.

The following menu appears:

Before you can access the QFabric system, you must complete the initial setup of the Director group by using the steps that follow.

If the initial setup procedure does not complete successfully, log out of the Director device and then log back in to restart this setup menu.

```
Continue?[y/n] y
Initial Configuration
```

You may enter the configuration manually or restore from a backup.

```
Specify a backup file? [y/n] : y
Please specify the full path of the configuration backup file. :
/mnt/myusb/mybackup-20110907
```

12. Enter **y** to continue.

13. Enter **y** and specify the path to the backup configuration file located on the external USB flash drive.

```
/mnt/myusb/mybackup-20110907
```

The following messages appear:

```
Saving temporary configuration...
Configuring peer...
connect error for 1.1.1.2:9001
Configuring local interfaces...
Configuring interface eth0 with [10.49.213.163/24:10.49.213.254]
Configured interface eth0 with [10.49.213.163/24:10.49.213.254]
Configuring QFabric software with initial pool of 4000 MAC addresses
[00:10:00:00:00:00 - 00:10:00:00:0f:3b]
Configuring QFabric address [10.49.213.50]
Reconfiguring QFabric software static configuration
Applying the new Director Device password
Applying the QFabric component password
```

```
First install initial configuration, generating and sharing SSH keys.
First install initial configuration, generating SSH keys.
connect error for 1.1.1.2:9001
Shared SSH keys.
Configuration complete. Director Group services will auto start within 30
seconds.
```

The Director device reboots from the local disk on which the software was just installed.  
Exit the Director device session and log in to the QFabric default partition CLI.

14. Issue the **request system software configuration-restore** command and specify the path to the backup configuration file located on the external USB flash drive to load the previously saved QFabric system configuration.

15. From the default partition, issue the **request system reboot node-group all** command to reboot all of the Node groups in the QFabric system to ensure that all Node devices are running the same version of software as the Director-group.

```
user@switch> request system reboot node-group all
```

16. From the default partition, issue the **request system reboot fabric** command to reboot the Interconnect devices and the other components in the fabric in the QFabric system to ensure that Interconnect devices are running the same version of software as the Director group.

```
user@switch> request system reboot fabric
```

17. Log in to the default partition and issue the **show version component all** command to verify that all components are running the same version of software.

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
```

```

JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

```

FC-1:

```

Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

```

DRE-0:

```

-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

```

FM-0:

```

-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

```

nodedevice1:

```

-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

```

interconnectdevice1:

```

-
Hostname: qfabric

```

```
Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
warning: from interconnectdevice0: Disconnected
```

- Related Documentation**
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group on page 1335](#)
  - [Upgrading Software on a QFabric System on page 134](#)
  - [request system software configuration-backup on page 402](#)
  - [request system software configuration-restore on page 403](#)

---

## Layer 3 Protocols

- [Troubleshooting Virtual Routing Instances on page 6536](#)

### Troubleshooting Virtual Routing Instances

- [Direct Routes Not Leaked Between Routing Instances on page 6536](#)

---

#### Direct Routes Not Leaked Between Routing Instances

**Problem**    **Description:** Direct routes are not exported (leaked) between virtual routing instances. For example, consider the following scenario:

- QFX switch with two virtual routing instances:
  - Routing instance 1 connects to downstream device through interface xe-0/0/1.
  - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the QFX switch connects to the upstream device over a direct route and these routes are not leaked between instances.



**NOTE:** You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

---

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the QFX switch over indirect routes.

**Solution**    This is expected behavior.

- Related Documentation**
- [Understanding Virtual Router Routing Instances on page 2822](#)
  - [Configuring Virtual Router Routing Instances on page 2829](#)
  - [rib-group on page 2935](#)

## Security

- [Troubleshooting Firewall Filter Configuration on page 6537](#)
- [Troubleshooting Policer Configuration on page 6543](#)

### Troubleshooting Firewall Filter Configuration

Use the following information to troubleshoot your firewall filter configuration.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 6537](#)
- [Filter Counts Previously Dropped Packet on page 6539](#)
- [Matching Packets Not Counted on page 6539](#)
- [Counter Reset When Editing Filter on page 6540](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 6540](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 6540](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 6541](#)
- [Egress Firewall Filters with Private VLANs on page 6541](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 6542](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 6542](#)
- [Invalid Statistics for Policer on page 6542](#)
- [Policers can Limit Egress Filters on page 6542](#)

#### Firewall Filter Configuration Returns a No Space Available in TCAM Message

**Problem** **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

**Solution** When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```



**NOTE:** The original filter is not deleted and is still available in the configuration.

---

### Filter Counts Previously Dropped Packet

- Problem** **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
  - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

**Solution** This is expected behavior.

### Matching Packets Not Counted

**Problem** **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet. For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **admin** VLAN, and interface xe-0/0/1 is a member of that VLAN.

- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

**Solution** This is expected behavior.

---

#### Counter Reset When Editing Filter

**Problem** **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

---

#### Cannot Include loss-priority and policer Actions in Same Term

**Problem** **Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

**Solution** This is expected behavior.

---

#### Cannot Egress Filter Certain Traffic Originating on QFX Switch

**Problem** **Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

**Solution** This is expected behavior.



### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

**Problem** **Description:** If you create a firewall filter that includes a match condition of `dot1q-tag` or `dot1q-user-priority` and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the `set dot1q-tunneling ethertype 0x8100` statement at the `[edit ethernet-switching-options]` hierarchy level. You must also configure the other end of the link to use the same Ethertype.

### Egress Firewall Filters with Private VLANs

**Problem** **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).

- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Filtering of L2PT Traffic Not Supported

---

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

---

#### Cannot Drop BGP Packets in Certain Circumstances

---

**Problem** **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

**Solution** This is expected behavior.

---

#### Invalid Statistics for Policer

---

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

---

#### Policers can Limit Egress Filters

---

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional

egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5038](#)
- [Configuring Firewall Filters on page 4747](#)
- [Verifying That Firewall Filters Are Operational on page 4848](#)

## Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 6544](#)
- [Counter Reset When Editing Filter on page 6544](#)
- [Invalid Statistics for Policer on page 6544](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 6544](#)

- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 6545](#)
- [Policers Can Limit Egress Filters on page 6546](#)

---

### Incomplete Count of Packet Drops

---

**Problem**    **Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

**Solution**    This is expected behavior.

---

### Counter Reset When Editing Filter

---

**Problem**    **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution**    This is expected behavior.

---

### Invalid Statistics for Policer

---

**Problem**    **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution**    This is expected behavior.

---

### Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

---

**Problem**    **Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate

might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

**Solution** This is expected behavior.

### Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem** **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and

reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

**Solution** To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 4658](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

---

### Policers Can Limit Egress Filters

---

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters. Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.

- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

## Services

- [Troubleshooting Port Mirroring on page 6547](#)

### Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 6547](#)
- [Egress Port Mirroring with VLAN Translation on page 6549](#)
- [Egress Port Mirroring with Private VLANs on page 6549](#)

#### Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 6547](#)
- [Remote Port Mirroring Only on page 6549](#)

##### **Local and Remote Port Mirroring**

The following constraints and limitations apply to local and remote port mirroring with the QFX Series:

- You can create a total of four port-mirroring configurations on a QFX Series standalone switch.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a QFX3500 device or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).



### ***Remote Port Mirroring Only***

The following constraints and limitations apply to remote port mirroring with the QFX Series:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

### **Egress Port Mirroring with VLAN Translation**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

### **Egress Port Mirroring with Private VLANs**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Port Mirroring on page 4887](#)
- [Example: Configuring Port Mirroring for Local Analysis on page 4895](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 4900](#)

---

## Storage

- [Troubleshooting Dropped FCoE Traffic on page 6550](#)
- [Troubleshooting Fibre Channel Interface Deletion on page 6553](#)
- [Troubleshooting Dropped FIP Traffic on page 6554](#)

## Troubleshooting Dropped FCoE Traffic

**Problem** **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

**Cause** There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.
5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.

6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Disabled
 100 Disabled
 101 Enabled 2500
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 101 5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “Example: Configuring CoS PFC for FCoE Traffic” on page 5113 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 5931](#)
- [show class-of-service forwarding-class-set on page 5938](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 5795](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5113](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 5428](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5064](#)

## Troubleshooting Fibre Channel Interface Deletion

**Problem**    **Description:** You deleted a Fibre Channel (FC) interface at the **[edit interfaces]** hierarchy level, but the commit check fails so the interface is not deleted.

**Cause**    You must first delete the FC interface from the FC fabric on the QFX Series before you can delete the FC interface at the **[edit interfaces]** hierarchy level. You must perform both operations to delete a FC interface.

**Solution** First delete the interface from the FC fabric and then delete the interface from the QFX Series:

1. Delete the FC interface from the FC fabric to which it belongs:

```
[edit]
user@switch# delete fc-fabrics fabric-name interface interface-name
```

For example, to delete the FC interface **fc-0/0/3.0** from an FC fabric named **sanfab1**:

```
[edit]
user@switch# delete fc-fabrics sanfab1 interface fc-0/0/3.0
```

2. Delete the FC interface at the **[edit interfaces]** hierarchy level:

```
[edit]
user@switch: delete interfaces interface-name
```

For example, to delete the interface **fc-0/0/3.0** from the switch:

```
[edit]
user@switch: delete interfaces fc-0/0/3.0
```

- Related Documentation**
- [fc-fabrics on page 5232](#)
  - [interface on page 5242](#)
  - [interfaces on page 2584](#)
  - [Understanding Interfaces on an FCoE-FC Gateway on page 4996](#)

## Troubleshooting Dropped FIP Traffic

**Problem** **Description:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) traffic such as FIP VLAN discovery and notification frames is dropped on the QFX Series.

**Cause** The interface on which the FIP traffic is dropped does not have a native VLAN configured. FIP VLAN discovery and notification messages are exchanged as untagged packets on the native VLAN. (After the FCoE session with the Fibre Channel switch is established, FCoE traffic uses the FCoE VLAN.)

**Solution** Check to ensure that every 10-Gigabit Ethernet interface that connects to an FCoE device includes a native VLAN. Configure a native VLAN on all 10-Gigabit Ethernet interfaces that connect to FCoE devices.



**NOTE:** Make sure that the native VLAN you are using on the QFX Series is the same native VLAN that the FCoE devices use for Ethernet traffic.

---

To configure a native VLAN on an interface:

1. Set the interface port mode to **tagged-access** if you have not already done so:

```
[edit]
```

```
user@switch# set interfaces interface unit unit family ethernet-switching port-mode
tagged-access
```

For example, to set the port mode to **tagged-access** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure a native VLAN on the interface:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching native-vlan-id
vlan-id
```

For example, to set the native VLAN ID to 1 for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
```

3. Configure the native VLAN if it does not already exist:

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

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

#### Related Documentation

- [interfaces on page 2584](#)
- [vlans on page 2145](#)
- [Understanding FIP Functions on page 4984](#)

## Traffic Management

- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 6555](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 6556](#)
- [Troubleshooting Egress Queue Bandwidth Impacted by Congestion on page 6557](#)
- [Troubleshooting an Unexpected Rewrite Value on page 6558](#)
- [Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic on page 6559](#)

### Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

**Problem**    **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth (shaping rate) configured for the queue.

**Cause**    When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.

The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

**Solution** When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

**Related  
Documentation**

- [shaping-rate on page 5904](#)
- [Example: Configuring Maximum Output Bandwidth on page 5689](#)
- [Example: Configuring Queue Schedulers on page 5674](#)
- [Understanding CoS Output Queue Schedulers on page 5486](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

**Problem** **Description:** The minimum bandwidth of a queue or a priority group when measured at the egress port exceeds the minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

**Cause** When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



**NOTE:** The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

---

**Solution** When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit



rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

- Related Documentation**
- [guaranteed-rate on page 5869](#)
  - [transmit-rate on page 5911](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 5684](#)
  - [Example: Configuring Queue Schedulers on page 5674](#)
  - [Understanding CoS Output Queue Schedulers on page 5486](#)

## Troubleshooting Egress Queue Bandwidth Impacted by Congestion

**Problem** **Description:** Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

**Cause** Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

**Solution** The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a tail-drop profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

- Related Documentation**
- [drop-profile on page 5851](#)
  - [Example: Configuring Tail-Drop Profiles on page 5663](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5606](#)
  - [Understanding CoS Tail-Drop Profiles on page 5545](#)

## Troubleshooting an Unexpected Rewrite Value

**Problem**    **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

**Cause**    If you configure a rewrite rule for a forwarding class on an egress port but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

**Solution**    If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



**TIP:** If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



**NOTE:** There are no default rewrite rules. You can bind one rewrite rule for each type (DSCP and IEEE 802.1) to a given interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value associations.

1. Assign a rewrite value to a forwarding class. Add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:  
  
[edit class-of-service rewrite-rules]

```
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **fcoe** is being randomly rewritten, and you want to use IEEE 802.1 code point **011** for the **fcoe** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class fcoe loss-priority high code-point
011
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

#### Related Documentation

- [interfaces on page 5880](#)
- [rewrite-rules on page 5897](#)
- [Defining CoS Rewrite Rules on page 5805](#)
- [Monitoring CoS Rewrite Rules on page 5918](#)

## Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

**Problem**    **Description:** In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.

**Cause**    Failure of a queue to transmit packets for 12 consecutive seconds may be due to:

- A strict-high priority queue consuming all of the port bandwidth
- Several queues consuming all of the port bandwidth
- Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
- Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

**Solution**    If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map
scheduler-map-name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name
output-traffic-control-profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues.

[Table 581 on page 6073](#) shows the topology for this example:

**Table 640: Components of the Rate Shaping Troubleshooting Example**

| Component                                  | Settings                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Affected interface                         | <b>shpnode:xe-0/0/10</b>                                                                                                                                                                                                                                                                                                                          |
| Scheduler (strict-high priority scheduler) | Name: <b>shp-sched</b><br>Shaping rate: <b>7g</b><br>Priority: <b>strict-high</b><br><br><b>NOTE:</b> This example assumes that the scheduler already exists and has been configured as <b>strict-high</b> priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied. |
| Scheduler map                              | Name: <b>shp-map</b><br>Forwarding class to associate with the <b>shp-sched</b> scheduler: <b>strict-high</b><br><br><b>NOTE:</b> This example assumes that a strict-high priority forwarding class has been configured and assigned the name <b>strict-high</b> .                                                                                |
| Traffic control profile                    | Name: <b>shp-tcp</b><br><br><b>NOTE:</b> This example does not describe how to define a complete traffic control profile.                                                                                                                                                                                                                         |

Table 640: Components of the Rate Shaping Troubleshooting Example (continued)

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding class set  | <div>Name: <b>shp-pg</b></div> <div><p>To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface <b>shpnode:xe-0/0/10</b> using the CLI:</p><ol style="list-style-type: none"><li>Specify the scheduler for the strict-high priority queue (<b>shp-sched</b>) with a maximum bandwidth of 7 Gbps:<br/><pre>[edit class-of-service schedulers] user@switch# set shp-sched shaping-rate 7g</pre></li><li>Configure a scheduler map (<b>shp-map</b>) that associates the scheduler (<b>shp-sched</b>) with the forwarding class (<b>strict-high</b>):<br/><pre>[edit class-of-service scheduler-maps] user@switch# set shp-map forwarding-class strict-high scheduler shp-sched</pre></li><li>Associate the scheduler map <b>shp-map</b> with a traffic control profile (<b>shp-tcp</b>):<br/><pre>[edit class-of-service traffic-control-profiles] user@switch# set shp-tcp scheduler-map shp-map</pre></li><li>Associate the traffic control profile <b>shp-tcp</b> with a forwarding class set (<b>shp-pg</b>) and the affected interface (<b>shpnode:xe-0/0/10</b>):<br/><pre>[edit class-of-service] user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg output-traffic-control-profile shp-tcp</pre></li></ol></div> |
| Related Documentation | <div><ul style="list-style-type: none"><li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5486</a></li><li>• <a href="#">Defining CoS Queue Scheduling Priority on page 5792</a></li><li>• <a href="#">Example: Configuring Queue Schedulers on page 5674</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 5682</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 5671</a></li><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5606</a></li></ul></div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



## PART 23

# Index

- [Index on page 6565](#)





# Index

## Symbols

|                                              |            |
|----------------------------------------------|------------|
| !                                            |            |
| regular expression operator.....             | 1686       |
| system logging.....                          | 6228       |
| #, comments in configuration statements..... | cxxiv      |
| \$                                           |            |
| regular expression operator.....             | 1686, 1687 |
| system logging.....                          | 6228       |
| ( )                                          |            |
| regular expression operator.....             | 1686, 1687 |
| system logging.....                          | 6228       |
| ( ), in syntax descriptions.....             | cxxiv      |
| *                                            |            |
| regular expression operator.....             | 1686       |
| system logging.....                          | 6228       |
| +                                            |            |
| regular expression operator.....             | 1686       |
| system logging.....                          | 6228       |
| .                                            |            |
| regular expression operator.....             | 1686       |
| system logging.....                          | 6228       |
| /var/log/mib2d file.....                     | 6393, 6495 |
| /var/log/snmpd file.....                     | 6393, 6495 |
| 128-bit IPv6 address.....                    | 2041       |
| 32-bit IPv4 address.....                     | 2041       |
| 802.3ad statement.....                       | 2549       |
| < >, in syntax descriptions.....             | cxxiv      |
| ?                                            |            |
| regular expression operator                  |            |
| system logging.....                          | 6228       |
| [ ]                                          |            |
| regular expression operator                  |            |
| system logging.....                          | 6228       |
| [ ], in configuration statements.....        | cxxiv      |
| \                                            |            |
| regular expression operator.....             | 1686, 1687 |
| ^                                            |            |
| regular expression operator.....             | 1686, 1687 |
| system logging.....                          | 6228       |
| { }, in configuration statements.....        | cxxiv      |

|                                     |       |
|-------------------------------------|-------|
|                                     |       |
| regular expression operator         |       |
| system logging.....                 | 6228  |
| (pipe), in syntax descriptions..... | cxxiv |

## A

|                                               |            |
|-----------------------------------------------|------------|
| ABRs See area border routers                  |            |
| accept-data statement.....                    | 2320       |
| usage guidelines.....                         | 2304       |
| accept-remote-nexthop statement.....          | 3485       |
| usage guidelines.....                         | 3367       |
| accept-remote-source statement                |            |
| usage guidelines.....                         | 4330       |
| access privilege levels                       |            |
| configuration example.....                    | 1729       |
| operational mode commands.....                | 1729       |
| configuring.....                              | 1692       |
| operational mode commands.....                | 178, 1720  |
| login classes.....                            | 1672       |
| user accounts.....                            | 1682       |
| access statement.....                         | 1749       |
| access-end statement                          |            |
| login class.....                              | 232        |
| access-start statement                        |            |
| login class.....                              | 232        |
| accounting statement.....                     | 233, 1750  |
| authentication                                |            |
| usage guidelines.....                         | 1697, 1714 |
| IGMP.....                                     | 4425       |
| IGMP interface.....                           | 4426       |
| accounting-options statement.....             | 1751       |
| accounting-port statement                     |            |
| RADIUS servers.....                           | 234        |
| accounting-server statement.....              | 1753       |
| accounting-stop-on-access-deny statement..... | 1754       |
| accounting-stop-on-failure statement.....     | 1755       |
| action statement.....                         | 4784       |
| actions                                       |            |
| routing policy.....                           | 3923       |
| activate command                              |            |
| usage guidelines.....                         | 81         |
| active routes.....                            | 3332       |
| active statement                              |            |
| aggregate routes.....                         | 2857       |
| generated routes.....                         | 2857       |
| static routes.....                            | 2857       |
| active-source-limit statement.....            | 4462       |
| usage guidelines.....                         | 4332       |

|                                            |            |
|--------------------------------------------|------------|
| add-path statement                         |            |
| BGP                                        |            |
| usage guidelines.....                      | 3380       |
| address statement.....                     | 2069, 2550 |
| anycast RPs.....                           | 4356       |
| usage guidelines.....                      | 4267, 4338 |
| local RPs.....                             | 4357       |
| SNMPv3.....                                | 6278       |
| static RPs.....                            | 4358       |
| usage guidelines.....                      | 4266       |
| usage guidelines.....                      | 2040       |
| address-mask statement.....                | 6279       |
| addresses                                  |            |
| IP addresses.....                          | 184        |
| router source addresses.....               | 151, 252   |
| administrative distance.....               | 3326       |
| BGP See preference statement               |            |
| advertise-external statement.....          | 3486       |
| usage guidelines.....                      | 3310       |
| advertise-inactive statement.....          | 3487       |
| usage guidelines.....                      | 3310, 3328 |
| advertise-interval statement.....          | 2321       |
| usage guidelines.....                      | 2299       |
| advertise-peer-as statement.....           | 3488       |
| usage guidelines.....                      | 3313       |
| advertisement intervals, VRRP.....         | 2299       |
| advertisement-interval statement.....      | 1756       |
| agent-address statement.....               | 1757, 6279 |
| agent-id statement.....                    | 6267       |
| aggregate routes.....                      | 2858       |
| graceful restart.....                      | 2269       |
| preferences.....                           | 3326       |
| aggregate statement.....                   | 2858       |
| aggregated-devices statement.....          | 2552       |
| aggregated-ether-options statement.....    | 2553       |
| aggregator statement.....                  | 2860       |
| AIGP                                       |            |
| BGP.....                                   | 3267       |
| aigp statement                             |            |
| BGP                                        |            |
| usage guidelines.....                      | 3267       |
| aigp-originate statement                   |            |
| BGP                                        |            |
| usage guidelines.....                      | 3267       |
| alarm conditions.....                      | 147        |
| alarm messages.....                        | 6488       |
| Alarm MIB.....                             | 6130       |
| alarm severity                             |            |
| major (red) .....                          | 6487       |
| See also major alarms                      |            |
| minor (yellow).....                        | 6487       |
| See also minor alarms                      |            |
| alarm statement.....                       | 2554       |
| RMON.....                                  | 6280       |
| STP.....                                   | 2071       |
| usage guidelines.....                      | 147        |
| alarms                                     |            |
| chassis.....                               | 6488       |
| interface.....                             | 6491       |
| major See major alarms                     |            |
| minor See minor alarms                     |            |
| overview.....                              | 6487       |
| red See major alarms                       |            |
| yellow See minor alarms                    |            |
| alarms, displaying                         |            |
| chassis.....                               | 452        |
| health monitor.....                        | 6445       |
| RMON.....                                  | 6455       |
| system.....                                | 914        |
| alert (system logging severity level)..... | 2909       |
| algorithm statement                        |            |
| BFD authentication.....                    | 4359       |
| alias option for static-host-mapping       |            |
| statement.....                             | 295        |
| alias statement.....                       | 295        |
| aliases statement                          |            |
| QFabric system aliases.....                | 1380       |
| all (tracing flag).....                    | 2950       |
| class of service.....                      | 5908       |
| Fibre Channel.....                         | 5262       |
| Fibre Channel fip.....                     | 5260, 5265 |
| Fibre Channel proxy.....                   | 5267       |
| all-failures (tracing flag)                |            |
| STP.....                                   | 2137       |
| allow-commands statement.....              | 234        |
| usage guidelines.....                      | 1675       |
| allow-configuration statement.....         | 235        |
| usage guidelines.....                      | 1675       |
| allow-transients statement.....            | 236, 6237  |
| allowed-days statement                     |            |
| login class.....                           | 235        |
| allowed-mac statement.....                 | 4803, 4821 |
| allowing commands to login classes.....    | 1675       |
| always compare, BGP MED option.....        | 3209       |
| always-compare-med option.....             | 3332       |
| Analyzer MIB.....                          | 6130, 6137 |

- 
- analyzer statement.....4911
  - annotate command.....81
  - announcement statement.....236
    - usage guidelines.....149
  - announcements
    - system login.....149
  - any (system logging facility).....6222
  - any (system logging severity level).....6223
  - any-sender statement
    - RIP.....4172
  - anycast RP.....4273
    - overview.....4217
  - anycast-pim statement.....4360
    - usage guidelines.....4267, 4338
  - application statement.....5217, 5823
  - apply-groups statement.....3491
  - apply-groups-except statement.....3491
  - apply-macro statement .....6238
  - apply-path statement
    - firewall filter match condition.....3441
  - archival configuration
    - displaying.....941, 1651
  - archival statement.....237, 1614, 1758
    - usage guidelines.....1611
  - archive router configuration.....1611
  - archive statement
    - all system log files.....1381, 6366, 6369
    - individual system log file.....6368
    - usage guidelines.....6220
  - archive-sites statement
    - configuration files.....1615, 1759
    - system log files.....6368
    - system logging
      - usage guidelines.....6220
    - usage guidelines.....1612
  - archiving files.....337, 1632
  - area border routers
    - backbone area See backbone area
    - description.....3814
    - overview.....3803
  - area statement .....3963
    - usage guidelines, backbone.....3815
    - usage guidelines, multiarea.....3817
  - area-range statement.....3965
    - usage guidelines.....3852
  - areas See area border routers; backbone area;
    - NSSAs; stub areas
    - overview.....3803
  - arithmetic and relational operators
    - for monitor traffic command.....6410
  - arp.....238
    - statistics, displaying.....2250
    - status information, displaying.....2250
  - arp statement.....238, 2072
  - arp-inspection statement.....4804
  - AS boundary routers
    - overview.....3804
  - AS external link advertisements.....3807
  - AS path
    - ignoring in route selection.....3336
  - AS paths
    - distribution of, displaying.....2960
    - domain information, displaying.....2964
    - matching regular expressions,
      - displaying.....2986
    - summary of, displaying.....2967
  - as-path (tracing flag).....3577
  - as-path statement.....2860
  - as-path-ignore
    - usage guidelines.....3332, 3336
  - as-path-ignore option.....3563
  - ASCII file, Junos OS, configuring using.....169
  - asm-override-ssm statement.....4426, 4482
  - ASN
    - BGP community routes, displaying.....2993
  - ASs
    - configuring.....2862
    - paths.....3146
      - aggregate routes.....2860
      - generated routes.....2860, 2875
      - operations, tracing.....3577
      - static routes.....2860
    - private, removing.....3342, 3343, 3568
  - ASs (autonomous systems)
    - area border routers.....3814
    - breaking into confederations.....3422
    - stub areas See stub areas
  - assert (tracing flag).....4419
  - assert timeout
    - configuring.....4284
  - assert-timeout statement.....4361
    - usage guidelines.....4284
  - asymmetric-hold-time statement.....2322
  - asynch (tracing flag)
    - class of service.....5908
  - auth (tracing flag).....4193

|                                              |                              |  |
|----------------------------------------------|------------------------------|--|
| authentication See of routes                 |                              |  |
| algorithm                                    |                              |  |
| BGP.....                                     | 3430                         |  |
| IS-IS.....                                   | 3631                         |  |
| BGP.....                                     | 3147, 3489, 3492             |  |
| IS-IS.....                                   | 3631                         |  |
| keychains                                    |                              |  |
| BGP.....                                     | 3430                         |  |
| IS-IS.....                                   | 3631                         |  |
| MD5                                          |                              |  |
| BGP.....                                     | 3430                         |  |
| NTP authentication keys.....                 | 152, 153                     |  |
| order.....                                   | 1676, 1694                   |  |
| RADIUS.....                                  | 1665, 1681, 1699, 1701       |  |
| RIPv2, MD5.....                              | 4121                         |  |
| RIPv2, plain-text passwords.....             | 4121                         |  |
| root password.....                           | 158, 1687, 1702              |  |
| shared user accounts.....                    | 1665, 1701                   |  |
| simple                                       |                              |  |
| RIP.....                                     | 4116                         |  |
| TACACS+.....                                 | 1665, 1681, 1701, 1711       |  |
| users.....                                   | 1681                         |  |
| authentication configuration                 |                              |  |
| BFD.....                                     | 2848, 3359, 3889, 4130, 4292 |  |
| authentication statement                     |                              |  |
| BFD.....                                     | 4363                         |  |
| BFD protocol.....                            | 4362                         |  |
| login.....                                   | 239                          |  |
| usage guidelines.....                        | 1681, 1692                   |  |
| authentication, VRRP.....                    | 2297                         |  |
| authentication-algorithm statement           |                              |  |
| BGP.....                                     | 3494                         |  |
| usage guidelines.....                        | 3430                         |  |
| IS-IS                                        |                              |  |
| usage guidelines.....                        | 3647                         |  |
| authentication-key statement.....            | 240, 2323                    |  |
| BGP.....                                     | 3495                         |  |
| usage guidelines.....                        | 3430                         |  |
| IS-IS.....                                   | 3694                         |  |
| usage guidelines.....                        | 3647, 3690                   |  |
| MSDP.....                                    | 4463                         |  |
| RIP.....                                     | 4173                         |  |
| usage guidelines.....                        | 4116                         |  |
| usage guidelines.....                        | 152, 153, 2297               |  |
| authentication-key-chain statement.....      | 3496, 3695                   |  |
| BGP                                          |                              |  |
| usage guidelines.....                        | 3430                         |  |
| IS-IS                                        |                              |  |
| usage guidelines.....                        | 3647                         |  |
| usage guidelines.....                        | 3430                         |  |
| authentication-key-chains statement.....     | 4920                         |  |
| authentication-md5 statement.....            | 6281                         |  |
| authentication-none statement.....           | 6282                         |  |
| authentication-order statement.....          | 241, 1760                    |  |
| usage guidelines.....                        | 1676, 1694                   |  |
| authentication-password statement.....       | 6283                         |  |
| authentication-server statement.....         | 1761                         |  |
| authentication-sha statement.....            | 6284                         |  |
| authentication-type statement.....           | 2324                         |  |
| IS-IS.....                                   | 3696                         |  |
| usage guidelines.....                        | 3690                         |  |
| RIP.....                                     | 4174                         |  |
| usage guidelines.....                        | 4116                         |  |
| usage guidelines.....                        | 2297                         |  |
| authorization See permissions                |                              |  |
| authorization (system logging facility)..... | 6222                         |  |
| option to facility-override statement.....   | 6224                         |  |
| authorization statement.....                 | 1762, 6285                   |  |
| usage guidelines.....                        | 6194                         |  |
| auto-load-rebalance statement.....           | 5217                         |  |
| autonomous system number See ASN             |                              |  |
| autonomous system paths See AS paths         |                              |  |
| autonomous-system statement.....             | 2862                         |  |
| auxiliary statement.....                     | 242                          |  |
| <b>B</b>                                     |                              |  |
| backbone area                                |                              |  |
| configuring.....                             | 3815                         |  |
| description.....                             | 3814                         |  |
| overview.....                                | 3803                         |  |
| backbone router                              |                              |  |
| overview.....                                | 3804                         |  |
| Backup and recovery                          |                              |  |
| QFabric system.....                          | 1460, 1467                   |  |
| backup router configuration.....             | 1611                         |  |
| backup-pe-group statement.....               | 2864                         |  |
| backup-spf-options statement                 |                              |  |
| OSPF.....                                    | 3968                         |  |
| backups statement.....                       | 2865                         |  |
| bandwidth statement.....                     | 2866, 4837                   |  |
| bandwidth-based metrics                      |                              |  |
| OSPF.....                                    | 3862                         |  |

- bandwidth-based-metrics statement.....3969
  - usage guidelines.....3867
- bandwidth-limit statement.....4784
- bandwidth-threshold statement.....2325
  - usage guidelines.....2302
- bb-sc-n statement
  - fibrechannel-options.....5218
- Beacon
  - operation of, controlling.....358
- best routes, displaying.....2988
- BFD
  - authentication
    - configuration.....2848, 3359, 3889, 4130, 4292
    - configuring.....3883
    - protocol.2831, 3348, 3349, 3881, 4122, 4123, 4290
    - traceoptions statement
      - usage guidelines.....2835
    - with IBGP.....3348, 3349
  - BFD authentication
    - algorithm statement.....4359
    - authentication statement.....4362
    - for IS-IS.....3665
    - key-chain statement.....4386
    - loose-check statement.....4389
    - OSPFv2.....3887
- bfd-liveness-detection statement
  - BGP.....3497, 3509
    - minimum-interval.....3545
    - threshold.....3576
    - transmit-interval.....3580
    - usage guidelines.....3349
  - IS-IS.....3697
  - OSPF.....3971
    - usage guidelines.....3883
  - PIM.....4363, 4368
    - minimum-interval.....4392
    - threshold.....4417
    - transmit-interval.....4418
    - usage guidelines.....4290
  - RIP.....4175
    - usage guidelines.....4123
  - static routes.....2867
    - usage guidelines.....2831
  - usage guidelines.....2840
- BGP See multipath
  - administrative distance.....3326
  - advertising multiple paths to a
    - destination.....3379, 3380
  - aggregator path attribute.....3558
  - AIGP attribute.....3267
  - AS numbers, peers.....3565
  - ASs See ASs
  - authentication.....3147, 3430, 3495
  - authentication algorithm.....3494
  - authentication keychain.....3496
  - best external route
    - advertising.....3310
  - BFD.3348,3349,3489,3492,3497,3529,3540,3571,3581
  - community ASN, displaying routes.....2993
  - community name, displaying routes.....2995
  - confederations.....2877
  - damping parameters
    - clearing.....3584
    - displaying.....3615
  - damping routes, displaying.....2997, 3617
  - description.....3173
  - EBGP IPv6 peering.....3157
  - enabling on router.....3501
  - external (EBGP).....3145
  - graceful restart.....1388, 2269, 2310, 3517
  - groups.....3149, 3208, 3518
    - general information, displaying.....3589
  - hold time.....3147, 3523
  - identifier.....3147
  - ignoring the AS path attribute in route
    - selection.....3336
  - indirect next hops.....3367
  - injecting OSPF routes into BGP.....3306, 3923
  - internal.....3172
  - internal (IBGP).....3145
  - IP address.....3147
  - IPv6
    - logical systems.....3157
  - keepalive messages.....3148, 3577
  - local address.....3173, 3531
  - local AS.....3246, 3251, 3261
  - local preference.....3194, 3195
  - MED.....3210, 3223, 3236
  - messages.....3147
  - monitoring.....3583
  - MP-BGP.....3513
  - MTU discovery.....3548
  - multihop sessions.....3550
  - multipath configuration.....3362, 3363
  - neighbors
    - clearing connections.....3585
    - displaying.....2346, 3596
  - neighbors BGP, peers See BGP, peers

|                                            |                  |                                                     |            |
|--------------------------------------------|------------------|-----------------------------------------------------|------------|
| NLRI.....                                  | 3148             | VPNs                                                |            |
| open messages.....                         | 3147, 3449, 3562 | preventing session flaps.....                       | 3449       |
| outbound route filter.....                 | 3313             | with BFD.....                                       | 3348, 3349 |
| outbound route filters                     |                  | BGP (Border Gateway Protocol)                       |            |
| interoperability.....                      | 2871, 3502       | confederations <i>See</i> BGP confederations        |            |
| overview.....                              | 3144             | export policy for CLNS.....                         | 3684       |
| packets, tracing.....                      | 3577             | internal peer session (configuration                |            |
| passive mode.....                          | 3449             | editor).....                                        | 3173       |
| path attributes.....                       | 3146, 3148       | peering sessions <i>See</i> BGP peers; BGP sessions |            |
| peers.....                                 | 3145, 3208, 3553 | point-to-point internal peer session                |            |
| point-to-point peer session (configuration |                  | logical systems.....                                | 3184       |
| editor).....                               | 3150             | route reflectors <i>See</i> BGP route reflectors    |            |
| policy, routing.....                       | 3512, 3525       | route-flap damping.....                             | 3456       |
| precedence.....                            | 3309             | BGP confederations                                  |            |
| preferences.....                           | 3326, 3327, 3567 | creating (configuration editor).....                | 3423       |
| private AS.....                            | 3342, 3343       | description.....                                    | 3422       |
| route reflection.....                      | 3504, 3559       | route-flap damping.....                             | 3456       |
| router identifier.....                     | 2938             | BGP groups                                          |            |
| routes.....                                | 3146             | confederations (configuration editor).....          | 3423       |
| routing tables                             |                  | BGP Monitoring Protocol.....                        | 2872       |
| delays in exchanging routes.....           | 3312, 3560       | BGP peers                                           |            |
| nonactive routes.....                      | 3310, 3328, 3487 | external (configuration editor).....                | 3150       |
| retaining routes.....                      | 3527             | internal.....                                       | 3184       |
| session drops.....                         | 3449             | internal (configuration editor).....                | 3173       |
| set local AS number.....                   | 3533             | point-to-point connections.....                     | 3149       |
| summary information, displaying.....       | 3610             | BGP route reflectors                                |            |
| system log messages.....                   | 3449             | cluster of clusters.....                            | 3406       |
| table                                      |                  | creating (configuration editor).....                | 3407       |
| clearing.....                              | 3587             | description.....                                    | 3405       |
| TCP.....                                   | 3144             | multiple clusters.....                              | 3406       |
| TCP segment size.....                      | 3444             | BGP routing information.....                        | 3583       |
| tracing operations.....                    | 3577             | BGP sessions                                        |            |
| 4-byte AS events.....                      | 3476             | internal.....                                       | 3184       |
| BFD protocol events.....                   | 3476             | internal (configuration editor).....                | 3173       |
| damping operations.....                    | 3476             | point-to-point external (configuration              |            |
| description.....                           | 3476             | editor).....                                        | 3150       |
| graceful restart.....                      | 3476             | sample peering session.....                         | 3150       |
| keepalive messages.....                    | 3476             | bgp statement.....                                  | 3501       |
| NSR synchronization.....                   | 3476             | bgp-orf-cisco-mode                                  |            |
| open PDUs.....                             | 3476             | usage guidelines.....                               | 3313       |
| policy processing.....                     | 3476             | bgp-orf-cisco-mode statement.....                   | 2871, 3502 |
| protocol task processing.....              | 3476             | Bidirectional Forwarding Detection <i>See</i> BFD   |            |
| protocol timer processing.....             | 3476             | block statement                                     |            |
| refresh PDUs.....                          | 3476             | STP.....                                            | 2073       |
| route information.....                     | 3476             | bmp statement.....                                  | 2872       |
| state transitions.....                     | 3476             | boilerplate                                         |            |
| update PDUs.....                           | 3476             | commit scripts.....                                 | 6095       |
| update messages.....                       | 3148             | op scripts.....                                     | 6097       |
| version supported.....                     | 3144             | boot messages, displaying.....                      | 923        |

- boot server
    - NTP .....180
  - boot-server statement
    - NTP .....243
    - usage guidelines.....180
  - BOOTP relay agent.....4907
  - bootp statement.....4940
  - bootstrap (tracing flag).....4419
  - bootstrap messages.....4217, 4276
  - bootstrap routers
    - overview.....4217
  - bootstrap routers, displaying.....4587
  - bootstrap statement.....4364
  - bootstrap-export statement.....4365
  - bootstrap-import statement.....4366
  - bootstrap-priority statement.....4367
  - Border Gateway Protocol *See* BGP
  - bpdu (tracing flag).....2137
  - bpdu-block statement
    - STP.....2074
  - bpdu-block-on-edge statement
    - STP.....2075
  - bpdu-timeout-action statement
    - STP.....2076
  - braces, in configuration statements.....cxxxiv
  - brackets
    - angle, in syntax descriptions.....cxxxiv
    - square, in configuration statements.....cxxxiv
  - bridge-detection-state-machine (tracing flag).....2138
  - bridge-priority statement.....2077
  - brief statement.....2873
    - system logging.....296, 6379
    - usage guidelines.....6215
  - broadcast
    - NTP.....52, 53, 153, 154, 156
    - synchronizing NTP.....160
  - broadcast messages, synchronizing NTP.....245
  - broadcast statement.....244, 2078, 4941
    - usage guidelines.....154, 2040
  - broadcast-client statement.....245
    - usage guidelines.....160
  - BSR
    - policy, import.....4380
  - bucket-size statement.....6286
    - ICMPv4.....260
    - usage guidelines.....151
  - buffer-partition statement.....5824, 5826
  - buffer-size statement.....5828
  - buffers, displaying system.....930
  - built-in Ethernet interfaces.....2434
    - See also* Fast Ethernet ports; Gigabit Ethernet ports
  - burst-size-limit statement.....4785
- ## C
- ca-name statement.....4922
  - cable-length statement.....5830
  - cables
    - console port,
      - connecting.....1159, 1718, 6478, 6525
    - Ethernet rollover,
      - connecting.....1159, 1718, 6478, 6525
  - cache (tracing flag).....4419
  - cache-size statement.....4921
  - cache-timeout-negative statement .....4922
  - categories statement.....1763, 6286
    - usage guidelines.....6195
  - CB
    - environmental information, displaying.....561
    - Ethernet switch, displaying port
      - information.....491, 1478
    - operation of, controlling.....360
  - CBT
    - issues.....4209
  - CCC, graceful restart.....2269
  - certificates
    - installed, displaying.....937
  - certificates statement.....4923
  - certification-authority statement.....4924
  - change-log (system logging facility).....6222
  - change-type statement.....245
    - usage guidelines.....1687
  - chassis
    - alarm conditions, displaying.....452
    - configuration
      - alarm conditions.....147
    - craft interface display messages
      - clearing the display of.....328
      - displaying.....1469
      - stopping the display of.....1469
    - environmental information, displaying.....508
    - Ethernet switch information,
      - displaying.....466, 491, 1478, 1495
    - firmware version, displaying.....631
    - installed hardware, displaying.....667
    - location, displaying.....808
    - MAC addresses, displaying.....812



|                                                 |                                 |                                                  |            |
|-------------------------------------------------|---------------------------------|--------------------------------------------------|------------|
| serial numbers, displaying.....                 | 667                             | clear fibre-channel fip enode command.....       | 5277       |
| temperature threshold settings,                 |                                 | clear fibre-channel fip statistics command.....  | 5278       |
| displaying.....                                 | 852                             | clear fibre-channel fip vn-port command.....     | 5279       |
| Chassis Definitions for Router Model MIB.....   | 6130                            | clear fibre-channel flogi statistics             |            |
| Chassis MIB.....                                | 6130, 6137                      | command.....                                     | 5280       |
| chassis software process.....                   | 68, 1662                        | clear fibre-channel proxy statistics             |            |
| chassis statement.....                          | 1382, 2557, 2558                | command.....                                     | 5281       |
| chassis-scheduler (tracing flag)                |                                 | clear fip snooping enode command.....            | 5282       |
| class of service.....                           | 5908                            | clear fip snooping statistics command.....       | 5283       |
| chassisd process.....                           | 68, 1662                        | clear fip snooping vlan command.....             | 5284       |
| check-zero statement.....                       | 4178                            | clear fip vlan-discovery statistics command..... | 5285       |
| checksum                                        |                                 | clear firewall command.....                      | 4860       |
| calculating for a                               |                                 | clear igmp membership command.....               | 4491       |
| file.....                                       | 339, 340, 341, 1634, 1635, 1636 | clear igmp statistics command.....               | 4495       |
| displaying values for all files.....            | 915                             | clear igmp-snooping membership                   |            |
| for IS-IS.....                                  | 3687                            | command.....                                     | 4494       |
| checksum statement.....                         | 246, 3699, 6239                 | clear igmp-snooping statistics command.....      | 4497       |
| circuit cross-connect See CCC                   |                                 | clear ipv6 neighbors command.....                | 2959       |
| Cisco non-deterministic, BGP MED option.....    | 3209                            | clear isis adjacency command.....                | 3756       |
| cisco-non-deterministic option.....             | 3332, 3563                      | clear isis database command.....                 | 3758       |
| class statement.....                            | 5835, 5836                      | clear isis overload command.....                 | 3760       |
| assigning to user.....                          | 248                             | clear isis statistics command.....               | 3762       |
| login.....                                      | 247                             | clear lldp neighbors command.....                | 1834       |
| usage guidelines.....                           | 1681, 1683, 1692                | clear lldp statistics command.....               | 1835       |
| Class-of-Service MIB.....                       | 6130, 6137                      | clear log command.....                           | 327, 1630  |
| class-of-service statement.....                 | 5831                            | clear msdp cache command.....                    | 4498       |
| classifiers statement.....                      | 5837                            | clear msdp statistics command.....               | 4499       |
| classifiers, CoS.....                           | 5915                            | clear multicast bandwidth-admission              |            |
| cleanup, storage space.....                     | 423                             | command.....                                     | 4500       |
| clear                                           |                                 | clear multicast scope command.....               | 4502       |
| snmp history.....                               | 6402                            | clear multicast sessions command.....            | 4503       |
| clear (ospf   ospf3) database command.....      | 4033                            | clear multicast statistics command.....          | 4504       |
| clear (ospf   ospf3) database-protection        |                                 | clear pim join command.....                      | 4505       |
| command.....                                    | 4036                            | clear pim register command.....                  | 4507       |
| clear (ospf   ospf3) io-statistics command..... | 4037                            | clear pim statistics command.....                | 4509       |
| clear (ospf   ospf3) neighbor command.....      | 4038                            | clear rip general-statistics command.....        | 4198       |
| clear (ospf   ospf3) overload command.....      | 4040                            | clear rip statistics command.....                | 4199       |
| clear (ospf   ospf3) statistics command.....    | 4041                            | clear sflow collector statistics command.....    | 6401       |
| clear arp inspection statistics command.....    | 4857                            | clear snmp history command.....                  | 6402       |
| clear bgp damping command.....                  | 3584                            | clear snmp statistics command.....               | 6403       |
| clear bgp neighbor command.....                 | 3585                            | clear spanning-tree statistics command.....      | 2190       |
| clear bgp table command.....                    | 3587                            | clear system commit command.....                 | 331, 1631  |
| clear chassis display message command.....      | 328                             | clear system reboot command.....                 | 332        |
| clear dhcp snooping binding command.....        | 4858                            | CLI                                              |            |
| clear ethernet-switching bpdu-error             |                                 | command history                                  |            |
| command.....                                    | 2185                            | displaying.....                                  | 881        |
| clear ethernet-switching port-error             |                                 | comparing configuration versions.....            | 1593, 1602 |
| command.....                                    | 4859                            | configuration mode                               |            |
| clear fibre-channel fc2 statistics command..... | 5276                            | description.....                                 | 80         |



- current working directory
  - displaying.....880
  - exiting to create UNIX-level shell.....1148
  - Junos OS, configuring using.....169, 170
  - permissions, displaying.....876
  - settings, displaying.....874
- client list
  - adding to SNMP community.....6196
- client mode, NTP.....52, 53, 153, 156
- client-list statement.....1763, 6287
  - usage guidelines.....6196
- client-list-name statement.....1764, 6287
  - usage guidelines.....6196
- client-response-ttl statement.....4941
- clients statement.....1764, 6288
  - usage guidelines.....6194
- CLNP (Connectionless Network Protocol).....3625
- CLNS (Connectionless Network Service) VPNs
  - BGP export policy.....3684
  - IS-IS.....3684
- cluster statement.....3504
  - usage guidelines.....3407
- clusters See BGP route reflectors
- code-point statement.....5839, 5840, 5841
  - class of service.....5838
- code-point-aliases statement.....5841
- code-points statement.....5842
- collector statement.....6268
- color statement
  - aggregate routes.....2918
  - generated routes.....2918
  - static routes.....2918
- color-aware statement.....4786
- color-blind statement.....4787
- command output
  - filtering
    - comparing configuration
      - versions.....1593, 1602
- command statement.....6240
- commands
  - allowing or denying to login classes.....1675
  - filenames, specifying.....73
  - options.....79
  - URLs, specifying.....73
- comments, in configuration statements.....cxxiv
- commit command.....322
  - usage guidelines.....81
- commit operations, pending
  - clearing.....331, 1631
  - displaying.....939, 1649
- commit scripts
  - boilerplate.....6095
  - commands for monitoring.....6398
  - flow of operation illustrated.....6085
  - Junos OS, configuring using.....169, 171
  - multiple.....6089
  - output, displaying.....6398
  - persistent configuration changes.....6090
  - transient configuration changes.....6090
  - using multiple.....6089
- commit statement.....249, 6241
- commit synchronize command.....322
- commit-delay statement.....1765, 6288
- committed-burst-size statement.....4788
- committed-information-rate statement.....4789
- communities
  - aggregate routes.....2875
  - static routes.....2875
- community ASN, displaying routes.....2993
- community name, displaying routes.....2995
- community statement
  - aggregate routes.....2875
  - generated routes.....2875
  - RMON.....6290
  - SNMP.....1766, 6289
    - usage guidelines.....6194
  - static routes.....2875
- community string, SNMP.....6194
- community-name statement.....6291
- compact flash
  - displaying usage.....317, 2629
- compare command
  - usage guidelines.....1593, 1602
- comparing files.....342, 1637
- complete sequence number PDUs, IS-IS.....3628
- compress-configuration-files statement.....250
  - usage guidelines.....1595
- compressing configuration files.....250, 1595
- compressing files.....337, 1632
- confederation statement.....2877
  - usage guidelines.....3423
- confederations.....2877, 3422 See BGP
  - c o n f e d e r a t i o n s
- config-internal (tracing flag).....2950

|                                                |                        |
|------------------------------------------------|------------------------|
| configuration                                  |                        |
| activating.....                                | 1600                   |
| comparing with previous.....                   | 1593, 1602             |
| deleting rescue configuration.....             | 373, 1647              |
| displaying                                     |                        |
| archival configuration.....                    | 941, 1651              |
| previous configuration.....                    | 1015, 1654             |
| rescue configuration.....                      | 942, 1652              |
| files See configuration files                  |                        |
| merging current and new.....                   | 1598                   |
| previous, displaying.....                      | 1601                   |
| replacing.....                                 | 1597                   |
| saving rescue configuration.....               | 374, 1648              |
| saving to file.....                            | 1605, 1608             |
| storage of previous.....                       | 1591                   |
| syntax, verifying.....                         | 1150, 1656             |
| configuration files                            |                        |
| compressing.....                               | 250, 1595              |
| filename, specifying.....                      | 73                     |
| saving to files.....                           | 1605, 1608             |
| URL, specifying.....                           | 73                     |
| Configuration Management MIB.....              | 6131, 6138             |
| configuration mode commands                    |                        |
| commit script.....                             | 6398                   |
| configuration mode, CLI                        |                        |
| commands                                       |                        |
| activate.....                                  | 81                     |
| annotate.....                                  | 81                     |
| commit.....                                    | 81                     |
| copy.....                                      | 81                     |
| deactivate.....                                | 81                     |
| delete.....                                    | 81                     |
| edit.....                                      | 81                     |
| exit.....                                      | 81                     |
| extension.....                                 | 81                     |
| help.....                                      | 81                     |
| insert.....                                    | 81                     |
| load.....                                      | 81                     |
| paste.....                                     | 82                     |
| quit.....                                      | 82                     |
| rollback.....                                  | 82, 1607               |
| run.....                                       | 82                     |
| save.....                                      | 82                     |
| set.....                                       | 82                     |
| show.....                                      | 82                     |
| status.....                                    | 82                     |
| top.....                                       | 82                     |
| up.....                                        | 82                     |
| update.....                                    | 82                     |
| configuration hierarchy, description.....      | 84                     |
| description.....                               | 80                     |
| identifier, description.....                   | 83                     |
| statement                                      |                        |
| container.....                                 | 84                     |
| description.....                               | 83                     |
| leaf.....                                      | 84                     |
| top level statements, interpreting.....        | 83                     |
| configuration statement.....                   | 1617, 1767             |
| usage guidelines.....                          | 1611                   |
| configuration statements                       |                        |
| specifying IP addresses in.....                | 49                     |
| configuration-name statement                   |                        |
| STP.....                                       | 2080                   |
| configure command                              |                        |
| usage guidelines.....                          | 76                     |
| configuring                                    |                        |
| QFX3500 device.....                            | 163                    |
| QFX3500 switch.....                            | 6104                   |
| VLANs.....                                     | 2048                   |
| conflict-log (system logging facility).....    | 6222                   |
| congestion-notification-profile statement..... | 5844                   |
| connection-limit statement.....                | 1768, 6230             |
| connections                                    |                        |
| IP sockets, displaying active.....             | 944                    |
| SSH, opening.....                              | 1863                   |
| testing                                        |                        |
| general connections.....                       | 354, 6415              |
| connectivity                                   |                        |
| unidirectional (RIP).....                      | 4107                   |
| console port                                   |                        |
| adapter.....                                   | 1159, 1717, 6478, 6525 |
| console statement                              |                        |
| physical port.....                             | 251                    |
| system logging.....                            | 6370                   |
| usage guidelines.....                          | 6211                   |
| contact statement.....                         | 1769, 6292             |
| container-devices statement.....               | 2560                   |
| device-count.....                              | 2560                   |
| context-identifier statement                   |                        |
| OSPF.....                                      | 3974                   |
| Control Board See CB                           |                        |
| conventions                                    |                        |
| text and syntax.....                           | cxxiii                 |
| copy command                                   |                        |
| usage guidelines.....                          | 76, 81                 |

- 
- core dump files
    - usage guidelines.....175
    - viewing.....181
  - core-dumps
    - usage information, displaying.....963
  - CoS
    - byte count by PIC type
      - displaying.....2712, 6020
    - classifiers, monitoring.....5915
    - code point aliases, monitoring.....5920
    - congestion-notification, displaying.....5931
    - forwarding class sets, displaying.....5938
    - forwarding classes, monitoring.....5916
    - forwarding table, displaying
      - classifier information.....5946
      - code point value to loss priority.....5944
      - code point value to queue number.....5944
      - configuration.....5940
      - RED drop profiles.....5948
      - rewrite rules.....5950
      - rewrite rules, table identifiers.....5952
      - scheduler map information.....5953
    - inputs and outputs overview.....5442
    - interfaces, displaying.....5955
    - interfaces, monitoring.....5917
    - loss priority.....5920
    - mapping, displaying
      - code point aliases to bit patterns.....5929
      - code point value to forwarding
        - class.....5927
      - code point value to loss priority.....5927
      - forwarding classes to queue
        - numbers.....5936
      - schedulers to forwarding classes.....5984
    - MIB.....6130, 6137
    - packet loss priority.....5920
    - RED profile information, displaying.....5933
    - rewrite rules, displaying.....5982
    - rewrite rules, monitoring.....5918
    - scheduler map information, displaying.....5984
    - scheduler maps, monitoring.....5919
    - shared buffer, displaying.....5986
    - traffic-control profile information,
      - displaying.....5988
  - cos-adjustment (tracing flag)
    - class of service.....5908
  - cost statement
    - STP.....2079
  - CPU utilization, displaying.....315, 2628
  - craft interface
    - alarm conditions
      - overview.....147
  - craft interface display messages
    - clearing.....328
    - displaying
      - on the craft interface display.....1469
    - stopping.....1469
  - craft-lockout statement.....2561
  - creating emergency boot device.....165, 1583, 6526
  - credibility-protocol-preference
    - traffic engineering
      - IS-IS.....3750
  - critical (system logging severity level).....2909
  - crl statement
    - ES PIC.....4924
  - csn (tracing flag).....3747
  - csnp-interval statement.....3700
  - curly braces, in configuration statements.....cxxiv
  - current time, displaying.....1067
  - current working directory
    - displaying.....880
  - customer support.....cxxv
    - contacting JTAC.....cxxv
  - cut-through statement.....2080
- ## D
- daemon (system logging facility).....6223
    - option to facility-override statement.....6224
  - daemons See processes, software
  - damping.....3506, 3577
  - damping (tracing flag).....3577
  - damping parameters, BGP
    - clearing.....3584
    - displaying.....3615
  - damping routes, BGP
    - displaying.....2997, 3617
  - damping statement.....3506
    - usage guidelines.....3457, 3466
  - data-encapsulation statement.....4464
    - usage guidelines.....4332
  - database description packets.....3806
  - database-protection statement
    - OSPF.....3975
  - dcbx statement.....5221, 5846
  - deactivate command
    - usage guidelines.....81
  - dead-interval statement.....3977
    - usage guidelines.....3876

|                                                 |                        |                                       |                  |
|-------------------------------------------------|------------------------|---------------------------------------|------------------|
| debug (system logging severity level).....      | 2909                   | destination statement.....            | 255              |
| default setting                                 |                        | usage guidelines.....                 | 1697, 1714, 2040 |
| CoS.....                                        | 5440, 5443             | destination-port statement            |                  |
| default-address-selection statement.....        | 252                    | SNMP.....                             | 6293             |
| usage guidelines.....                           | 151                    | usage guidelines.....                 | 6195             |
| default-lsa statement.....                      | 3978                   | destination-profile statement         |                  |
| usage guidelines.....                           | 3826                   | usage guidelines.....                 | 2040             |
| default-metric statement                        |                        | detection-time statement              |                  |
| usage guidelines.....                           | 3822, 3826             | BFD.....                              | 2867             |
| default-peer statement.....                     | 4465                   | BGP.....                              | 3497, 3509       |
| usage guidelines.....                           | 4332                   | IS-IS.....                            | 3697             |
| defaults statement                              |                        | PIM.....                              | 4368             |
| aggregate statement.....                        | 2858                   | device-count statement.....           | 2563             |
| generate statement.....                         | 2888                   | dfc (system logging facility).....    | 6223             |
| static statement.....                           | 2943                   | DHCP                                  |                  |
| delay statement                                 |                        | relay agent.....                      | 4907             |
| OSPF.....                                       | 4018                   | dhcp-snooping-file statement.....     | 4808             |
| delay-med-update statement                      |                        | dhcp-trusted statement.....           | 4806, 4808, 4821 |
| usage guidelines.....                           | 3236                   | diagnosis                             |                  |
| delete command                                  |                        | verifying RIP host reachability ..... | 4115             |
| usage guidelines.....                           | 81                     | direct-access statement.....          | 256, 6242        |
| deleting                                        |                        | director-device statement             |                  |
| current rescue configuration (CLI configuration |                        | QFabric system component aliases..... | 1384             |
| editor).....                                    | 1609                   | directories                           |                  |
| files.....                                      | 345, 1640              | usage information, displaying.....    | 976              |
| licenses (CLI).....                             | 94                     | working, displaying.....              | 880              |
| software packages.....                          | 404                    | disable (DCBX) statement.....         | 5225, 5849       |
| demand-circuit statement                        |                        | disable statement.....                | 2307, 3577       |
| OSPF                                            |                        | BGP.....                              | 3510             |
| usage guidelines.....                           | 3846                   | IGMP.....                             | 4427             |
| deny-commands statement.....                    | 253                    | usage guidelines.....                 | 4316             |
| usage guidelines.....                           | 1675                   | IGMP snooping.....                    | 4450             |
| deny-configuration statement.....               | 254                    | IS-IS.....                            | 3701             |
| usage guidelines.....                           | 1675                   | LLDP.....                             | 1769             |
| denying commands to login classes.....          | 1675                   | MSDP.....                             | 4466             |
| description statement.....                      | 2081, 2562, 2878, 3508 | MVRP.....                             | 2148             |
| helper service or interface.....                | 4939, 4942             | OSPF.....                             | 3979             |
| op script arguments.....                        | 6242                   | PIM family.....                       | 4369             |
| op scripts.....                                 | 6242                   | PIM interfaces.....                   | 4369             |
| RMON.....                                       | 6293                   | PIM protocol.....                     | 4369             |
| SNMP.....                                       | 6292                   | STP.....                              | 2082             |
| usage guidelines.....                           | 3173                   | traceoption.....                      | 2950             |
| designated router.....                          | 4215                   | usage guidelines                      |                  |
| configuring.....                                | 3812                   | BGP.....                              | 1376, 2282, 2286 |
| controlling election.....                       | 3812                   | ES-IS.....                            | 2287             |
| OSPF.....                                       | 3809                   | global.....                           | 1375, 2282, 2285 |
| designated router, IS-IS                        |                        | IS-IS.....                            | 2287             |
| about.....                                      | 3687                   | OSPF .....                            | 1377, 2283, 2288 |
| configuring election priority.....              | 3687                   | OSPFv3.....                           | 1377, 2283, 2288 |

PIM.....2289  
 RIP.....2289  
 RIPng.....2289  
 disable-timeout statement  
   port-error-disable.....4809  
   STP.....2083  
 discard statement  
   aggregate routes.....2879  
   generated routes.....2879  
 disk space available, displaying.....1061  
 disk-failure-action statement.....2564  
 distance-vector routing protocols.....4103  
   *See also* RIP  
 distribution trees  
   RPT.....4220  
   shared.....4220  
 DNS  
   hostnames, displaying.....882  
 DNS name servers.....146  
 documentation  
   comments on.....cxxv  
 Domain Name System *See* DNS  
 domain names on routers.....146  
 domain-name statement  
   router.....257  
   usage guidelines.....146  
 domain-search statement.....257  
   usage guidelines.....147  
 domains to be searched.....147, 257  
 dr-election-on-p2p statement.....4370  
   PIM  
     usage guidelines.....4254  
 dr-register-policy statement.....4370  
   usage guidelines.....4282  
 drop profiles, displaying data points.....5933  
 drop-profile statement.....5851  
 drop-profile-map statement.....5851  
 drop-profiles statement.....5852  
 dscp ipv6 statement.....5855  
 dscp statement.....5853  
 dscp-code-point statement.....5856  
 DVMRP  
   groups, displaying.....4584  
 dynamic (tracing flag)  
   class of service.....5908  
 dynamic IGMP statements  
   promiscuous-mode  
     interface.....4438  
 dynamic overload bit, resetting for IS-IS.....3760

## E

EBGIP *See* BGP  
 EBGIP (external BGP)  
   route-flap damping.....3456  
 EBGIP IPv6 peering, BGP.....3157  
 edge statement.....2085  
 edit command  
   usage guidelines.....81  
 egress statement.....5857  
   port mirroring.....4912  
 embedded-rp statement.....4371  
 emergency (system logging severity level).....2909  
 emergency boot device  
   creating.....165, 1583, 6526  
 enable IGMP static group membership.....4305  
 enable statement  
   routing options.....2895  
 encoding statement.....4925  
 encrypted passwords.....158, 1687, 1702  
 encrypted-password option.....158, 1687, 1702  
 engine-id statement  
   SNMPv3.....6294  
 enhanced-transmission-service  
   statement.....5226, 5858  
 enrollment-retry statement.....4925  
 enrollment-url statement.....4926  
 environmental information  
   CB, displaying.....561  
   chassis, displaying.....510  
   FPC, displaying.....579  
   PEMs, displaying.....604  
   Routing Engines, displaying.....613  
 erasing user data.....433  
 error (system logging severity level).....2909  
 error (tracing flag)  
   IS-IS.....3747  
   OSPF.....4023  
   RIP.....4193  
 ES-IS  
   graceful restart.....2269  
 ether-options statement.....2567  
 Ethernet interfaces  
   status information, displaying  
     Gigabit  
       Ethernet.....2217, 2651, 2672, 2689, 2765, 2783  
 Ethernet MAC MIB.....6131  
 Ethernet rollover cable, connecting the switch to a  
   management device.....1159, 1718, 6478, 6525  
 ethernet statement.....2564, 2565, 2571, 2597

|                                           |                              |
|-------------------------------------------|------------------------------|
| Ethernet switch information,              |                              |
| displaying.....                           | 466, 491, 1478, 1495         |
| ethernet-switching-options                |                              |
| statement.....                            | 1770, 2086, 4810, 4838, 4913 |
| eui-64 statement.....                     | 2088, 2568                   |
| usage guidelines.....                     | 2040                         |
| Event MIB.....                            | 6131                         |
| event recording                           |                              |
| IGMP.....                                 | 4312                         |
| event statement.....                      | 6295                         |
| event threshold                           |                              |
| symbol period .....                       | 2593                         |
| events (tracing flag)                     |                              |
| STP.....                                  | 2138                         |
| events statement.....                     | 259                          |
| usage guidelines.....                     | 1697, 1714                   |
| EX4200 configuration                      |                              |
| procedures.....                           | 1333                         |
| examine-dhcp statement.....               | 4812                         |
| examine-fip statement.....                | 4813, 5228                   |
| examine-vn2vn statement.....              | 5219, 5229                   |
| example                                   |                              |
| tracing ospf traffic.....                 | 3956                         |
| excess-burst-size statement.....          | 4790                         |
| exclude statement                         |                              |
| IGMP.....                                 | 4427                         |
| usage guidelines.....                     | 4305                         |
| exclude-cmd-attribute statement.....      | 1821                         |
| exit command                              |                              |
| from configuration mode.....              | 189                          |
| usage guidelines.....                     | 81                           |
| explicit-priority statement.....          | 258, 6371                    |
| usage guidelines                          |                              |
| single-chassis system.....                | 6213                         |
| export route information, displaying..... | 3019                         |
| export statement                          |                              |
| BGP.....                                  | 3512                         |
| usage guidelines.....                     | 3310                         |
| forwarding table.....                     | 2880                         |
| usage guidelines.....                     | 3363                         |
| IS-IS.....                                | 3702                         |
| MSDP.....                                 | 4467                         |
| OSPF.....                                 | 3981                         |
| PIM.....                                  | 4372                         |
| configuring.....                          | 4280                         |
| PIM RP                                    |                              |
| usage guidelines.....                     | 4276                         |
| RIP.....                                  | 4179                         |
| usage guidelines.....                     | 4110                         |

|                                             |      |
|---------------------------------------------|------|
| export statement, for routing policies..... | 3305 |
| export-rib statement.....                   | 2881 |
| extension command                           |      |
| usage guidelines.....                       | 81   |
| external-preference statement               |      |
| IS-IS.....                                  | 3703 |
| OSPF.....                                   | 3982 |
| usage guidelines.....                       | 3869 |
| external-router-id option.....              | 3332 |

## F

|                                      |                        |
|--------------------------------------|------------------------|
| fabric (tracing flag)                |                        |
| Fibre Channel.....                   | 5262                   |
| fabric statement.....                | 1385                   |
| fabric-id statement.....             | 5230                   |
| fabric-type statement.....           | 5230                   |
| facilities (system logging)          |                        |
| alternate for remote machine.....    | 6224                   |
| default for remote machine.....      | 6224                   |
| for local machine.....               | 6222                   |
| facility-override statement.....     | 6371                   |
| system logging                       |                        |
| usage guidelines.....                | 6225                   |
| failover-delay statement.....        | 2326                   |
| falling-event-index statement.....   | 6296                   |
| falling-threshold statement          |                        |
| health monitor.....                  | 1772, 6297             |
| RMON.....                            | 6298                   |
| falling-threshold-interval statement |                        |
| RMON.....                            | 6299                   |
| family statement.....                | 2566, 2569, 2579, 2580 |
| BGP.....                             | 3513                   |
| bootstrap.....                       | 4373                   |
| firewall filters.....                | 4776                   |
| IS-IS.....                           | 3704                   |
| local RP.....                        | 4375                   |
| PIM protocol.....                    | 4374                   |
| Fast Ethernet interfaces             |                        |
| loopback mode.....                   | 2535                   |
| fast-interval statement.....         | 2327                   |
| usage guidelines.....                | 2299                   |
| fate-sharing                         |                        |
| signaled LSPs.....                   | 2882                   |
| fate-sharing statement.....          | 2882                   |
| fault tolerance                      |                        |
| advertising multiple paths to a      |                        |
| destination.....                     | 3379, 3380             |
| fc-fabrics statement.....            | 5232                   |
| fc-map statement.....                | 4814, 5234             |

- 
- fc-options statement.....5235
  - fc2 (tracing flag)
    - Fibre Channel.....5262
  - fc2 statement.....5231
  - fcoe-trusted statement.....4815, 5236
  - feature licenses See licenses
  - FEB
    - firmware version, displaying.....631
  - fibre channel (interfaces) statement.....5237
  - fibre-channel (Port) statement.....5238
  - fibrechannel-options statement.....5238
  - file archive command.....337, 1632
  - file checksum md5 command.....339, 1634
  - file checksum sha-256 command.....341, 1636
  - file checksum sha1 command.....340, 1635
  - file command.....336
    - logical systems.....3477
    - usage guidelines.....76
  - file compare command.....342, 1637
  - file delete command.....345, 1640
  - file list command.....346, 1641
  - file rename command.....348, 1643
  - file show command.....350, 1645, 6212, 6385, 6499
  - file statement
    - commit scripts.....6243
    - op scripts.....6244
    - security certificate.....4926
    - system logging.....1386, 6372, 6373
    - usage guidelines.....6209
  - file systems
    - checksum values, displaying.....915
    - free disk space, displaying.....1061
  - filenames, specifying in commands.....73
  - files
    - archiving.....337, 1632
    - calculating
      - checksum.....339, 340, 341, 1634, 1635, 1636
    - comparing.....342, 1637
    - compressing.....337, 1632
    - configuration files, compressing.....250
    - configuration, compressing.....1595
    - contents, displaying.....350, 1645
    - deleting.....345, 1640
    - list of, displaying.....346, 1641
    - log file, clearing.....327, 1630
    - renaming.....348, 1643
    - saving configurations to files.....1605, 1608
    - system log messages, archiving.....6220
  - files statement.....6374
    - archiving of all system log files.....6366
    - archiving of individual system log file.....6368
    - system logging
      - usage guidelines.....6220
  - filter statement
    - firewall filters.....4777
    - interfaces.....4778
    - vlan.....2088, 4779
  - filter-duplicates statement.....1772, 6299
  - filter-interfaces statement.....6300
  - filter-specific statement.....4791
  - fip (tracing flag)
    - Fibre Channel.....5262
  - fip statement.....5239
  - firewall (system logging facility).....6223
  - Firewall MIB.....6131
  - firewall statement.....4780, 4792
  - firmware
    - chassis, displaying.....632
  - fka-adv-period statement.....5240
  - flags
    - login class.....1672
    - user permissions.....1672
  - flap damping.....3456
    - parameters.....3457
  - flash (tracing flag).....2950
  - logi (tracing flag)
    - Fibre Channel.....5262
  - flood reduction.....3860
  - flood-reduction statement.....3983
  - flooding (tracing flag).....4023
  - flow routes.....2883
  - flow statement.....2883, 3513
  - flow-control statement.....2572, 5861
  - flow-control-queue statement.....5862
  - flow-map statement.....2884
  - font conventions.....cxxxiii
  - format statement.....259
  - forward-delay statement.....2090
  - forwarding software process.....68, 1662
  - forwarding table
    - aggregate routes.....2873
    - CoS information, displaying
      - classifier information.....5946
      - code point value to loss priority.....5944
      - code point value to queue number.....5944
      - configuration.....5940
      - loss priority per classifier.....5944



|                                            |                  |
|--------------------------------------------|------------------|
| RED drop profiles.....                     | 5948             |
| rewrite rules.....                         | 5950             |
| scheduler map.....                         | 5953             |
| generated routes.....                      | 2873             |
| multicast information, displaying.....     | 4569             |
| policy, routing.....                       | 2880             |
| route entries, displaying.....             | 3039             |
| static routes.....                         | 2857, 2893, 2929 |
| forwarding-cache statement                 |                  |
| flow maps.....                             | 2885             |
| multicast.....                             | 2886             |
| forwarding-class statement.....            | 4816             |
| class of service.....                      | 5863, 5864, 5866 |
| forwarding-class-set statement             |                  |
| class of service.....                      | 5866             |
| forwarding-classes statement.....          | 5867             |
| forwarding-database (tracing flag)         |                  |
| Fibre Channel.....                         | 5262             |
| forwarding-table statement.....            | 2887             |
| usage guidelines.....                      | 3363             |
| FPC                                        |                  |
| environmental information, displaying..... | 579              |
| Ethernet switch, displaying port           |                  |
| information.....                           | 466, 1495        |
| firmware version, displaying.....          | 631              |
| installed, displaying list.....            | 668              |
| operation of, controlling.....             | 363              |
| status, displaying.....                    | 642              |
| fpc statement.....                         | 2574, 2575       |
| fragmentation                              |                  |
| avoiding.....                              | 3444             |
| free disk space, displaying.....           | 1061             |
| freeing up storage space.....              | 423              |
| from statement.....                        | 2914, 4781       |
| ftp (system logging facility).....         | 6223             |
| option to facility-override statement..... | 6225             |
| full mesh requirement                      |                  |
| fulfilling with confederations.....        | 3422             |
| fulfilling with route reflectors.....      | 3405             |
| full names, in user accounts.....          | 1681             |
| full statement.....                        | 2873             |
| full-name statement.....                   | 1773             |
| usage guidelines.....                      | 1681, 1692       |
| full-neighbors statement                   |                  |
| OSPF                                       |                  |
| usage guidelines.....                      | 3883             |
| fwdd process.....                          | 68, 1662         |

## G

|                                     |                                    |
|-------------------------------------|------------------------------------|
| ge-0/0/0, management interface..... | 2434                               |
| See also Gigabit Ethernet ports     |                                    |
| general (tracing flag).....         | 2950                               |
| generate statement.....             | 2888                               |
| generated routes.....               | 2888                               |
| preferences.....                    | 3326                               |
| Gigabit Ethernet interfaces         |                                    |
| diagnostics information,            |                                    |
| displaying.....                     | 883, 2645                          |
| loopback mode.....                  | 2535                               |
| status information,                 |                                    |
| displaying.....                     | 2217, 2651, 2672, 2689, 2765, 2783 |
| global tracing operations.....      | 6081                               |
| graceful restart                    |                                    |
| commands, operational mode.....     | 2343                               |
| concepts.....                       | 2269                               |
| configuration procedure             |                                    |
| routing protocols.....              | 2285                               |
| disabling.....                      | 3894                               |
| disabling strict LSA checking.....  | 3904                               |
| enabling.....                       | 3894                               |
| grace period interval.....          | 3894                               |
| OSPFv2 helper mode                  |                                    |
| disabling.....                      | 3898                               |
| re-enabling.....                    | 3898                               |
| OSPFv3 helper mode                  |                                    |
| disabling.....                      | 3901                               |
| re-enabling.....                    | 3901                               |
| overview.....                       | 3892                               |
| protocol support.....               | 2269                               |
| trace options.....                  | 1378, 2284, 2291                   |
| verifying operation of.....         | 2343                               |
| graceful-restart (tracing flag)     |                                    |
| IS-IS.....                          | 3747                               |
| OSPF.....                           | 4023                               |
| graceful-restart                    |                                    |
| statement.....                      | 1387, 1388, 2309, 2310, 3517       |
| BGP.....                            | 1388, 2310, 3517                   |
| OSPF.....                           | 1389, 2311, 3984                   |
| usage guidelines.....               | 3892                               |
| BGP.....                            | 1376, 2282, 2286                   |
| global.....                         | 1375, 2282, 2285                   |
| grace period.....                   | 3894                               |
| helper mode.....                    | 3898, 3901                         |
| IS-IS.....                          | 2287                               |
| OSPF.....                           | 1377, 2283, 2288                   |
| OSPFv3.....                         | 1377, 2283, 2288                   |
| PIM.....                            | 2289                               |



- RIP.....2289
- RIPng.....2289
- strict LSA checking.....3904
- graft (tracing flag)
  - PIM.....4419
- gratuitous-arp-reply.....2575
- group joins
  - limiting.....4313
- group membership
  - SSM maps.....4350
- group statement
  - BGP.....3518
  - IGMP.....4428
    - usage guidelines.....4305
  - IGMP snooping.....4450
  - MSDP.....4468
  - PIM RPF selection.....4376
  - RIP.....4180
  - SNMPv3 (for configuring).....6301
  - SNMPv3 (for security name).....6300
- group-count statement
  - IGMP.....4429
    - usage guidelines.....4305
- group-increment statement
  - IGMP.....4429
    - usage guidelines.....4305
- group-limit statement.....2091
  - configuring.....4313
  - IGMP interface.....4430
- group-policy statement
  - IGMP.....4431
    - usage guidelines.....4301
- group-ranges statement.....4377
  - usage guidelines.....4266
- groups
  - BGP
    - general information, displaying.....3589
  - DVMRP, displaying.....4584
  - IGMP membership, displaying.....4525
  - OSPF areas.....3817
  - PIM
    - general information, displaying.....4592
    - usage information, displaying.....4584
  - RIP routers.....4110
  - SSM.....4484
- groups statement.....4933
- guaranteed-rate statement.....5869

## H

- halts
  - pending
    - clearing.....332
    - displaying.....1010
    - requesting.....375, 1451
- hardware
  - major (red) alarm conditions on.....6487
- hardware, installed, displaying.....667
- hardware-database (tracing flag)
  - class of service.....5908
- health monitor alarms, displaying.....6445
- health-monitor statement.....1773, 6302
- hello (tracing flag)
  - IS-IS.....3747
  - PIM.....4419
- hello packets
  - IS-IS.....3628
  - OSPF.....3806
- hello-authentication-key statement.....3705
- hello-authentication-key-chain statement.....3706
- hello-authentication-type statement.....3707
- hello-interval statement
  - IS-IS.....3708
  - OSPF.....3986
    - usage guidelines.....3876
  - PIM.....4378
    - usage guidelines.....4246
- hello-padding statement.....3709
- hello-time statement.....2092
- help command
  - usage guidelines.....81
- helper-disable statement.....2313, 3987
  - usage guidelines
    - IS-IS.....2287
    - OSPF .....1377, 2283, 2288
    - OSPFv3.....1377, 2283, 2288
- history statement
  - RMON.....6303
- history, CLI commands
  - displaying.....881
- hold-down-interval
  - BGP.....3521
- hold-multiplier statement.....1774
- hold-time statement.....2328
  - BGP.....3523
  - IS-IS.....3711
  - Physical Interface.....2577
  - PIM.....4379

|                                       |            |                                              |            |
|---------------------------------------|------------|----------------------------------------------|------------|
| holddown (tracing flag).....          | 4193       | enabling.....                                | 4297, 4432 |
| holddown statement                    |            | event recording.....                         | 4312       |
| OSPF.....                             | 4018       | group membership                             |            |
| RIP.....                              | 4182       | SSM maps for different groups to different   |            |
| usage guidelines.....                 | 4160       | sources.....                                 | 4350       |
| holddown-interval statement           |            | group membership, displaying.....            | 4525       |
| BFD                                   |            | host-query message interval.....             | 4299, 4438 |
| static routes.....                    | 2867       | interface group limit.....                   | 4430       |
| hop count, maximizing.....            | 4106       | interfaces, displaying.....                  | 4529       |
| <i>See also</i> RIP                   |            | last-member query interval.....              | 4299, 4439 |
| host reachability                     |            | overview.....                                | 4230       |
| verifying for RIP network hosts.....  | 4115       | PIM-to-IGMP message translation information, |            |
| Host Resources MIB.....               | 6131, 6138 | displaying.....                              | 4565       |
| host statement.....                   | 6375       | query response interval.....                 | 4303, 4440 |
| system logging                        |            | robustness variable.....                     | 4304, 4441 |
| usage guidelines for single-chassis   |            | static group membership.....                 | 4305       |
| system.....                           | 6210       | statistics, displaying.....                  | 4533       |
| host-name statement.....              | 260        | tracing operations.....                      | 4314       |
| usage guidelines.....                 | 147        | version.....                                 | 4298, 4448 |
| host-outbound traffic statement.....  | 5870       | IGMP snooping                                |            |
| hostnames                             |            | group statement.....                         | 4450       |
| IS-IS, displaying.....                | 3777       | static statement.....                        | 4450       |
| hostnames, DNS, displaying.....       | 882        | igmp statement.....                          | 4432       |
| hosts, reachability                   |            | usage guidelines.....                        | 4297       |
| general connections.....              | 354, 6415  | IGMP statements                              |            |
| <b>I</b>                              |            | promiscuous-mode                             |            |
| IBGP <i>See</i> BGP                   |            | interface.....                               | 4438       |
| overview.....                         | 3172       | igmp-snooping statement.....                 | 4452       |
| IBGP (internal BGP)                   |            | IGMPv3.....                                  | 4231       |
| full mesh (configuration editor)..... | 3150       | interoperability with older versions.....    | 4231       |
| icmpv4-rate-limit statement.....      | 260        | ignore-attached-bit statement.....           | 3712       |
| usage guidelines.....                 | 151        | ignore-lsp-metrics statement                 |            |
| identifiers                           |            | OSPF.....                                    | 3988       |
| BGP <i>See</i> BGP, identifier        |            | IGP plus MED, BGP option.....                | 3209       |
| idle time, displaying.....            | 317, 2629  | igp shortcuts                                |            |
| idle timeout values                   |            | overview.....                                | 3908       |
| login classes.....                    | 162        | immediate-leave statement                    |            |
| idle-timeout statement.....           | 261, 1775  | IGMP.....                                    | 4434       |
| usage guidelines.....                 | 162        | usage guidelines.....                        | 4300       |
| ieee-802.1 statement.....             | 5873, 5874 | import statement.....                        | 5875       |
| class of service.....                 | 5872       | BGP.....                                     | 3525       |
| ieee-802.1 statement.....             | 5871       | usage guidelines.....                        | 3310       |
| if-exceeding statement.....           | 4793       | bootstrap.....                               | 4380       |
| ifd process.....                      | 68, 1662   | usage guidelines.....                        | 4276       |
| IGMP.....                             | 4523       | MSDP.....                                    | 4469       |
| configuration statements.....         | 4295       | OSPF.....                                    | 3989       |
| configuring.....                      | 4295       | PIM.....                                     | 4381       |
| disabling.....                        | 4316       | usage guidelines.....                        | 4281       |

- 
- RIP.....4183
    - usage guidelines.....4136
    - route resolution.....2889
  - import statement, for routing policies.....3305
  - import-policy statement.....2890
  - import-rib statement.....2891
  - include-mp-next-hop statement.....3526
  - incoming metric (RIP)
    - description.....4142
  - indirect next hop.....2892
  - indirect-next-hop statement.....2892
  - inet protocol family
    - interface addresses.....2041
  - inet statement.....295
    - usage guidelines.....168
  - inet6 protocol family
    - interface addresses.....2041
  - infinity statement.....4382
    - usage guidelines.....4287
  - info (system logging severity level).....2909
  - informs SNMP *See* SNMP informs
  - ingress statement.....4915, 5876
  - init (tracing flag)
    - class of service.....5908
  - input statement.....4916
    - class of service.....5878
  - insecure statement.....251, 6177
  - insert command
    - usage guidelines.....81
  - install statement.....2893
  - installation
    - licenses (CLI).....93
  - installing software.....394, 395
  - inte statement.....1391, 1396, 1398, 2581, 2607, 2608
  - inter-area-prefix-export statement
    - OSPFv3.....3990
    - usage guidelines.....3938
  - inter-area-prefix-import statement
    - OSPFv3.....3991
    - usage guidelines.....3946
  - interactive-commands (system logging facility).....6223
  - interconnect-device statement
    - QFabric system component aliases.....1392
  - interface
    - broadcast
      - configuring.....3839
      - monitoring.....317, 2629
    - point-to-multipoint
      - configuring.....3844
    - point-to-point
      - configuring.....3839
  - interface (DCBX) statement.....5241, 5879
  - interface (FIP) statement.....5243
  - interface (Storm Control)
    - statement.....2095, 2132, 4836, 4840, 4843
  - interface (tracing flag)
    - Fibre Channel.....5262
    - Fibre Channel proxy.....5267
  - interface addresses
    - logical interfaces.....2040
  - interface alarms
    - messages.....6491
  - Interface MIB.....6132, 6138
  - interface range
    - interfaces.....2432
  - interface software process.....68, 1662
  - interface statement.....4817
    - BOOTP.....4943
    - IGMP.....4435
      - usage guidelines.....4297
    - IGMP snooping.....4453
    - IS-IS.....3713
    - LLDP.....1776
    - multicast.....2897
    - multicast via static routes.....2895
    - MVRP.....2149
    - OSPF.....3992
      - usage guidelines.....3839, 3841
    - PIM.....4383
    - port mirroring.....4917, 4918
    - RMON history.....6305
    - SNMP.....6304
      - usage guidelines.....6197
    - STP.....2093, 2094, 2096
    - VRRP.....2329
      - usage guidelines.....2302
  - interface statistics, real-time, displaying.....2635
  - interface-range statement.....2582
  - interface-routes statement.....2898
  - interface-specific statement.....4782
  - interface-type statement.....3995
    - usage guidelines.....3841
  - interfaces
    - broadcast.....3838
    - demand circuit.....3838
    - descriptive text.....2562

|                                                  |                        |
|--------------------------------------------------|------------------------|
| monitor interface traffic                        |                        |
| command.....                                     | 6386, 6387, 6501       |
| NBMA.....                                        | 3838                   |
| overview.....                                    | 3838                   |
| passive.....                                     | 3838                   |
| passive traffic engineering mode.....            | 3838                   |
| peer.....                                        | 3838                   |
| point-to-multipoint.....                         | 3838                   |
| point-to-point.....                              | 3838                   |
| system logging.....                              | 6391, 6500             |
| tracing operations.....                          | 6082                   |
| interfaces (fibre-channel fabric) statement..... | 5242                   |
| interfaces descriptive text.....                 | 2878                   |
| interfaces limiting SNMP access.....             | 6197                   |
| interfaces statement.....                        | 2584, 6268             |
| class of service.....                            | 5880                   |
| internal routers                                 |                        |
| description.....                                 | 3803                   |
| overview.....                                    | 3804                   |
| Internet Group Management Protocol See IGMP      |                        |
| internet-options statement.....                  | 261                    |
| usage guidelines.....                            | 151                    |
| interpolate statement.....                       | 5881                   |
| interval statement                               |                        |
| health monitor.....                              | 1777, 6305             |
| RMON alarm.....                                  | 6306                   |
| RMON history.....                                | 6306                   |
| invalid routes, displaying.....                  | 3068                   |
| IP addresses.....                                | 168                    |
| 128-bit.....                                     | 2041                   |
| 32-bit.....                                      | 2041                   |
| router mapping.....                              | 184                    |
| router names, mapping.....                       | 168                    |
| specifying in configuration statements.....      | 49                     |
| IP multicast                                     |                        |
| announced sessions, displaying.....              | 4581                   |
| bandwidth admission                              |                        |
| clearing.....                                    | 4500                   |
| flow map information, displaying.....            | 4556                   |
| forwarding table, displaying.....                | 4569                   |
| interface information, displaying.....           | 4558                   |
| network information, displaying.....             | 4560                   |
| next-hop table, displaying.....                  | 4562                   |
| PIM-to-IGMP message translation information,     |                        |
| displaying.....                                  | 4565                   |
| PIM-to-MLD message translation information,      |                        |
| displaying.....                                  | 4567                   |
| RPF calculations, displaying.....                | 4575                   |
| scope, clearing.....                             | 4502                   |
| scoped information, displaying.....              | 4579                   |
| sessions, clearing.....                          | 4503                   |
| statistics                                       |                        |
| clearing.....                                    | 4504                   |
| tracing routes                                   |                        |
| from the receiver to the source.....             | 4512                   |
| from the source to the gateway                   |                        |
| router.....                                      | 4520                   |
| from the source to the receiver.....             | 4515                   |
| listen for responses.....                        | 4518                   |
| IP packets                                       |                        |
| router source addresses.....                     | 151, 252               |
| IP sockets, displaying active.....               | 944                    |
| IPv4 Protocol family                             |                        |
| interface addresses.....                         | 2040                   |
| ipv4-multicast statement.....                    | 3715                   |
| usage guidelines.....                            | 3670                   |
| ipv4-multicast-metric statement.....             | 3716                   |
| usage guidelines.....                            | 3670                   |
| IPv6                                             |                        |
| EBGP link-local peering.....                     | 3157                   |
| logical systems.....                             | 3157                   |
| neighbor cache information                       |                        |
| clearing.....                                    | 2959                   |
| displaying.....                                  | 2969                   |
| router advertisements                            |                        |
| displaying.....                                  | 2971                   |
| ipv6-multicast statement.....                    | 3716                   |
| ipv6-multicast-metric statement.....             | 3717                   |
| ipv6-unicast statement.....                      | 3718                   |
| ipv6-unicast-metric statement.....               | 3719                   |
| IS-IS                                            |                        |
| addresses.....                                   | 3626                   |
| adjacency database entries, clearing.....        | 3756                   |
| authentication.....                              | 3631, 3690, 3694, 3729 |
| CSNP.....                                        | 3729                   |
| hello.....                                       | 3730                   |
| hitless keychain.....                            | 3647                   |
| PSNP.....                                        | 3734                   |
| authentication keychain.....                     | 3695, 3706             |
| authentication, displaying.....                  | 3768                   |
| BFD.....                                         | 3697                   |
| complete sequence number                         |                        |
| PDUs.....                                        | 3628, 3700, 3747       |
| designated router.....                           | 3743                   |
| disabling.....                                   | 3701                   |
| IPv4 multicast topology.....                     | 3670                   |
| IPv4 unicast topology.....                       | 3670                   |
| IPv6 multicast topology.....                     | 3670                   |

- dynamic overload bit, resetting.....3760
- enabling.....3720
- errored packets.....3747
- graceful restart.....2269
- hello
  - interval.....3708
  - packet authentication.....3707
  - packet authentication key.....3705
  - PDUs.....3628, 3747
- hold time.....3711
- hold-down timer
  - disabling.....3728
- hostname database, displaying.....3777
- interfaces.....3713
- interfaces, displaying.....3779
- IPv4 unicast topology.....3735
- IPv6 unicast topology.....3718, 3734
- level properties
  - global.....3721
- link-state database entries
  - clearing.....3758
  - displaying.....3770
- link-state PDUs.....3628, 3747
  - interval.....3723
  - lifetime.....3724
  - tracing.....3747
- loose authentication.....3691, 3722
- mesh groups.....3726
- metrics.....3744
  - IPv6.....3719
  - multicast.....3716, 3717
  - normal.....3727
  - wide.....3753
- multicast topologies.....3669, 3670, 3715, 3716
  - IPv4.....3730
  - IPv6.....3732
- multilevel.....3639
- neighbors, displaying.....3764
- network PDUs.....3626
- NSAP.....3626
- overloaded, marking router as.....3736
- padding.....3709
- partial sequence number PDUs.....3628, 3747
- PDUs.....3628
- point-to-point interface.....3740
- policy, routing.....3702
- preferences.....3326, 3703, 3741
- prefix limit.....3742
- redistributing OSPF routes into.....3651
- routes, displaying.....3787
- SPF delay calculations.....3747
- topology.....3746
- tracing operations.....3747
- traffic engineering support.....3701, 3750
- traffic statistics
  - clearing.....3762
  - displaying.....3791
- IS-IS (Intermediate System-to-Intermediate System)
  - about designated routers.....3687
  - BFD authentication.....3665
  - checksum.....3687
  - configuring.....3633
  - configuring designated router election
    - priority.....3687
    - for CLNS route exchange.....3684
- isis statement.....3720
- ISO
  - addresses.....3626
  - system identifier.....3626
- ISO Protocol family.....2040
- J**
  - J-Web graphical user interface (GUI)
    - Junos OS, configuring using.....169, 170
  - join (tracing flag).....4419
  - join states, clearing PIM.....4505
  - join-load-balance statement.....4384
    - usage guidelines.....4256
  - join-prune-timeout statement.....4385
  - join-timer statement
    - MVRP.....2150
  - Juniper-Allow-Commands attribute
    - (RADIUS).....1669
  - Juniper-Allow-Configuration attribute
    - (RADIUS).....1669
  - Juniper-Configuration-Change attribute
    - (RADIUS).....1669
  - Juniper-Deny-Commands attribute
    - (RADIUS).....1669
  - Juniper-Deny-Configuration attribute
    - (RADIUS).....1669
  - Juniper-Interactive-Command attribute
    - (RADIUS).....1669
  - Juniper-Local-User-Name attribute
    - (RADIUS).....1668
  - Juniper-User-Permissions attribute
    - (RADIUS).....1670

|                                     |                           |                                            |                 |
|-------------------------------------|---------------------------|--------------------------------------------|-----------------|
| juniper.conf file, compressing..... | 250, 1595                 | Junos Space                                |                 |
| JUNOS Internet software             |                           | preparing for.....                         | 6104            |
| version, displaying.....            | 316, 2628                 | Junos XML management protocol              |                 |
| Junos OS                            |                           | Junos OS, configuring using.....           | 169, 171        |
| alarms, displaying.....             | 914                       | Junos-FIPS                                 |                 |
| boot messages, displaying.....      | 923                       | password requirements.....                 | 159, 1683, 1703 |
| buffers, displaying.....            | 930                       |                                            |                 |
| checksum values, displaying.....    | 915                       | <b>K</b>                                   |                 |
| core-dumps, displaying.....         | 963                       | keep statement.....                        | 3527            |
| directory usage, displaying.....    | 976                       | usage guidelines.....                      | 3312            |
| disk space, displaying.....         | 1061                      | keepalive (tracing flag)                   |                 |
| halt, requesting a.....             | 375, 1451                 | BGP.....                                   | 3577            |
| loaded extensions, displaying.....  | 1018                      | MSDP.....                                  | 4479            |
| methods for configuring.....        | 169                       | keepalive messages.....                    | 3148            |
| ASCII file.....                     | 169, 170                  | kernel (system logging facility).....      | 6223            |
| CLI.....                            | 169, 170                  | option to facility-override statement..... | 6225            |
| CLI, ASCII file, J-Web GUI.....     | 169                       | kernel (tracing flag).....                 | 2950            |
| commit scripts.....                 | 169, 171                  | kernel memory usage, displaying.....       | 1077            |
| J-Web GUI.....                      | 169, 170                  | key statement.....                         | 4927            |
| Junos XML management                |                           | key-chain statement.....                   | 4928            |
| protocol.....                       | 169, 171                  | BFD authentication.....                    | 4386            |
| NETCONF XML management              |                           | key-chain-name                             |                 |
| protocol.....                       | 169, 171                  | BGP.....                                   | 3529            |
| pending reboots                     |                           | keyboard sequences                         |                 |
| clearing.....                       | 332                       | used with monitor interface command.....   | 2635            |
| displaying.....                     | 1010                      | used with monitor interface traffic        |                 |
| processes, displaying.....          | 983                       | command.....                               | 2636            |
| reinstalling                        |                           | keychain                                   |                 |
| using removable media.....          | 111, 1581, 6528           | BGP.....                                   | 3430            |
| SRC client, displaying.....         | 1017                      | IS-IS.....                                 | 3631            |
| system statistics, displaying.....  | 1026                      | overview.....                              | 3429            |
| uptime, displaying.....             | 1067                      | krts (tracing flag)                        |                 |
| version, displaying                 |                           | Fibre Channel.....                         | 5263            |
| general.....                        | 1135                      |                                            |                 |
| virtual memory, displaying.....     | 1077                      | <b>L</b>                                   |                 |
| JUNOS software                      |                           | l3-interface statement.....                | 2098            |
| overview.....                       | 67, 1661                  | laptop See management device               |                 |
| Packet Forwarding Engine.....       | 68, 1661                  | ldap-url statement.....                    | 4929            |
| processes.....                      | 68, 1662                  | LDP                                        |                 |
| Routing Engine.....                 | 68, 1661                  | authentication algorithm.....              | 3494            |
| JUNOS Software                      |                           | graceful restart.....                      | 2269            |
| bundles, deleting.....              | 404                       | leave (tracing flag)                       |                 |
| packages, deleting.....             | 404                       | IGMP.....                                  | 4446            |
| powering off.....                   | 385                       | leave-timer statement                      |                 |
| rolling back.....                   | 416                       | MVRP.....                                  | 2151            |
| upgrade                             |                           | leaveall-timer statement                   |                 |
| performing.....                     | 394, 408, 410, 1461, 2393 | MVRP.....                                  | 2152            |
| validating candidate.....           | 420                       |                                            |                 |

- 
- level statement
    - IS-IS
      - protocol.....3721
      - usage guidelines.....3639
  - lib (tracing flag)
    - Fibre Channel.....5263
  - license infringement
    - verifying license usage.....97
    - verifying licenses installed.....96
  - license keys
    - components.....88, 1253
  - licenses.....88, 1253
    - adding.....381
    - adding (CLI).....93
    - deleting.....382
    - deleting (CLI).....94
    - displaying.....980
    - displaying (CLI).....96
    - displaying usage.....97
    - overview.....86, 1250
    - saving.....383
    - saving (CLI).....95
    - verifying.....96
  - lif (tracing flag)
    - Fibre Channel.....5263
  - link-mode statement.....2594
  - link-speed statement.....2595
  - link-state acknowledgment packets *See* OSPF, link-state acknowledgment packets
  - link-state advertisements *See* OSPF, link-state advertisements
  - link-state PDUs.....3628
  - lldp statement.....1778
  - lldp-configuration-notification-interval
    - statement.....1779
  - lo0 interface.....151, 252
  - load balancing
    - advertising multiple paths to a destination.....3379, 3380
    - for PIM join.....4256
  - load command.....352, 402, 403
    - usage guidelines.....81
  - load merge command
    - usage guidelines.....1598
  - load override command
    - usage guidelines.....1597
  - load set command
    - usage guidelines.....1599
  - load-balance statement
    - usage guidelines.....3363
  - load-balance-algorithm statement.....5244
  - load-key-file command
    - usage guidelines.....159, 1681, 1692, 1702
  - load-key-file statement.....262
    - usage guidelines.....158, 1681, 1687, 1692, 1702
  - local AS
    - BGP.....3246, 3251, 3261
  - local password authentication.....1701
  - local statement.....4930
    - PIM.....4387
      - usage guidelines.....4264
  - local user
    - template accounts.....1695
  - local-address statement.....2899
    - BFD.....2867
      - usage guidelines.....2831
    - BGP.....3531
    - MSDP group.....4470
    - MSDP peer.....4470
    - PIM.....4388
      - usage guidelines.....3173
  - local-as statement.....3533
    - usage guidelines.....3251, 3261
  - local-engine statement.....6307
  - local-interface statement
    - usage guidelines.....3157
  - local-ip-address statement
    - master.....2596
  - local-preference statement.....3536
    - usage guidelines.....3195
  - local0 - local7 (options to facility-override statement).....6225
  - location statement.....263, 4818
    - SNMP.....1779, 6308
      - usage guidelines.....156
  - location, chassis.....808
  - log files
    - clearing contents of.....327, 1630
    - contents, displaying.....889, 1567, 2360, 6437
    - specifying properties.....6220
  - log-out-on-disconnect statement.....251, 6177
  - log-prefix statement
    - system logging.....6377
      - usage guidelines.....6209
  - log-updown statement.....3537
    - BGP
      - usage guidelines.....3449

|                                                               |                        |
|---------------------------------------------------------------|------------------------|
| logging                                                       |                        |
| routing protocol process.....                                 | 2909                   |
| logging in as root.....                                       | 1710, 6179             |
| logging operations                                            |                        |
| tracing operations.....                                       | 6081                   |
| logging out users.....                                        | 384                    |
| logical interfaces                                            |                        |
| interface addresses.....                                      | 2040                   |
| logical operators                                             |                        |
| for monitor traffic command.....                              | 6408                   |
| logical systems                                               |                        |
| EBGP                                                          |                        |
| with IPv6 interfaces.....                                     | 3157                   |
| internal BGP.....                                             | 3184                   |
| viewing system files on.....                                  | 3477                   |
| logical-system statement.....                                 | 6309                   |
| login announcements, system.....                              | 149                    |
| login classes                                                 |                        |
| access privilege levels.....                                  | 1672                   |
| commands, allowing or denying.....                            | 1675                   |
| defining.....                                                 | 1345, 1683, 1737       |
| idle timeout values.....                                      | 162                    |
| QFabric system component.....                                 | 1345, 1737             |
| security configuration example.....                           | 1736                   |
| login messages, system.....                                   | 149                    |
| login statement.....                                          | 264                    |
| usage guidelines.....                                         | 1681, 1683, 1692, 1716 |
| login time, displaying.....                                   | 317, 2629              |
| login-alarms statement.....                                   | 265                    |
| usage guidelines.....                                         | 161                    |
| login-tip statement.....                                      | 265                    |
| logout, users.....                                            | 384                    |
| loopback mode.....                                            | 2535                   |
| loopback statement                                            |                        |
| Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet..... | 2597                   |
| Fast Ethernet interfaces                                      |                        |
| usage guidelines.....                                         | 2535                   |
| fibrechannel-options.....                                     | 5245                   |
| Gigabit Ethernet interfaces                                   |                        |
| usage guidelines.....                                         | 2535                   |
| loops statement                                               |                        |
| BGP address family.....                                       | 3538                   |
| loose-authentication-check statement                          |                        |
| IS-IS.....                                                    | 3722                   |
| usage guidelines.....                                         | 3691                   |
| loose-check                                                   |                        |
| BGP.....                                                      | 3540                   |
| loose-check statement                                         |                        |
| BFD authentication.....                                       | 4389                   |
| loss priority, CoS.....                                       | 5920                   |
| loss-priority statement                                       |                        |
| class of service.....                                         | 5882                   |
| system drop profiles.....                                     | 5883                   |
| loss-priority statement (rewrite rules)                       |                        |
| class of service.....                                         | 5884                   |
| lsa-refresh-interval statement.....                           | 3996                   |
| LSAs See OSPF, link-state advertisements                      |                        |
| lsp (tracing flag).....                                       | 3747                   |
| lsp-generation (tracing flag).....                            | 3747                   |
| lsp-interval statement.....                                   | 3723                   |
| lsp-lifetime statement.....                                   | 3724                   |
| LSPs                                                          |                        |
| fate-sharing.....                                             | 2882                   |
| <b>M</b>                                                      |                        |
| MAC addresses                                                 |                        |
| displaying.....                                               | 812                    |
| mac statement.....                                            | 4818                   |
| mac-limit statement.....                                      | 2099, 2874, 4819       |
| mac-move-limit statement.....                                 | 4820                   |
| mac-notification statement.....                               | 2100, 2111             |
| mac-table-aging-time statement.....                           | 2101                   |
| major (red) alarms                                            |                        |
| description.....                                              | 6487                   |
| management access, sample task.....                           | 1697                   |
| management device                                             |                        |
| recovering root password                                      |                        |
| from.....                                                     | 1159, 1717, 6478, 6524 |
| management interfaces                                         |                        |
| overview.....                                                 | 2434                   |
| management software process.....                              | 68, 1662               |
| manuals                                                       |                        |
| comments on.....                                              | cxxv                   |
| mappings                                                      |                        |
| SSM.....                                                      | 4485                   |
| martian addresses.....                                        | 2900                   |
| martians statement.....                                       | 2900                   |
| martians, displaying.....                                     | 3068                   |
| match conditions                                              |                        |
| for monitor traffic command.....                              | 6407                   |
| routing policy.....                                           | 3922                   |
| match statement.....                                          | 6377                   |
| usage guidelines.....                                         | 6227                   |
| max-age statement.....                                        | 2102                   |
| max-areas statement.....                                      | 3725                   |
| max-configurations-on-flash statement.....                    | 266                    |



- 
- max-hops statement.....2103
  - max-login-sessions statement.....5246
  - max-login-sessions-per-node statement.....5247
  - max-sessions-per-enode statement.....5248
  - maximum hop count, RIP.....4106
  - maximum statement
    - MSDP.....4471
    - usage guidelines.....4332
  - maximum-bandwidth statement.....2901
  - maximum-certificates statement.....4931
  - maximum-hop-count statement.....4944
  - maximum-length statement.....266
    - usage guidelines.....1687
  - maximum-paths statement.....2902
  - maximum-prefixes statement.....2904
  - maximum-rps statement.....4390
  - maximum-transmit-rate statement
    - IGMP.....4436
    - usage guidelines.....4305
  - MBGP MVPNs.....3466
  - mc-ae statement.....2559, 2598, 2599, 2602, 2616
  - MD5 authentication.....3430
    - BGP.....3430
  - MD5 checksum, calculating.....339, 1634
  - MED See BGP
  - MED (multiple exit discriminator)
    - always compare option.....3209
    - Cisco non-deterministic option.....3209
    - plus IGP option.....3209
  - med-igp-update-interval statement.....2905
    - usage guidelines.....3236
  - med-plus-igp statement.....3563
    - usage guidelines.....3332
  - member statement.....2599
  - member-range statement.....2600
  - members statement.....2104
    - usage guidelines.....3423
  - memory utilization, displaying.....315, 2628
  - mesh groups.....3726
    - MSDP.....4332
  - mesh-group statement.....3726
  - message statement.....267
    - usage guidelines.....149
  - message-processing-model statement.....6310
  - message-size statement.....4184
  - messages
    - boot, displaying.....923
    - broadcast messages, NTP.....160, 245
    - multicast, NTP.....160
    - redirect.....148
    - system login.....149
    - user screens, displaying on.....372
  - metric
    - traffic engineering.....3908
  - metric statement
    - aggregate routes.....2906
    - BGP
      - usage guidelines.....3223
    - generated routes.....2906
    - IS-IS.....3727
    - OSPF.....3997
      - usage guidelines.....3863
    - static routes.....2906
  - metric-in statement
    - RIP.....4185
    - usage guidelines.....4146
  - metric-out statement
    - BGP.....3542
    - usage guidelines.....3210
    - RIP.....4186
  - metrics
    - IS-IS.....3744
    - OSPF.....3861, 4014
    - RIP.....4186
  - mgd process.....68, 1662
  - MIBs
    - Alarm.....6130
    - Analyzer.....6130, 6137
    - Chassis.....6130, 6137
    - Chassis Definitions for Router Model.....6130
    - Class-of-Service.....6130, 6137
    - Configuration Management.....6131, 6138
    - Ethernet MAC.....6131
    - Event.....6131
    - Firewall.....6131
    - Host Resources.....6131, 6138
    - Interface.....6132, 6138
    - Ping.....6132
    - Power Supply.....6139
    - RMON Events and Alarms .....6132
    - SNMP object values, displaying.....6452
    - standards documents.....6125
    - Structure of Management Information.....6132
    - System Log.....6132
    - Utility.....6116
    - views
      - SNMP.....6197
    - VLAN.....6133

|                                     |            |                                             |                  |
|-------------------------------------|------------|---------------------------------------------|------------------|
| minimum-changes statement.....      | 267        | MLD                                         |                  |
| usage guidelines.....               | 1687       | group membership                            |                  |
| minimum-interval                    |            | SSM maps for different groups to different  |                  |
| BGP.....                            | 3544, 3545 | sources.....                                | 4350             |
| PIM.....                            | 4392       | PIM-to-MLD message translation information, |                  |
| usage guidelines.....               | 3349       | displaying.....                             | 4567             |
| minimum-interval statement          |            | mode (STP) statement.....                   | 2105             |
| BFD.....                            | 2867       | mode statement.....                         | 5222, 5847       |
| usage guidelines.....               | 2831       | MSDP.....                                   | 4472             |
| BGP.....                            | 3497       | usage guidelines.....                       | 4332             |
| IS-IS.....                          | 3697       | PIM.....                                    | 4394             |
| OSPF.....                           | 3971       | usage guidelines.....                       | 4255             |
| usage guidelines.....               | 3883       | monitor                                     |                  |
| PIM.....                            | 4391       | interface command output fields,            |                  |
| usage guidelines.....               | 4290       | table.....                                  | 6388, 6503       |
| RIP.....                            | 4175       | monitor interface command.....              | 2635             |
| usage guidelines.....               | 4123       | monitor interface traffic                   |                  |
| minimum-length statement.....       | 268        | command.....                                | 6386, 6387, 6501 |
| usage guidelines.....               | 1687       | monitor traffic command.....                | 6405             |
| minimum-links statement.....        | 2601       | monitoring                                  |                  |
| minimum-lower-cases statement.....  | 269        | interface.....                              | 317, 2629        |
| minimum-numeric statement.....      | 270        | routing tables.....                         | 2955             |
| minimum-punctuations statement..... | 271        | system process information.....             | 315, 2627        |
| minimum-receive-interval            |            | system properties.....                      | 316, 2628        |
| BGP.....                            | 3547       | Monitoring                                  |                  |
| minimum-receive-interval statement  |            | port security.....                          | 4847             |
| BFD.....                            | 2867       | monitoring tools                            |                  |
| usage guidelines.....               | 2831       | tracing operations.....                     | 6081             |
| BFD (BGP)                           |            | MP-BGP.....                                 | 3513             |
| usage guidelines.....               | 3349       | MPLS                                        |                  |
| BGP.....                            | 3497       | labels, displaying routes.....              | 3064             |
| IS-IS.....                          | 3697       | protocol family.....                        | 2040             |
| OSPF.....                           | 3971       | MPLS, graceful restart.....                 | 2269             |
| PIM.....                            | 4363, 4393 | mru statement.....                          | 5886             |
| usage guidelines.....               | 4290       | ms-chapv2                                   |                  |
| RIP.....                            | 4175       | changing password ms-chapv2.....            | 1700             |
| usage guidelines.....               | 4123       | MSDP                                        |                  |
| minimum-receive-ttl statement       |            | active source limit.....                    | 4462             |
| BFD.....                            | 2867       | maximum.....                                | 4471             |
| BFD (static routes)                 |            | per-source.....                             | 4477             |
| usage guidelines.....               | 2831       | threshold.....                              | 4478             |
| minimum-upper-cases statement.....  | 272        | authentication.....                         | 4463             |
| minimum-wait-time statement.....    | 4944       | cache entries, clearing.....                | 4498             |
| minor (yellow) alarms               |            | configuration statements.....               | 4327             |
| description.....                    | 6487       | configuring.....                            | 4327             |
|                                     |            | data-encapsulation.....                     | 4464             |
|                                     |            | default peer.....                           | 4332, 4465       |
|                                     |            | enabling.....                               | 4473             |
|                                     |            | general information, displaying.....        | 4545             |

- 
- groups.....4468
  - local address.....4470
  - message source information, displaying.....4547
  - mode.....4472
  - peer statistics
    - clearing.....4499
    - displaying.....4552
  - policy, routing.....4467, 4469
  - remote source.....4330
  - routing tables.....4476
  - source-active cache, displaying.....4549
  - tracing operations.....4328
  - msdp statement.....4473
  - msti statement.....2106
  - MSTP
    - configuration
      - displaying.....2246
  - mstp statement.....2107
  - mt (tracing flag).....4419
  - mtrace (tracing flag)
    - IGMP.....4314
  - mtrace command.....4512
  - mtrace from-source command.....4515
  - mtrace monitor command.....4518
  - mtrace to-gateway command.....4520
  - mtu statement.....2603
  - mtu-discovery statement.....3548
  - multi-chassis link aggregation interface
    - displaying.....2687
  - multi-chassis-protection
    - peer statement.....2611
  - multi-chassis-protection
    - statement.....2578, 2583, 2604
  - multi-destination statement.....5885
  - multiarea adjacency
    - OSPF.....3831
  - multiarea network.....3817
  - multicast
    - anycast RP.....4217
    - bootstrap router.....4217
    - NTP messages.....160
    - protocols
      - group membership.....4229
    - scoping.....2940
    - SSM groups.....4484
    - SSM mapping.....4485
  - multicast filters.....4218
    - MAC filters.....4219
    - MSDP SA messages.....4236
    - RP/DR register messages.....4219
    - configuring.....4282
  - multicast group joins
    - limiting.....4313
  - Multicast Source Discovery Protocol *See* MSDP
  - multicast statement.....2907
    - QFabric system routing options.....1393
  - multicast-client statement.....272
    - usage guidelines.....160
  - multicast-router-interface statement
    - IGMP snooping.....4454
  - multihop
    - BGP.....3317
  - multihop statement.....3550
    - usage guidelines.....3317
  - multilevel IS-IS.....3639
  - multipath statement
    - usage guidelines.....3363
  - multiple commit scripts.....6089
  - multiplier
    - BGP.....3552
  - multiplier statement
    - BFD.....2867
      - usage guidelines.....2831
    - BFD (BGP)
      - usage guidelines.....3349
    - BGP.....3497
    - IS-IS.....3697
    - OSPF.....3971
      - usage guidelines.....3883
    - PIM.....4363, 4394
      - usage guidelines.....4290
    - RIP.....4175
      - usage guidelines.....4123
  - multiprotocol BGP (MP-BGP).....3513
- ## N
- name servers, DNS.....146
  - name statement.....1780, 6310
  - name-server statement.....273
    - usage guidelines.....146
  - names
    - domain names on routers.....146
    - names.....184
    - router .....168
  - nas-ip-address statement .....1781

|                                                 |          |                                                |            |
|-------------------------------------------------|----------|------------------------------------------------|------------|
| native-vlan-id statement.....                   | 2108     | resolution database, displaying.....           | 3103       |
| neighbor statement                              |          | routes sent to, displaying.....                | 3070       |
| BGP.....                                        | 3553     | next-generation RIP See RIPng                  |            |
| OSPF                                            |          | NLRI, BGP.....                                 | 3148       |
| usage guidelines.....                           | 3844     | nlri-route-type statement                      |            |
| RIP.....                                        | 4187     | usage guidelines.....                          | 3466       |
| neighbor-policy statement.....                  | 4395     | no-accept-data statement.....                  | 2320       |
| usage guidelines.....                           | 4279     | usage guidelines.....                          | 2304       |
| neighbors                                       |          | no-accounting statement                        |            |
| BGP.....                                        | 3145     | IGMP.....                                      | 4426       |
| OSPF.....                                       | 3841     | no-adaptation                                  |            |
| NETCONF XML management protocol                 |          | BFD (BGP)                                      |            |
| Junos OS, configuring using.....                | 169, 171 | usage guidelines.....                          | 3349       |
| network                                         |          | BGP.....                                       | 3556       |
| masks.....                                      | 49       | PIM.....                                       | 4396       |
| network layer reachability information See BGP, |          | no-adaptation statement                        |            |
| NLRI                                            |          | BFD.....                                       | 2867       |
| network link advertisements.....                | 3807     | BFD (static routes)                            |            |
| network PDUs.....                               | 3626     | usage guidelines.....                          | 2831       |
| network service access point.....               | 3626     | BGP.....                                       | 3497       |
| Network Time Protocol See NTP                   |          | IS-IS.....                                     | 3697       |
| network-domain statement                        |          | OSPF.....                                      | 3971       |
| Director software partitions.....               | 1393     | RIP.....                                       | 4175       |
| network-summary-export statement                |          | usage guidelines.....                          | 4123       |
| usage guidelines.....                           | 3938     | no-adjacency-holddown statement.....           | 3728       |
| network-summary-import statement                |          | no-advertise-peer-as statement.....            | 3557       |
| usage guidelines.....                           | 3946     | usage guidelines.....                          | 3313       |
| networks                                        |          | no-aggregator-id statement.....                | 3558       |
| sample BGP confederations.....                  | 3423     | no-allow-url statement                         |            |
| sample BGP MED use.....                         | 3209     | op scripts.....                                | 6245       |
| sample BGP peer session.....                    | 3150     | no-authentication-check statement.....         | 3729       |
| sample BGP route reflector (one                 |          | usage guidelines.....                          | 3690       |
| cluster).....                                   | 3406     | no-broadcast statement.....                    | 4841       |
| sample BGP route reflectors (cluster of         |          | no-check-zero statement.....                   | 4178       |
| clusters).....                                  | 3407     | no-client-reflect statement.....               | 3559       |
| sample BGP route reflectors (multiple           |          | no-cmd-attribute-value statement.....          | 1821       |
| clusters).....                                  | 3406     | no-compress-configuration-files statement..... | 250        |
| sample distance-vector routing.....             | 4104     | usage guidelines.....                          | 1595       |
| sample multiarea OSPF routing.....              | 3814     | no-csnp-authentication statement.....          | 3729       |
| sample OSPF network with stubs and              |          | usage guidelines.....                          | 3690       |
| NSSAs.....                                      | 3821     | no-delegate-processing statement               |            |
| sample poison reverse routing.....              | 4107     | periodic packet management.....                | 2916       |
| sample split horizon routing.....               | 4106     | no-eligible-backup statement                   |            |
| sample unidirectional routing.....              | 4107     | OSPF.....                                      | 3999       |
| next hop, displaying.....                       | 2956     | no-fabric-wwn-verify statement.....            | 5249       |
| next hops                                       |          | no-gratuitous-arp-request statement.....       | 2606, 4822 |
| multicast entries, displaying.....              | 4562     | no-hello-authentication statement.....         | 3730       |
| PFE, displaying.....                            | 6051     | usage guidelines.....                          | 3690       |
|                                                 |          | no-install statement.....                      | 2893       |

- 
- no-ipv4-multicast statement.....3730
  - no-ipv4-routing statement.....3731
  - no-ipv6-multicast statement.....3732
  - no-ipv6-routing statement.....3733
  - no-ipv6-unicast statement.....3734
  - no-loopback statement.....2597
    - usage guidelines.....2535
  - no-make-before-break statement
    - QFabric system.....1394
  - no-multicast-echo statement.....274
    - PIM
      - usage guidelines.....4247
  - no-nssa-abr statement.....4000
    - usage guidelines.....3826
  - no-ping-record-route statement.....275
  - no-ping-time-stamp statement.....275
  - no-preempt statement.....2330
  - no-prepend-global-as statement
    - usage guidelines.....3251
  - no-psnp-authentication statement.....3734
    - usage guidelines.....3690
  - no-qos-adjust statement.....2908
  - no-readvertise statement.....2924
  - no-redirects statement.....276
    - usage guidelines.....148
  - no-remote-tracing.....6231
  - no-retain statement.....2929
  - no-rfc-1583 statement.....4001
    - usage guidelines.....3836
  - no-root-port statement
    - STP.....2110
  - no-saved-core-context statement.....287
    - usage guidelines.....175
  - no-strict-lsa-checking statement.....2314, 4002
    - usage guidelines
      - OSPF.....1377, 2283, 2288
      - OSPFv3.....1377, 2283, 2288
  - no-traps statement.....2621
  - no-unicast-topology statement.....3735
  - no-unknown-unicast statement.....4842
  - no-world-readable statement
    - archiving of all system log files.....6366
    - archiving of individual system log file.....6368
    - system logging
      - usage guidelines.....6220
  - node-device statement
    - QFabric system component aliases.....1395
  - node-group statement
    - Node group assignment.....1399
  - node-link-protection statement
    - OSPF.....4003
  - nodes statement
    - Node group assignment.....1397
  - nonvolatile statement.....1781, 6311
  - normal (tracing flag).....2950
  - not-so-stubby areas *See* NSSAs
    - configuring.....3826
  - notice (system logging severity level).....2909
  - notify statement.....6311
  - notify-duration statement.....2315, 4004
    - usage guidelines.....1377, 2283, 2288
  - notify-filter statement
    - for applying to target.....6312
    - for configuring.....6312
  - notify-view statement.....6313
  - NP\_Port load rebalancing
    - operation of, controlling.....5286
  - NPDUs.....3626
  - NSAP.....3626
  - nsr-synchronization (tracing flag).....4420
  - nssa
    - configuring.....3826
  - nssa statement.....4005
    - usage guidelines.....3826
  - NSSAs (not-so-stubby areas)
    - description.....3821
    - overview.....3804
  - NTP
    - authentication keys.....152, 153
    - boot server.....180
    - broadcast mode.....52, 53, 153, 154
    - client mode.....52, 53, 153, 156
    - configuration example.....185
    - configuring.....180
    - listening
      - for broadcast messages.....160, 245
      - for multicast messages.....160
    - peer status, displaying.....892
    - peer values, displaying.....894
    - security configuration example.....188
    - server mode.....155
    - symmetric active mode.....52, 53, 153, 154
  - ntp statement.....277
    - usage guidelines.....180

## O

|                                             |                              |
|---------------------------------------------|------------------------------|
| oid statement                               |                              |
| SNMP.....                                   | 1782, 6313                   |
| usage guidelines.....                       | 6197                         |
| SNMPv3.....                                 | 6314                         |
| oif-map statement                           |                              |
| IGMP.....                                   | 4436                         |
| on-disk-failure statement.....              | 2609                         |
| op scripts                                  |                              |
| boilerplate.....                            | 6097                         |
| flow of operation illustrated.....          | 6097                         |
| using.....                                  | 6096                         |
| op statement.....                           | 6246                         |
| open messages, BGP.....                     | 3147                         |
| Open Shortest Path First See OSPF See OSPF  |                              |
| operating system See JUNOS software         |                              |
| operational mode commands                   |                              |
| displaying output from commit scripts.....  | 6398                         |
| operators, regular expression.....          | 1686                         |
| system logging.....                         | 6228                         |
| optional statement.....                     | 278, 6247                    |
| options statement.....                      | 2909                         |
| order statement.....                        | 1783                         |
| ORF                                         |                              |
| BGP.....                                    | 3313                         |
| OSPF                                        |                              |
| adjacencies.....                            | 3963                         |
| area border routers.....                    | 3814, 3852 See area          |
| border routers                              |                              |
| areas.....                                  | 3814                         |
| configuring.....                            | 3963                         |
| See also area border routers; backbone      |                              |
| area; NSSAs; stub areas                     |                              |
| AS external link advertisements.....        | 3807                         |
| authentication.....                         | 3967                         |
| backbone.....                               | 3815, 3963                   |
| backbone area See backbone area             |                              |
| backup coverage                             |                              |
| displaying.....                             | 4043                         |
| backup neighbor paths.....                  | 4046                         |
| backup-spf-options statement.....           | 3968                         |
| bandwidth-based metrics.....                | 3969                         |
| configuring.....                            | 3867                         |
| BFD.....                                    | 3881, 3971                   |
| configuring, router identifier.....         | 3810                         |
| context identifier, displaying.....         | 4048                         |
| controlling designated router election..... | 3812                         |
| cost See OSPF, metrics                      |                              |
| database description packets.....           | 3806                         |
| database protection                         |                              |
| configuring.....                            | 3920                         |
| overview.....                               | 3919                         |
| statement.....                              | 3975                         |
| dead interval.....                          | 3876                         |
| default route.....                          | 3804                         |
| demand circuits.....                        | 3846                         |
| designated router.....                      | 3809, 3841, 4012             |
| enabling.....                               | 3838, 3992, 4006             |
| enabling, description.....                  | 3810, 3815, 3817             |
| error packets.....                          | 4023                         |
| flood-reduction statement.....              | 3983                         |
| graceful restart.....                       | 1389, 2269, 2311, 3892, 3984 |
| hello interval.....                         | 3876, 3986                   |
| hello packets.....                          | 3806                         |
| interface types.....                        | 3995                         |
| interfaces, displaying.....                 | 4069                         |
| link-state                                  |                              |
| acknowledgment packets.....                 | 3807                         |
| advertisements.....                         | 3806, 3807, 3876, 4015       |
| flooding packets.....                       | 4023                         |
| request packets.....                        | 3806                         |
| update packets.....                         | 3806                         |
| link-state database entries, displaying     |                              |
| version 2.....                              | 4050                         |
| version 3.....                              | 4058                         |
| lsa-refresh-interval statement.....         | 3996                         |
| LSAs See OSPF, link-state advertisements    |                              |
| metrics.....                                | 3861, 3997, 4014             |
| multiarea adjacency                         |                              |
| configuring.....                            | 3832                         |
| overview.....                               | 3831                         |
| multiarea network.....                      | 3817                         |
| NBMA networks.....                          | 3841                         |
| neighbors.....                              | 3841, 4005                   |
| clearing connections.....                   | 4038                         |
| displaying.....                             | 4080                         |
| network link advertisements.....            | 3807                         |
| no-eligible-backup statement.....           | 3999                         |
| node-link-protection statement.....         | 4003                         |
| nonbroadcast, multiaccess networks.....     | 3841                         |
| NSSAs.....                                  | 3978                         |
| overload bit.....                           | 4007                         |
| overview                                    |                              |
| displaying.....                             | 2363, 4086                   |
| packets.....                                | 3805, 3807, 4023             |
| passive mode.....                           | 3848, 4009                   |

- policy, routing.....3981, 3989
  - network summaries.....3938, 3946
  - network summaries, overview.....3937
  - route install priority.....3933
- preferences.....3326, 3982, 4010
- prefix limit.....3858, 4011
- redistributing routes into IS-IS.....3651
- refresh.....3860
- route cost *See* OSPF, metrics
- route preference.....3800
- route selection.....3861
- route summarization.....3852, 3965
- router dead interval.....3977
- router identifier.....2938, 3810
- router link advertisements.....3807
- routing algorithm.....3800
- routing table entries, displaying.....4091
- single-area network.....3815
- SPF.....3800, 4023
- SPF calculations, displaying.....4077
- statistics, general
  - clearing.....4041
  - displaying.....4097
- statistics, I/O
  - clearing.....4037
  - displaying.....4075
- stub areas.....3978
- summary link advertisements.....3807
- tags
  - aggregate routes.....2946
  - generated routes.....2946
  - static routes.....2946
- timers.....3875
- topological database.....3798
- tracing operations.....4023
  - database description PDUs.....3954
  - demand circuit extensions.....3954
  - error PDUs.....3954
  - event.....3954
  - graceful restart.....3954
  - hello PDUs.....3954
  - LDP synchronization.....3954
  - link-state acknowledgement PDUs.....3954
  - link-state analysis PDUs.....3954
  - link-state PDUs.....3954
  - link-state request PDUs.....3954
  - link-state updates PDUs.....3954
  - NSR synchronization.....3954
  - packet dump.....3954
  - policy processing.....3954
  - protocol task processing.....3954
  - protocol timer processing.....3954
  - restart-signaling.....3954
  - route information.....3954
  - SPF calculations.....3954
  - state transitions.....3954
- traffic control.....3861
- traffic engineering
  - features.....4026
  - lsp metrics.....3988
  - support.....4017
- transmission delay.....3876, 4029
- transmit interval.....4028
- OSPF (Open Shortest Path First)
  - injecting OSPF routes into BGP.....3306, 3923
  - monitoring.....4031
- OSPF database protection.....3919
- OSPF metric
  - configuring.....3863
- OSPF reference bandwidth
  - configuring.....3863
- OSPF routing information.....4031
- ospf statement.....4006
- ospf3 statement.....4006
- OSPFv2
  - restart-signaling.....4023
- OSPFv3
  - overview.....3802
- other-routing-engine option to host
  - statement.....6375
- out-delay statement.....3560
  - usage guidelines.....3312
- out-of-band management interfaces.....2434
- outbound-route-filter
  - usage guidelines.....3313
- outbound-route-filter statement
  - BGP.....3561
- outgoing metric (RIP)
  - description.....4142
- output control keys
  - for monitor interface command.....2635
  - for monitor interface traffic command.....2636
- output statement
  - port mirroring.....4918
- output-traffic-control-profile statement.....5888
- overload bit, resetting for IS-IS.....3760

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| overload statement                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| IS-IS.....                                   | 3736                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| OSPF.....                                    | 4007                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| usage guidelines.....                        | 3872                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| overloaded                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| configuring routing devices.....             | 3872                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| routing devices.....                         | 3871                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| override-interval                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| PIM.....                                     | 4397                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| override-interval statement                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| usage guidelines.....                        | 4259                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| overview                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| CoS inputs and outputs.....                  | 5442                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| owner statement.....                         | 6314                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| <b>P</b>                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| Packet Forwarding Engine.....                | 68, 1661                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |  |
| packet headers, transmitted, displaying..... | 6405                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| packet loss priority, CoS.....               | 5920                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| packet-dump (tracing flag).....              | 4023                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| packet-rate statement                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| ICMPv4.....                                  | 260                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| usage guidelines.....                        | 151                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| packets                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| router source addresses.....                 | 151, 252                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |  |
| packets (tracing flag)                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| BGP.....                                     | 3577                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| IGMP.....                                    | 4446                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| IS-IS.....                                   | 3747                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| OSPF.....                                    | 4023                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| PIM.....                                     | 4420                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| RIP.....                                     | 4193                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| parameters statement.....                    | 6315                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| parentheses, in syntax descriptions.....     | cxxiv                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |  |
| parse (tracing flag).....                    | 2950                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| partial sequence number PDUs.....            | 3628                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| passive ARP learning, VRRP.....              | 2304                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| passive statement.....                       | 4009                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| aggregate routes.....                        | 2857                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| BGP.....                                     | 3562                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| usage guidelines.....                        | 3449                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| generated routes.....                        | 2857                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| IGMP.....                                    | 4437                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| IS-IS.....                                   | 3739                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| MVRP.....                                    | 2153                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| OSPF                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| usage guidelines.....                        | 3848                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| static routes.....                           | 2857                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| passive traffic-engineering mode             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| support.....                                 | 3916                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| password                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| for RIPv2 authentication.....                | 4121                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| root, setting .....                          | 189                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| ssh public string.....                       | 189                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| password statement                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| login.....                                   | 278                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |  |
| passwords                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| RADIUS.....                                  | 1699                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| root.....                                    | 158, 1687, 1702                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| root password,                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| recovering.....                              | 1159, 1717, 6478, 6524                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |  |
| shared user.....                             | 1701                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| passwords statement                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| usage guidelines.....                        | 1687                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| paste command                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| usage guidelines.....                        | 82                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |  |
| path attributes, BGP.....                    | 3146, 3148                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |  |
| path cost metrics                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| for RIP routes, description.....             | 4142                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| for RIP routes, modifying.....               | 4143, 4144                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |  |
| path-count statement                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| BGP                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| usage guidelines.....                        | 3380                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| path-length statement.....                   | 4931                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| path-selection statement.....                | 3563                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| usage guidelines.....                        | 3332                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| PC See management device                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| PDUs.....                                    | 3628                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| peak-burst-size statement.....               | 4795                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| peak-information-rate statement.....         | 4796                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| peer interfaces                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| configuring.....                             | 3850                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| peer                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| statement.....                               | 2925, 3555, 3562, 3566, 3602, 3604, 3605, 3606, 3607, 3608, 3609, 3610, 3611, 3612, 3613, 3614, 3615, 3616, 3617, 3618, 3619, 3620, 3621, 3622, 3623, 3624, 3625, 3626, 3627, 3628, 3629, 3630, 3631, 3632, 3633, 3634, 3635, 3636, 3637, 3638, 3639, 3640, 3641, 3642, 3643, 3644, 3645, 3646, 3647, 3648, 3649, 3650, 3651, 3652, 3653, 3654, 3655, 3656, 3657, 3658, 3659, 3660, 3661, 3662, 3663, 3664, 3665, 3666, 3667, 3668, 3669, 3670, 3671, 3672, 3673, 3674, 3675, 3676, 3677, 3678, 3679, 3680, 3681, 3682, 3683, 3684, 3685, 3686, 3687, 3688, 3689, 3690, 3691, 3692, 3693, 3694, 3695, 3696, 3697, 3698, 3699, 3700, 3701, 3702, 3703, 3704, 3705, 3706, 3707, 3708, 3709, 3710, 3711, 3712, 3713, 3714, 3715, 3716, 3717, 3718, 3719, 3720, 3721, 3722, 3723, 3724, 3725, 3726, 3727, 3728, 3729, 3730, 3731, 3732, 3733, 3734, 3735, 3736, 3737, 3738, 3739, 3740, 3741, 3742, 3743, 3744, 3745, 3746, 3747, 3748, 3749, 3750, 3751, 3752, 3753, 3754, 3755, 3756, 3757, 3758, 3759, 3760, 3761, 3762, 3763, 3764, 3765, 3766, 3767, 3768, 3769, 3770, 3771, 3772, 3773, 3774, 3775, 3776, 3777, 3778, 3779, 3780, 3781, 3782, 3783, 3784, 3785, 3786, 3787, 3788, 3789, 3790, 3791, 3792, 3793, 3794, 3795, 3796, 3797, 3798, 3799, 3800, 3801, 3802, 3803, 3804, 3805, 3806, 3807, 3808, 3809, 3810, 3811, 3812, 3813, 3814, 3815, 3816, 3817, 3818, 3819, 3820, 3821, 3822, 3823, 3824, 3825, 3826, 3827, 3828, 3829, 3830, 3831, 3832, 3833, 3834, 3835, 3836, 3837, 3838, 3839, 3840, 3841, 3842, 3843, 3844, 3845, 3846, 3847, 3848, 3849, 3850, 3851, 3852, 3853, 3854, 3855, 3856, 3857, 3858, 3859, 3860, 3861, 3862, 3863, 3864, 3865, 3866, 3867, 3868, 3869, 3870, 3871, 3872, 3873, 3874, 3875, 3876, 3877, 3878, 3879, 3880, 3881, 3882, 3883, 3884, 3885, 3886, 3887, 3888, 3889, 3890, 3891, 3892, 3893, 3894, 3895, 3896, 3897, 3898, 3899, 3900, 3901, 3902, 3903, 3904, 3905, 3906, 3907, 3908, 3909, 3910, 3911, 3912, 3913, 3914, 3915, 3916, 3917, 3918, 3919, 3920, 3921, 3922, 3923, 3924, 3925, 3926, 3927, 3928, 3929, 3930, 3931, 3932, 3933, 3934, 3935, 3936, 3937, 3938, 3939, 3940, 3941, 3942, 3943, 3944, 3945, 3946, 3947, 3948, 3949, 3950, 3951, 3952, 3953, 3954, 3955, 3956, 3957, 3958, 3959, 3960, 3961, 3962, 3963, 3964, 3965, 3966, 3967, 3968, 3969, 3970, 3971, 3972, 3973, 3974, 3975, 3976, 3977, 3978, 3979, 3980, 3981, 3982, 3983, 3984, 3985, 3986, 3987, 3988, 3989, 3990, 3991, 3992, 3993, 3994, 3995, 3996, 3997, 3998, 3999, 4000, 4001, 4002, 4003, 4004, 4005, 4006, 4007, 4008, 4009, 4010, 4011, 4012, 4013, 4014, 4015, 4016, 4017, 4018, 4019, 4020, 4021, 4022, 4023, 4024, 4025, 4026, 4027, 4028, 4029, 4030, 4031, 4032, 4033, 4034, 4035, 4036, 4037, 4038, 4039, 4040, 4041, 4042, 4043, 4044, 4045, 4046, 4047, 4048, 4049, 4050, 4051, 4052, 4053, 4054, 4055, 4056, 4057, 4058, 4059, 4060, 4061, 4062, 4063, 4064, 4065, 4066, 4067, 4068, 4069, 4070, 4071, 4072, 4073, 4074, 4075, 4076, 4077, 4078, 4079, 4080, 4081, 4082, 4083, 4084, 4085, 4086, 4087, 4088, 4089, 4090, 4091, 4092, 4093, 4094, 4095, 4096, 4097, 4098, 4099, 4100, 4101, 4102, 4103, 4104, 4105, 4106, 4107, 4108, 4109, 4110, 4111, 4112, 4113, 4114, 4115, 4116, 4117, 4118, 4119, 4120, 4121, 4122, 4123, 4124, 4125, 4126, 4127, 4128, 4129, 4130, 4131, 4132, 4133, 4134, 4135, 4136, 4137, 4138, 4139, 4140, 4141, 4142, 4143, 4144, 4145, 4146, 4147, 4148, 4149, 4150, 4151, 4152, 4153, 4154, 4155, 4156, 4157, 4158, 4159, 4160, 4161, 4162, 4163, 4164, 4165, 4166, 4167, 4168, 4169, 4170, 4171, 4172, 4173, 4174, 4175, 4176, 4177, 4178, 4179, 4180, 4181, 4182, 4183, 4184, 4185, 4186, 4187, 4188, 4189, 4190, 4191, 4192, 4193, 4194, 4195, 4196, 4197, 4198, 4199, 4200, 4201, 4202, 4203, 4204, 4205, 4206, 4207, 4208, 4209, 4210, 4211, 4212, 4213, 4214, 4215, 4216, 4217, 4218, 4219, 4220, 4221, 4222, 4223, 4224, 4225, 4226, 4227, 4228, 4229, 4230, 4231, 4232, 4233, 4234, 4235, 4236, 4237, 4238, 4239, 4240, 4241, 4242, 4243, 4244, 4245, 4246, 4247, 4248, 4249, 4250, 4251, 4252, 4253, 4254, 4255, 4256, 4257, 4258, 4259, 4260, 4261, 4262, 4263, 4264, 4265, 4266, 4267, 4268, 4269, 4270, 4271, 4272, 4273, 4274, 4275, 4276, 4277, 4278, 4279, 4280, 4281, 4282, 4283, 4284, 4285, 4286, 4287, 4288, 4289, 4290, 4291, 4292, 4293, 4294, 4295, 4296, 4297, 4298, 4299, 4300, 4301, 4302, 4303, 4304, 4305, 4306, 4307, 4308, 4309, 4310, 4311, 4312, 4313, 4314, 4315, 4316, 4317, 4318, 4319, 4320, 4321, 4322, 4323, 4324, 4325, 4326, 4327, 4328, 4329, 4330, 4331, 4332, 4333, 4334, 4335, 4336, 4337, 4338, 4339, 4340, 4341, 4342, 4343, 4344, 4345, 4346, 4347, 4348, 4349, 4350, 4351, 4352, 4353, 4354, 4355, 4356, 4357, 4358, 4359, 4360, 4361, 4362, 4363, 4364, 4365, 4366, 4367, 4368, 4369, 4370, 4371, 4372, 4373, 4374, 4375, 4376, 4377, 4378, 4379, 4380, 4381, 4382, 4383, 4384, 4385, 4386, 4387, 4388, 4389, 4390, 4391, 4392, 4393, 4394, 4395, 4396, 4397, 4398, 4399, 4400, 4401, 4402, 4403, 4404, 4405, 4406, 4407, 4408, 4409, 4410, 4411, 4412, 4413, 4414, 4415, 4416, 4417, 4418, 4419, 4420, 4421, 4422, 4423, 4424, 4425, 4426, 4427, 4428, 4429, 4430, 4431, 4432, 4433, 4434, 4435, 4436, 4437, 4438, 4439, 4440, 4441, 4442, 4443, 4444, 4445, 4446, 4447, 4448, 4449, 4450, 4451, 4452, 4453, 4454, 4455, 4456, 4457, 4458, 4459, 4460, 4461, 4462, 4463, 4464, 4465, 4466, 4467, 4468, 4469, 4470, 4471, 4472, 4473, 4474, 4475, 4476, 4477, 4478, 4479, 4480, 4481, 4482, 4483, 4484, 4485, 4486, 4487, 4488, 4489, 4490, 4491, 4492, 4493, 4494, 4495, 4496, 4497, 4498, 4499, 4500, 4501, 4502, 4503, 4504, 4505, 4506, 4507, 4508, 4509, 4510, 4511, 4512, 4513, 4514, 4515, 4516, 4517, 4518, 4519, 4520, 4521, 4522, 4523, 4524, 4525, 4526, 4527, 4528, 4529, 4530, 4531, 4532, 4533, 4534, 4535, 4536, 4537, 4538, 4539, 4540, 4541, 4542, 4543, 4544, 4545, 4546, 4547, 4548, 4549, 4550, 4551, 4552, 4553, 4554, 4555, 4556, 4557, 4558, 4559, 4560, 4561, 4562, 4563, 4564, 4565, 4566, 4567, 4568, 4569, 4570, 4571, 4572, 4573, 4574, 4575, 4576, 4577, 4578, 4579, 4580, 4581, 4582, 4583, 4584, 4585, 4586, 4587, 4588, 4589, 4590, 4591, 4592, 4593, 4594, 4595, 4596, 4597, 4598, 4599, 4600, 4601, 4602, 4603, 4604, 4605, 4606, 4607, 4608, 4609, 4610, 4611, 4612, 4613, 4614, 4615, 4616, 4617, 4618, 4619, 4620, 4621, 4622, 4623, 4624, 4625, 4626, 4627, 4628, 4629, 4630, 4631, 4632, 4633, 4634, 4635, 4636, 4637, 4638, 4639, 4640, 4641, 4642, 4643, 4644, 4645, 4646, 4647, 4648, 4649, 4650, 4651, 4652, 4653, 4654, 4655, 4656, 4657, 4658, 4659, 4660, 4661, 4662, 4663, 4664, 4665, 4666, 4667, 4668, 4669, 4670, 4671, 4672, 4673, 4674, 4675, 4676, 4677, 4678, 4679, 4680, 4681, 4682, 4683, 4684, 4685, 4686, 4687, 4688, 4689, 4690, 4691, 4692, 4693, 4694, 4695, 4696, 4697, 4698, 4699, 4700, 4701, 4702, 4703, 4704, 4705, 4706, 4707, 4708, 4709, 4710, 4711, 4712, 4713, 4714, 4715, 4716, 4717, 4718, 4719, 4720, 4721, 4722, 4723, 4724, 4725, 4726, 4727, 4728, 4729, 4730, 4731, 4732, 4733, 4734, 4735, 4736, 4737, 4738, 4739, 4740, 4741, 4742, 4743, 4744, 4745, 4746, 4747, 4748, 4749, 4750, 4751, 4752, 4753, 4754, 4755, 4756, 4757, 4758, 4759, 4760, 4761, 4762, 4763, 4764, 4765, 4766, 4767, 4768, 4769, 4770, 4771, 4772, 4773, 4774, 4775, 4776, 4777, 4778, 4779, 4780, 4781, 4782, 4783, 4784, 4785, 4786, 4787, 4788, 4789, 4790, 4791, 4792, 4793, 4794, 4795, 4796, 4797, 4798, 4799, 4800, 4801, 4802, 4803, 4804, 4805, 4806, 4807, 4808, 4809, 4810, 4811, 4812, 4813, 4814, 4815, 4816, 4817, 4818, 4819, 4820, 4821, 4822, 4823, 4824, 4825, 4826, 4827, 4828, 4829, 4830, 4831, 4832, 4833, 4834, 4835, 4836, 4837, 4838, 4839, 4840, 4841, 4842, 4843, 4844, 4845, 4846, 4847, 4848, 4849, 4850, 4851, 4852, 4853, 4854, 4855, 4856, 4857, 4858, 4859, 4860, 4861, 4862, 4863, 4864, 4865, 4866, 4867, 4868, 4869, 4870, 4871, 4872, 4873, 4874, 4875, 4876, 4877, 4878, 4879, 4880, 4881, 4882, 4883, 4884, 4885, 4886, 4887, 4888, 4889, 4890, 4891, 4892, 4893, 4894, 4895, 4896, 4897, 4898, 4899, 4900, 4901, 4902, 4903, 4904, 4905, 4906, 4907, 4908, 4909, 4910, 4911, 4912, 4913, 4914, 4915, 4916, 4917, 4918, 4919, 4920, 4921, 4922, 4923, 4924, 4925, 4926, 4927, 4928, 4929, 4930, 4931, 4932, 4933, 4934, 4935, 4936, 4937, 4938, 4939, 4940, 4941, 4942, 4943, 4944, 4945, 4946, 4947, 4948, 4949, 4950, 4951, 4952, 4953, 4954, 4955, 4956, 4957, 4958, 4959, 4960, 4961, 4962, 4963, 4964, 4965, 4966, 4967, 4968, 4969, 4970, 4971, 4972, 4973, 4974, 4975, 4976, 4977, 4978, 4979, 4980, 4981, 4982, 4983, 4984, 4985, 4986, 4987, 4988, 4989, 4990, 4991, 4992, 4993, 4994, 4995, 4996, 4997, 4998, 4999, 5000, 5001, 5002, 5003, 5004, 5005, 5006, 5007, 5008, 5009, 5010, 5011, 5012, 5013, 5014, 5015, 5016, 5017, 5018, 5019, 5020, 5021, 5022, 5023, 5024, 5025, 5026, 5027, 5028, 5029, 5030, 5031, 5032, 5033, 5034, 5035, 5036, 5037, 5038, 5039, 5040, 5041, 5042, 5043, 5044, 5045, 5046, 5047, 5048, 5049, 5050, 5051, 5052, 5053, 5054, 5055, 5056, 5057, 5058, 5059, 5060, 5061, 5062, 5063, 5064, 5065, 5066, 5067, 5068, 5069, 5070, 5071, 5072, 5073, 5074, 5075, 5076, 5077, 5078, 5079, 5080, 5081, 5082, 5083, 5084, 5085, 5086, 5087, 5088, 5089, 5090, 5091, 5092, 5093, 5094, 5095, 5096, 5097, 5098, 5099, 5100, 5101, 5102, 5103, 5104, 5105, 5106, 5107, 5108, 5109, 5110, 5111, 5112, 5113, 5114, 5115, 5116, 5117, 5118, 5119, 5120, 5121, 5122, 5123, 5124, 5125, 5126, 5127, 5128, 5129, 5130, 5131, 5132, 5133, 5134, 5135, 5136, 5137, 5138, 5139, 5140, 5141, 5142, 5143, 5144, 5145, 5146, 5147, 5148, 5149, 5150, 5151, 5152, 5153, 5154, 5155, 5156, 5157, 5158, 5159, 5160, 5161, 5162, 5163, 5164, 5165, 5166, 5167, 5168, 5169, 5170, 5171, 5172, 5173, 5174, 5175, 5176, 5177, 5178, 5179, 5180, 5181, 5182, 5183, 5184, 5185, 5186, 5187, 5188, 5189, 5190, 5191, 5192, 5193, 5194, 5195, 5196, 5197, 5198, 5199, 5200, 5201, 5202, 5203, 5204, 5205, 5206, 5207, 5208, 5209, 5210, 5 |  |



- permission flags
  - login class.....1672
  - user.....1672
- permissions statement.....280
  - usage guidelines.....1672
- permissions, CLI, displaying.....876
- persistent configuration changes
  - compared to transient changes.....6091
  - overview.....6090
- persistent-learning statement.....4822
- pfc input statement.....5889
- PFE
  - next hops, displaying.....6051
  - routing table, displaying.....6056
  - terse information, displaying.....6062
  - version, displaying.....6064
- pfe (system logging facility).....6223
- physical interfaces
  - descriptive text.....2562, 2878
  - loopback mode.....2535
- physical interfaces, descriptive text.....2562, 2878
- pic (Port) statement.....1400
- pic statement.....2612
- PICs
  - installed, displaying list.....668
  - status
    - displaying for a specific PIC.....818
    - displaying FPCs and PICs.....641
- PIM
  - anycast RP.....4360, 4411
  - assert timeout.....4361, 4414
    - configuring.....4284
  - background.....4209
  - BFD.....4290, 4363, 4391, 4393, 4394, 4422
  - bootstrap messages import and export.....4276
  - bootstrap routers.....4217, 4276
  - bootstrap routers, displaying.....4587
  - configuring.....4212
  - designated router.....4215
  - embedded RP.....4371
  - enabling.....4398
  - filters See multicast filters
  - groups
    - general information, displaying.....4592
    - usage information, displaying.....4584
  - hello interval.....4246
  - hold-time period.....4379
  - incoming join filter policy, applying.....4281
  - interfaces
    - displaying.....4589
  - join load balancing
    - configuring.....4256
  - join states, clearing.....4505
  - join suppression
    - configuring.....4259
  - join-prune-timeout.....4385
  - maximum RPs.....4390
  - neighbors, displaying.....4601
  - network components.....4211
  - outgoing join filter policy, applying.....4280
  - overview.....4209
  - PIM-to-IGMP message translation information,
    - displaying.....4565
  - PIM-to-MLD message translation information,
    - displaying.....4567
  - policy, routing.....4381
  - prune states, clearing.....4505
  - register
    - clearing.....4507
  - rendezvous point tree.....4221
  - routing tables.....4407
  - RPF, displaying source state.....4612
  - RPs.....4214, 4220, 4264, 4408
    - anycast.....4360
    - anycast RP.....4217
    - displaying.....4605
    - embedded.....4371
    - mapping options.....4215
    - maximum.....4390
    - source registration.....4222
    - SPT cutover control.....4229
  - sparse mode.....4213, 4255
  - SSM.....4238, 4239, 4347
  - statistics
    - clearing.....4509
    - displaying.....4615
    - version.....4246, 4255, 4423
- pim statement.....4398
  - usage guidelines.....4212
- PIM, graceful restart.....2269
- PIM-RP
  - SPT
    - configuring threshold cutover
      - policy.....4287
- pim-to-igmp-proxy statement.....2910
- ping command.....354, 6392, 6415, 6503
- Ping MIB.....6132

|                                                  |                  |                                                 |                  |
|--------------------------------------------------|------------------|-------------------------------------------------|------------------|
| plain-text passwords.....                        | 159, 1702        | port-mode statement.....                        | 2112, 2613       |
| for user accounts.....                           | 1682             | port-range statement.....                       | 5253             |
| root password.....                               | 158, 1687, 1702  | port-receive-state-machine (tracing flag)       |                  |
| plain-text-password option.....                  | 158, 1687, 1702  | STP.....                                        | 2138             |
| point-to-point statement.....                    | 3740             | port-role-select-state-machine (tracing flag)   |                  |
| poison reverse technique.....                    | 4106             | STP.....                                        | 2138             |
| policer statement.....                           | 4797             | port-role-transit-state-machine (tracing flag)  |                  |
| policer, single-rate two-color                   |                  | STP.....                                        | 2138             |
| example.....                                     | 4351             | port-state-transit-state-machine (tracing flag) |                  |
| policy See routing policies                      |                  | STP.....                                        | 2138             |
| policy (tracing flag).....                       | 2950             | port-transmit-state-machine (tracing flag)      |                  |
| policy statement                                 |                  | STP.....                                        | 2138             |
| aggregate routes.....                            | 2911             | ports                                           |                  |
| flow map.....                                    | 2912             | RADIUS servers.....                             | 1699             |
| generated routes.....                            | 2911             | ports statement.....                            | 281              |
| SSM map.....                                     | 4483             | postinheritance, defined.....                   | 6085             |
| policy, import                                   |                  | Power Entry Module.....                         | 604              |
| BSR.....                                         | 4380             | Power Supply Unit MIB.....                      | 6139             |
| policy, routing                                  |                  | powering off routing software                   |                  |
| BGP.....                                         | 3512, 3525       | requesting a system power off.....              | 385              |
| forwarding table.....                            | 2880             | ppm statement.....                              | 2916, 2917       |
| IS-IS.....                                       | 3702             | ppmd (tracing flag)                             |                  |
| MSDP.....                                        | 4467, 4469       | STP.....                                        | 2138             |
| OSPF.....                                        | 3981, 3989       | preempt statement.....                          | 2330             |
| network summaries.....                           | 3938, 3946       | preference statement                            |                  |
| network summaries, overview.....                 | 3937             | aggregate routes.....                           | 2918             |
| PIM.....                                         | 4381             | BGP.....                                        | 3567             |
| PIM join filter.....                             | 4280, 4281       | usage guidelines.....                           | 3327             |
| precedence.....                                  | 3309             | generated routes.....                           | 2918             |
| RIP.....                                         | 4179, 4183       | IS-IS.....                                      | 3741             |
| policy-options statement.....                    | 2913, 5251, 5890 | OSPF.....                                       | 4010             |
| policy-statement statement.....                  | 2914             | usage guidelines.....                           | 3869             |
| poll-interval statement                          |                  | RIP.....                                        | 4188             |
| usage guidelines.....                            | 3841             | static routes.....                              | 2918             |
| polling-interval statement.....                  | 6269             | preferences                                     |                  |
| Port security                                    |                  | active routes.....                              | 3332             |
| monitoring.....                                  | 4847             | aggregate routes.....                           | 2918             |
| port statement                                   |                  | generated routes.....                           | 3326             |
| RADIUS.....                                      | 1784             | default.....                                    | 3326             |
| SNMPv3.....                                      | 6315             | IS-IS.....                                      | 3326, 3703, 3741 |
| TACACS+.....                                     | 280              | modifying                                       |                  |
| usage guidelines.....                            | 1711             | with configuration statements.....              | 3326             |
| usage guidelines.....                            | 1699             | OSPF.....                                       | 3982, 4010       |
| port-error-disable statement.....                | 4823             | RIP.....                                        | 3326             |
| port-information-state-machine (tracing          |                  | static routes.....                              | 2918, 3326       |
| flag).....                                       | 2138             | preferred statement.....                        | 2113             |
| port-migration-state-machine (tracing flag)..... | 2138             | usage guidelines.....                           | 2041             |
| port-mode (fibre channel interfaces)             |                  |                                                 |                  |
| statement.....                                   | 5252             |                                                 |                  |

- prefix limit
  - IS-IS.....3742
  - OSPF.....3858, 4011
- prefix list
  - adding to SNMP community.....6196
- prefix list statement
  - firewall filter match condition.....3441
- prefix statement.....2919
- prefix-based
  - usage guidelines.....3313
- prefix-export-limit statement
  - IS-IS.....3742
  - OSPF.....4011
  - usage guidelines.....3858
- prefix-list statement
  - PIM RPF selection.....4401
- prefix-policy statement
  - BGP
    - usage guidelines.....3380
- prefixes
  - specifying in configuration statements.....49
- primary statement
  - address on interface.....2114
- priorities
  - system logging, including in log message
    - for single-chassis system.....6213
- priority
  - PIM RPs.....4404
- priority statement.....2331, 5254
  - bootstrap.....4402
  - class of service.....5891
  - IS-IS.....3743
  - OSPF.....4012
  - usage guidelines.....3812
  - PIM.....4403
  - usage guidelines.....4253
  - STP.....2115
  - usage guidelines.....2296, 4276
- priority-cost statement.....2332
  - usage guidelines.....2302
- priority-flow-control statement.....5255, 5892
- priority-hold-time statement.....2333
  - usage guidelines.....2301, 2302
- privacy-3des statement.....6316
- privacy-aes128 statement.....6317
- privacy-des statement.....6318
- privacy-none statement.....6318
- privacy-password statement.....6319
- private statement
  - usage guidelines.....3261
- process (tracing flag)
  - class of service.....5909
- process ID, displaying.....315, 2627
- process information, system, monitoring.....315, 2627
- process owner, displaying.....315, 2627
- process start time, displaying.....316, 2628
- process state, displaying.....315, 2628
- processes
  - displaying information.....983
  - restarting.....438, 5288
- processes, software
  - chassis process.....68, 1662
  - forwarding process.....68, 1662
  - interface process.....68, 1662
  - management process.....68, 1662
  - routing protocol process.....68, 1662
- profile statement.....1785
- promiscuous-mode statement
  - IGMP
    - interface.....4438
    - usage guidelines.....4302
- propagation, suppressing.....3456
- propagation-delay statement.....4405
  - usage guidelines.....4259
- properties, system, monitoring.....316, 2628
- Protocol Independent Multicast *See* PIM *See* PIM
- protocol statement.....5894
- protocol-specific tracing operations.....6081
- protocol-version statement.....1799, 6232
  - usage guidelines.....1710, 1711, 6179, 6180
- protocols
  - distance vector *See* RIP
  - group membership.....4229
  - originating, displaying.....2956
  - OSPF, monitoring.....4031
  - redirect messages.....148
  - RIP *See* RIP
  - RIP, monitoring.....4197
  - routing protocols, monitoring.....3583
- protocols (FIP) statement.....5257
- protocols statement.....1786, 2116, 2920
- proxy statement.....5258
- proxy-arp statement.....2129
- prune (tracing flag)
  - PIM.....4420
- prune states, clearing PIM.....4505
- psn (tracing flag).....3747

PSNP IS-IS.....3628

## Q

### QFabric system

#### backup and recovery

performing.....1460, 1467

### QFabric system control plane Virtual Chassis

configuring.....1263

### QFX3000-M QFabric system copper-based control

#### plane

configuring.....1309

qualified-next-hop statement.....2922

usage guidelines.....2840

### query-interval statement

IGMP.....4438

usage guidelines.....4299

### query-last-member-interval statement

IGMP.....4439

usage guidelines.....4299

### query-response-interval statement

IGMP.....4440

usage guidelines.....4303

queue-num statement.....5895

### queues

forwarding class mapping, displaying.....5936

quit command.....76

usage guidelines.....82

## R

RADIUS accounting.....1697

RADIUS authentication.....1681, 1699

security configuration example.....1730

TACACS+ .....1701

RADIUS authorization See RADIUS authentication

radius statement.....1801

accounting.....282

### RADIUS templates

security configuration example.....1745

radius-options statement .....1802

radius-server statement.....1803

usage guidelines.....1699

random early detection.....5933

### rapid-runs statement

OSPF.....4018

rate-limit statement.....1804, 6233

### reachability

verifying for a RIP network.....4115

read-view statement.....6320

readvertise statement.....2924

### real-time monitoring

interfaces.....2635

IP multicast paths.....4512

traffic.....6405

realm statement.....4013

### rebooting router software

#### pending reboots

clearing.....332

displaying.....1010

### receive statement

#### BGP

usage guidelines.....3380

RIP.....4189

usage guidelines.....4150

recommendation-tlv statement.....5258, 5896

### recovery software installation

procedures.....111, 1581, 6528

red alarm conditions.....147

RED drop profiles, displaying.....5933

### redirect messages

disabling.....148

redirected routes.....3326

### redistributing routes

OSPF into IS-IS.....3651

redundant-sources statement.....2925

reference-bandwidth statement.....4014

IS-IS.....3744

usage guidelines.....3863

reflective-relay statement.....2130, 2614

### refresh statement

commit scripts.....283, 6248

op scripts.....6249

### refresh-from statement

commit scripts.....283, 6250

op scripts.....6251

regex-parse (tracing flag).....2950

register (tracing flag).....4420

regular expression operators.....1686

system logging.....6228

### regular expressions

AS paths, displaying matching routes.....2986

#### IP multicast scope

clearing.....4502

#### IP multicast sessions

clearing.....4503

displaying.....4581

reinstalling Junos OS.....111, 1581, 6528

reject option to static statement.....2943

relative option.....1599

- relay agents
  - DHCP and BOOTP.....4907
- remote
  - template account.....1701
- Remote Monitoring.....6455
- remote-debug-permission statement.....1401, 1805
- remote-engine statement.....6321
- removable media
  - reinstalling Junos OS, using.....111, 1581, 6528
- remove-private statement.....3568
  - usage guidelines.....3343
- removing
  - files.....345, 1640
  - software packages.....404
- renaming files.....348, 1643
- rendezvous points *See* RPs *See* PIM and RP
- replace option.....1598
- report (tracing flag)
  - IGMP.....4447
- request chassis beacon command.....358
- request chassis cb command.....360
- request chassis device-mode command.....1441
- request chassis fabric fpc command.....1443
- request chassis fpc command.....363
- request chassis routing-engine master
  - command.....367
- request command
  - usage guidelines.....76
- request component login command.....1444, 1836
- request fabric administration director-group
  - change-master command.....1446
- request fabric administration remove
  - command.....1447
- request fabric administration system mac-pool add
  - command.....1449
- request fabric administration system mac-pool
  - delete command.....1450
- request fibre-channel proxy load-rebalance fabric
  - command.....5286
- request message command.....372
- request snmp spoof-trap command.....6424
- request system configuration rescue delete
  - command.....373, 1597, 1604, 1609, 1647
- request system configuration rescue save
  - command.....374, 1596, 1604, 1648
- request system halt command.....375, 1451
- request system license add command.....93, 381
- request system license add terminal
  - command.....93
- request system license delete command.....382
- request system license save command.....95, 383
- request system logout command.....384
- request system power-off command.....385
- request system software add
  - command.....394, 408, 410, 1461, 2393
- request system software delete command.....404
- request system software format-qfabric-backup
  - command.....1460
- request system software rollback command.....416
- request system software system-backup
  - command.....1467
- request system software validate command.....420
- request system storage cleanup command.....423
- request system zeroize command.....433
- request-type statement.....6322
- rescue configuration
  - deleting.....373, 1647
  - displaying.....942, 1652
  - saving.....374, 1648
- reset-tracking-bit statement.....4406
  - usage guidelines.....4259
- resolution statement.....2926
- resolution-ribs statement.....2927
- resolve statement.....2928
- Resource Reservation Protocol *See* RSVP
- resources statement
  - Node group assignment.....1402
- restart (tracing flag)
  - class of service.....5909
- restart command.....438, 5288
  - usage guidelines.....76
- restart-duration statement.....2316
  - usage guidelines
    - ES-IS.....2287
    - global.....1375, 2282, 2285
    - IS-IS.....2287
    - OSPF .....1377, 2283, 2288
    - OSPFv3.....1377, 2283, 2288
    - PIM.....2289
- restart-signaling (tracing flag)
  - OSPFv2.....4023
- restart-time statement.....2317, 3570
  - usage guidelines
    - BGP.....1376, 2282, 2286
    - RIP.....2289
    - RIPng.....2289
- restarting
  - software processes.....438, 5288

|                                           |            |                                            |            |
|-------------------------------------------|------------|--------------------------------------------|------------|
| retain statement.....                     | 2929       | general statistics                         |            |
| retransmit-interval statement.....        | 4015       | clearing.....                              | 4198       |
| usage guidelines.....                     | 3876       | displaying.....                            | 4200       |
| retry statement.....                      | 284, 1806  | graceful restart.....                      | 2269       |
| usage guidelines.....                     | 1699       | hold-down timer.....                       | 4182       |
| retry-count statement.....                | 6323       | maximum hop count.....                     | 4106       |
| retry-options statement.....              | 285        | metrics.....                               | 4185, 4186 |
| usage guidelines.....                     | 1716       | neighbors.....                             | 4187       |
| reverse path forwarding See RPF           |            | displaying.....                            | 4202       |
| reverse-oif-mapping statement.....        | 2930       | overview.....                              | 4103, 4109 |
| reverting to earlier software.....        | 416        | packets.....                               | 4105       |
| revision-level statement                  |            | path cost metrics See path cost metrics    |            |
| STP.....                                  | 2130       | poison reverse technique.....              | 4106       |
| rewrite rules, displaying                 |            | policy, routing.....                       | 4179, 4183 |
| information.....                          | 5982       | preferences.....                           | 3326, 4188 |
| mapping of code point value to loss       |            | reserved fields.....                       | 4178       |
| priority.....                             | 5950       | rib-group messages.....                    | 4155       |
| mapping of code point value to queue      |            | rib-group statement                        |            |
| number.....                               | 5950       | usage guidelines.....                      | 4155       |
| table identifiers.....                    | 5952       | route timeout.....                         | 4191       |
| rewrite-rules statement.....              | 5897       | routing policy (configuration editor)..... | 4110       |
| rewrite-value statement                   |            | split horizon technique.....               | 4106       |
| class of service.....                     | 5898       | statistics                                 |            |
| RFC 1583                                  |            | clearing.....                              | 4199       |
| disabling.....                            | 3836       | displaying.....                            | 4204       |
| OSPFv2.....                               | 3836       | tracing operations                         |            |
| RFC 5185.....                             | 3832       | authentication.....                        | 4166       |
| rib statement                             |            | description.....                           | 4166, 4167 |
| route resolution.....                     | 2934       | error PDUs.....                            | 4166       |
| routing tables.....                       | 2932       | hold-down processing.....                  | 4166       |
| rib-group statement.....                  | 2935       | NSR synchronization.....                   | 4166       |
| IS-IS.....                                | 3745       | policy processing.....                     | 4166       |
| MSDP.....                                 | 4476       | protocol task processing.....              | 4166       |
| OSPF.....                                 | 4016       | protocol timer processing.....             | 4166       |
| PIM.....                                  | 4407       | request PDUs.....                          | 4166       |
| RIP.....                                  | 4190       | route expiration processing.....           | 4166       |
| usage guidelines.....                     | 4155       | route information.....                     | 4166, 4167 |
| rib-groups statement.....                 | 2936       | state transitions.....                     | 4166       |
| RIP                                       |            | trigger updates.....                       | 4166       |
| authentication.....                       | 4116, 4173 | update PDUs.....                           | 4166       |
| authentication (RIPv2 only).....          | 4116       | traffic control with metrics See path cost |            |
| basic network (configuration editor)..... | 4110       | metrics                                    |            |
| BFD.....                                  | 4122, 4123 | traffic control with metrics,              |            |
| disabling address checks.....             | 4172       | configuring.....                           | 4143, 4144 |
| distance vector protocol.....             | 4103       | UDP, use of.....                           | 4104       |
| efficiency techniques.....                | 4106       | unidirectional limitations.....            | 4107       |
| enabling.....                             | 4190       | update interval.....                       | 4196       |
|                                           |            | update messages.....                       | 4150, 4184 |
|                                           |            | verifying host reachability .....          | 4115       |

- 
- RIP (Routing Information Protocol)
    - authentication (RIPv2 only), configuring.....4121
  - RIP routing information.....4197
  - rip statement.....4190
  - RIPng
    - graceful restart.....2269
  - rising-event-index statement.....6324
  - rising-threshold statement
    - health monitor.....1807, 6325
    - RMON.....6326
  - RJ-45-to-DB-9 serial port
    - adapter.....1159, 1717, 6478, 6525
  - rlogin service, configuring.....1809
  - rmon
    - history.....6459
  - RMON alarms.....6121, 6122
  - RMON alarms and events, displaying.....6455
  - RMON events.....6121
  - RMON Events and Alarms MIB.....6132
  - RMON history control.....6122
  - rmon statement.....6327
  - robust-count statement.....4455
    - IGMP.....4441
    - usage guidelines.....4304
  - rollback
    - displaying.....1015, 1654
    - requesting.....416
  - rollback command.....449, 1607
    - usage guidelines.....82
  - rollover cable, connecting the console
    - port.....1159, 1718, 6478, 6525
  - root
    - password .....189
  - root password.....158, 1687, 1702
  - root password recovery.....1159, 1717, 6478, 6524
  - root-authentication statement.....286, 293
    - usage guidelines.....158, 1687, 1702
  - root-login statement.....1808
    - usage guidelines.....1710, 6179
  - route
    - static statement.....2943
  - route (tracing flag)
    - MSDP.....4479
    - routing.....2950
  - route advertisements
    - stub areas and NSSAs, to control.....3821
  - route authentication
    - peering sessions.....3429
  - route injection.....3306, 3923, 3927, 3929, 3933
  - route limit, configuring
    - paths.....2902
    - prefix.....2904
  - route preference
    - external routes.....3863
    - internal routes.....3863
  - route prefixes.....49
  - route recording.....2937
  - route redistribution.....3306, 3923, 3927, 3929, 3933
  - route reflectors *See* BGP route reflectors
  - BGP.....3407
  - route selection
    - OSPF.....3861
    - preference.....3863
    - static routes, defining.....2840
  - route statement
    - aggregate statement.....2858
    - generate statement.....2888
    - VRRP.....2334
    - usage guidelines.....2301
  - route summarization
    - configuring.....3852
  - route, displaying
    - next-hop.....3070
  - route-flap damping.....3456
    - parameters.....3457
  - route-record statement.....2937
  - route-socket (tracing flag)
    - class of service.....5909
  - route-timeout statement
    - RIP.....4191
    - usage guidelines.....4160
  - router advertisements
    - IPv6
      - displaying.....2971
  - router functionality.....3803
  - router identifier.....2938
    - configuring.....3810
  - router link advertisements.....3807
  - router-id statement.....2938
  - routers
    - DNS name servers, configuring.....146
    - domain names.....146
    - domains to be searched.....147, 257
    - login classes.....1683
    - names
      - configuring.....184
      - mapping to IP addresses.....168, 184
    - NTP.....180



|                                                |                                    |
|------------------------------------------------|------------------------------------|
| physical system location.....                  | 156                                |
| ports                                          |                                    |
| RADIUS servers.....                            | 1699                               |
| redirect .....                                 | 148                                |
| root login, controlling.....                   | 1710, 6179                         |
| source addresses.....                          | 151, 252                           |
| time zone setting.....                         | 172                                |
| user accounts.....                             | 1692                               |
| routes, displaying                             |                                    |
| active.....                                    | 2974                               |
| active path.....                               | 2979                               |
| all.....                                       | 2984                               |
| AS paths                                       |                                    |
| distribution of.....                           | 2960                               |
| domain information.....                        | 2964                               |
| regular expressions, matching.....             | 2986                               |
| summary of.....                                | 2967                               |
| best.....                                      | 2988                               |
| brief information.....                         | 2991                               |
| community ASN.....                             | 2993                               |
| community name.....                            | 2995                               |
| damping, BGP.....                              | 2997, 3617                         |
| detailed information.....                      | 3002                               |
| extensive information.....                     | 3022                               |
| flow validation.....                           | 3037                               |
| in a prefix range.....                         | 3091                               |
| in a specific routing table.....               | 3124                               |
| in the forwarding table.....                   | 3039                               |
| inactive path.....                             | 3052                               |
| inactive prefix.....                           | 3055                               |
| instances.....                                 | 1855, 3057                         |
| learned from a specific address.....           | 3114                               |
| learned from a specific protocol.....          | 3079                               |
| learned from snooping.....                     | 3106                               |
| LSP.....                                       | 3066                               |
| martian.....                                   | 3068                               |
| matching the specified address.....            | 3017                               |
| MPLS labels.....                               | 3064                               |
| next-hop resolution.....                       | 3103                               |
| not associated with a community.....           | 3076                               |
| policy-based route export.....                 | 3019                               |
| received through a neighbor.....               | 3095                               |
| summary statistics.....                        | 3120                               |
| terse information.....                         | 3135                               |
| to specified network host.....                 | 1151                               |
| routing                                        |                                    |
| in one AS with RIP.....                        | 4109                               |
| RIP See RIP                                    |                                    |
| Routing Engine                                 |                                    |
| software component.....                        | 68, 1661                           |
| Routing Engines                                |                                    |
| connections, displaying.....                   | 944                                |
| environmental information, displaying.....     | 613                                |
| operation of, controlling.....                 | 367                                |
| status, displaying.....                        | 832                                |
| Routing Information Protocol See RIP           |                                    |
| routing policies                               |                                    |
| applying.....                                  | 3305                               |
| BGP export, for CLNS.....                      | 3684                               |
| configuration                                  |                                    |
| tasks.....                                     | 3306, 3457, 3923, 3927, 3929, 3933 |
| export statement.....                          | 3305                               |
| import statement.....                          | 3305                               |
| injecting routes from one protocol into        |                                    |
| another.....                                   | 3306, 3923                         |
| OSPF import policy.....                        | 3929, 3933                         |
| redistributing static routes into OSPF.....    | 3927                               |
| reducing update messages with flap             |                                    |
| damping.....                                   | 3456                               |
| RIP routing policy (configuration editor)..... | 4110                               |
| route                                          |                                    |
| redistribution.....                            | 3306, 3923, 3927, 3929, 3933       |
| route-flap damping.....                        | 3456                               |
| routing policy                                 |                                    |
| actions.....                                   | 3923                               |
| match conditions.....                          | 3922                               |
| terms.....                                     | 3922                               |
| routing protocol software process.....         | 68, 1662                           |
| routing solutions                              |                                    |
| BGP confederations, for scaling                |                                    |
| problems.....                                  | 3423                               |
| BGP route reflectors, for scaling              |                                    |
| problems.....                                  | 3407                               |
| controlling RIP traffic with the incoming      |                                    |
| metric.....                                    | 4143                               |
| controlling RIP traffic with the outgoing      |                                    |
| metric.....                                    | 4144                               |
| NSSAs, to control route advertisement.....     | 3821                               |
| poison reverse, for traffic reduction.....     | 4106                               |
| reducing update messages with flap             |                                    |
| damping.....                                   | 3456                               |
| split horizon, for traffic reduction.....      | 4106                               |
| stub areas, to control route                   |                                    |
| advertisement.....                             | 3821                               |
| routing table                                  |                                    |
| sample distance-vector routing.....            | 4104                               |
| updates, limitations in RIP.....               | 4107                               |



- 
- routing tables
    - BGP, delays in exchanging routes.....3312
    - creating.....2932
    - group.....2881, 2891, 2935, 2936, 4016
    - import policy.....2890
    - MSDP.....4476
    - nonactive routes, exchanging with
      - BGP.....3310, 3328, 3487
      - PIM.....4407
  - routing-engine statement
    - reboot on disk failure.....2614
  - routing-instance statement
    - SNMP.....6328
    - usage guidelines.....1699
  - routing-options statement.....2939
    - QFabric system.....1402
  - RP
    - anycast.....4360
    - embedded.....4371
  - rp (tracing flag).....4420
  - rp statement.....4408
  - rp-register-policy statement.....4410
    - usage guidelines.....4282
  - rp-set statement.....4411
    - usage guidelines.....4267, 4338
  - rp-d process.....68, 1662
  - RPF
    - calculations, displaying.....4575
    - PIM source state, displaying.....4612
  - RPF check, multicast
    - RPF policy.....2931
  - rpf-check-policy statement.....2931
  - rpf-selection statement
    - PIM.....4412
  - RPs
    - displaying.....4605
    - maximum.....4390
  - RPT.....4220
  - rstp statement.....2131
  - RSVP
    - graceful restart.....2269
    - preferences.....3326
  - RSVP LSP metrics
    - ignoring.....3908
  - run command
    - usage guidelines.....82
  - S**
  - sample-rate statement.....6270
  - sample-type statement.....6329
  - save command.....450
    - usage guidelines.....82, 1605, 1608
  - saved-core-context statement.....287
    - usage guidelines.....175
  - saved-core-files statement.....287
    - usage guidelines.....175
  - saving licenses (CLI).....95
  - SCB
    - firmware version, displaying.....631
  - scc-master option to host statement.....6375
  - scheduler-map statement.....5901
  - scheduler-maps statement.....5902
  - schedulers statement.....5903
  - scope statement.....2940
  - scope-policy statement.....2941
  - scoping, multicast.....2940
    - with scope policy.....2941
  - SCP.....1707
  - scripts statement.....6252
  - secondary statement
    - usage guidelines.....3831, 3832
  - secret statement.....4932
    - authentication.....288
      - usage guidelines, RADIUS.....1699
      - usage guidelines, TACACS+.....1711
  - secure copy See SCP
  - secure-access-port statement.....4826
  - security
    - MD5 authentication for RIPv2.....4121
    - password authentication for RIPv2.....4121
  - security-level statement
    - for access privileges.....6330
    - for SNMP notifications.....6331
  - security-model statement
    - for access privileges.....6332
    - for groups.....6333
    - for SNMP notifications.....6334
  - security-name statement
    - for community string.....6335
    - for security group.....6336
    - for SNMP notifications.....6337
  - security-to-group statement.....6338
  - send statement
    - BGP
      - usage guidelines.....3380
    - RIP.....4192
      - usage guidelines.....4150

|                                                |                        |  |
|------------------------------------------------|------------------------|--|
| serial number                                  |                        |  |
| routing platform.....                          | 316, 2628              |  |
| serial numbers, displaying.....                | 667                    |  |
| server mode, usage guidelines.....             | 155                    |  |
| server statement                               |                        |  |
| DHCP and BOOTP service .....                   | 4945                   |  |
| NTP.....                                       | 289                    |  |
| RADIUS accounting.....                         | 290                    |  |
| TACPLUS+.....                                  | 290                    |  |
| usage guidelines.....                          | 153                    |  |
| service-name statement.....                    | 1821                   |  |
| services statement.....                        | 1809                   |  |
| Session and Resource Control.....              | 1017                   |  |
| session-mode statement                         |                        |  |
| BGP.....                                       | 3571                   |  |
| set chassis display message command.....       | 1469                   |  |
| set command                                    |                        |  |
| usage guidelines.....                          | 82                     |  |
| set date ntp.....                              | 177                    |  |
| set option.....                                | 1599                   |  |
| set system root-authentication command.....    | 189                    |  |
| set system root-authentication                 |                        |  |
| encrypted-password command.....                | 189                    |  |
| setting date and time (CLI).....               | 177                    |  |
| setting root password.....                     | 189                    |  |
| sFlow.....                                     | 6105, 6165, 6189       |  |
| sflow statement.....                           | 6271                   |  |
| sflow technology                               |                        |  |
| tracing operations.....                        | 6273                   |  |
| SFM                                            |                        |  |
| firmware version, displaying.....              | 631                    |  |
| sha-256 checksum, calculating.....             | 341, 1636              |  |
| shaping-rate statement.....                    | 5904                   |  |
| shared trees.....                              | 4220                   |  |
| shared-buffer statement.....                   | 5906                   |  |
| SHA-1 checksum, calculating.....               | 340, 1635              |  |
| shortcuts statement                            |                        |  |
| OSPF.....                                      | 4017                   |  |
| Shortest Path First See SPF algorithm          |                        |  |
| shortest-path trees.....                       | 4225                   |  |
| See also SPT                                   |                        |  |
| show (ospf   ospf3) backup coverage            |                        |  |
| command.....                                   | 4043                   |  |
| show (ospf   ospf3) backup neighbor.....       | 4046                   |  |
| show (ospf   ospf3) interface command.....     | 4069                   |  |
| show (ospf   ospf3) io-statistics command..... | 4075                   |  |
| show (ospf   ospf3) log command.....           | 4077                   |  |
| show (ospf   ospf3) neighbor command.....      | 4080                   |  |
| show (ospf   ospf3) overview                   |                        |  |
| command.....                                   | 2363, 4086             |  |
| show (ospf   ospf3) route command.....         | 4091                   |  |
| show (ospf   ospf3) statistics command.....    | 4097                   |  |
| show (tracing flag)                            |                        |  |
| class of service.....                          | 5909                   |  |
| show administrator inventory infrastructure    |                        |  |
| command.....                                   | 1553                   |  |
| show analyzer command.....                     | 4948                   |  |
| show arp inspection statistics command.....    | 4861                   |  |
| show as-path command.....                      | 2960                   |  |
| show as-path domain command.....               | 2964                   |  |
| show as-path summary command.....              | 2967                   |  |
| show bgp group command.....                    | 3589                   |  |
| show bgp neighbor                              |                        |  |
| command.....                                   | 2346, 3426, 3583, 3596 |  |
| explanation.....                               | 3427                   |  |
| show bgp summary command.....                  | 3428, 3583, 3610       |  |
| explanation.....                               | 3428                   |  |
| show chassis alarms command.....               | 452                    |  |
| show chassis device-mode command.....          | 1475                   |  |
| show chassis environment cb command.....       | 561                    |  |
| show chassis environment command.....          | 508                    |  |
| show chassis environment fpc command.....      | 579                    |  |
| show chassis environment pem command.....      | 604                    |  |
| show chassis environment routing-engine        |                        |  |
| command.....                                   | 613                    |  |
| show chassis ethernet-switch                   |                        |  |
| command.....                                   | 466, 491, 1478, 1495   |  |
| show chassis fabric connectivity command.....  | 1520                   |  |
| show chassis fabric device command.....        | 1527                   |  |
| show chassis fan command.....                  | 618                    |  |
| show chassis firmware command.....             | 631                    |  |
| show chassis fpc command.....                  | 641                    |  |
| show chassis hardware command.....             | 667                    |  |
| show chassis location command.....             | 808                    |  |
| show chassis mac-addresses command.....        | 812                    |  |
| show chassis pic command.....                  | 818                    |  |
| show chassis routing-engine command.....       | 831                    |  |
| show chassis temperature-thresholds            |                        |  |
| command.....                                   | 852                    |  |
| show class-of-service classifier               |                        |  |
| command.....                                   | 5915, 5927             |  |
| show class-of-service code-point-aliases       |                        |  |
| command.....                                   | 5920, 5929             |  |
| show class-of-service command.....             | 5923                   |  |
| show class-of-service congestion-notification  |                        |  |
| command.....                                   | 5931                   |  |

|                                                     |            |
|-----------------------------------------------------|------------|
| show class-of-service drop-profile                  |            |
| command.....                                        | 5933       |
| show class-of-service forwarding-class              |            |
| command.....                                        | 5916, 5936 |
| show class-of-service forwarding-class-set          |            |
| command.....                                        | 5938       |
| show class-of-service forwarding-table classifier   |            |
| command.....                                        | 5944       |
| show class-of-service forwarding-table classifier   |            |
| mapping command.....                                | 5946       |
| show class-of-service forwarding-table              |            |
| command.....                                        | 5940       |
| show class-of-service forwarding-table drop-profile |            |
| command.....                                        | 5948       |
| show class-of-service forwarding-table rewrite-rule |            |
| command.....                                        | 5950       |
| show class-of-service forwarding-table rewrite-rule |            |
| mapping command.....                                | 5952       |
| show class-of-service forwarding-table              |            |
| scheduler-map command.....                          | 5953       |
| show class-of-service interface command.....        | 5955       |
| show class-of-service multi-destination             |            |
| command.....                                        | 5981       |
| show class-of-service rewrite-rule                  |            |
| command.....                                        | 5982       |
| show class-of-service scheduler-map                 |            |
| command.....                                        | 5984       |
| show class-of-service shared-buffer                 |            |
| command.....                                        | 5986       |
| show class-of-service traffic-control-profile       |            |
| command.....                                        | 5988       |
| show cli authorization command.....                 | 876        |
| show cli command.....                               | 874        |
| show cli directory command.....                     | 880        |
| show cli history command.....                       | 881        |
| show command                                        |            |
| usage guidelines.....                               | 82         |
| show dcbx command.....                              | 5299, 5992 |
| show dcbx neighbors command.....                    | 5300, 5993 |
| show dhcp snooping binding command.....             | 4862       |
| show ethernet-switching interfaces                  |            |
| command.....                                        | 1838, 2191 |
| show ethernet-switching mac-learning-log            |            |
| command.....                                        | 2202       |
| show ethernet-switching statistics aging            |            |
| command.....                                        | 2205       |
| show ethernet-switching statistics mac-learning     |            |
| command.....                                        | 2207       |
| show ethernet-switching table command.....          | 2211       |
| show fabric administration inventory                |            |
| command.....                                        | 1543       |
| show fabric administration inventory director-group |            |
| status command.....                                 | 1548       |
| show fabric administration inventory                |            |
| interconnect-devices command.....                   | 1556       |
| show fabric administration inventory node-devices   |            |
| command.....                                        | 1558       |
| show fabric administration inventory node-groups    |            |
| command.....                                        | 1560       |
| show fabric administration system mac-pool          |            |
| command.....                                        | 1562       |
| show fabric inventory command.....                  | 1563       |
| show fabric session-host command.....               | 1566       |
| show fibre-channel fabric command.....              | 5322       |
| show fibre-channel fc2 sessions command.....        | 5324       |
| show fibre-channel fc2 statistics command.....      | 5326       |
| show fibre-channel fip command.....                 | 5328       |
| show fibre-channel fip enode command.....           | 5333       |
| show fibre-channel fip fabric command.....          | 5337       |
| show fibre-channel fip fcf command.....             | 5340       |
| show fibre-channel fip interface command.....       | 5343       |
| show fibre-channel fip statistics command.....      | 5346       |
| show fibre-channel flogi fport command.....         | 5350       |
| show fibre-channel flogi nport command.....         | 5352       |
| show fibre-channel flogi statistics                 |            |
| command.....                                        | 5354       |
| show fibre-channel interfaces command.....          | 5357       |
| show fibre-channel next-hops command.....           | 5360       |
| show fibre-channel proxy fabric-state               |            |
| command.....                                        | 5364       |
| show fibre-channel proxy login-table                |            |
| command.....                                        | 5368       |
| show fibre-channel proxy np-port command.....       | 5371       |
| show fibre-channel proxy statistics                 |            |
| command.....                                        | 5374       |
| show fibre-channel routes command.....              | 5362       |
| show fip snooping command.....                      | 5377       |
| show fip snooping enode command.....                | 5381       |
| show fip snooping fcf command.....                  | 5385       |
| show fip snooping interface command.....            | 5388       |
| show fip snooping statistics command.....           | 5391       |
| show fip snooping vlan command.....                 | 5394       |
| show fip vlan-discovery command.....                | 5398       |
| show firewall command.....                          | 4864       |
| show firewall policer command.....                  | 4868       |
| show host command.....                              | 882        |
| show igmp group command.....                        | 4525       |
| show igmp interface command.....                    | 4529       |

|                                              |                                               |                                                  |            |
|----------------------------------------------|-----------------------------------------------|--------------------------------------------------|------------|
| show igmp statistics command.....            | 4533                                          | show multicast pim-to-mld-proxy                  |            |
| show igmp-snooping membership                |                                               | command.....                                     | 4567       |
| command.....                                 | 4536                                          | show multicast route command.....                | 4569       |
| show igmp-snooping route command.....        | 4539                                          | show multicast rpf command.....                  | 4575       |
| show igmp-snooping statistics command.....   | 4541                                          | show multicast scope command.....                | 4579       |
| show igmp-snooping vlans command.....        | 4543                                          | show multicast sessions command.....             | 4581       |
| show interfaces (10-Gigabit Ethernet)        |                                               | show multicast usage command.....                | 4584       |
| command.....                                 | 2217, 2765, 2783                              | show ntp associations command.....               | 892        |
| show interfaces diagnostics optics           |                                               | show ntp status command.....                     | 894        |
| command.....                                 | 883, 2645                                     | show ospf context-identifier command.....        | 4048       |
| show interfaces fabric command.....          | 2651, 2689                                    | show ospf database command.....                  | 4050       |
| show interfaces filters command.....         | 4870                                          | show ospf interfaces command.....                | 4031       |
| show interfaces ge- (Gigabit Ethernet)       |                                               | show ospf neighbors command.....                 | 4031       |
| command.....                                 | 2672                                          | show ospf statistics command.....                | 4031       |
| show interfaces mc-ae.....                   | 2687                                          | show ospf3 database command.....                 | 4058       |
| show interfaces queue                        |                                               | show pfe next-hop command.....                   | 6051       |
| command.....                                 | 2707, 2743, 6015                              | show pfe route command.....                      | 6056       |
| show ipv6 neighbors command.....             | 2969                                          | show pfe terse command.....                      | 6062       |
| show ipv6 router-advertisement command.....  | 2971                                          | show pfe version command.....                    | 6064       |
| show isis adjacency brief command.....       | 3637                                          | show pim bootstrap command.....                  | 4587       |
| show isis adjacency command.....             | 3764                                          | show pim interfaces command.....                 | 4589       |
| show isis adjacency extensive command.....   | 3637                                          | show pim join command.....                       | 4592       |
| explanation.....                             | 3638                                          | show pim neighbors command.....                  | 4601       |
| show isis authentication command.....        | 3768                                          | show pim rps command.....                        | 4605       |
| show isis database command.....              | 3770                                          | show pim source command.....                     | 4612       |
| show isis hostname command.....              | 3777                                          | show pim statistics command.....                 | 4615       |
| show isis interface command.....             | 3779                                          | show policy damping command.....                 | 3615       |
| show isis overview command.....              | 3784                                          | show protocols igmp command.....                 | 4523       |
| show isis route command.....                 | 3787                                          | show rip general-statistics command.....         | 4200       |
| show isis statistics command.....            | 3791                                          | show rip neighbor command.....                   | 4197, 4202 |
| show lacp statistics interfaces command..... | 2806                                          | show rip statistics command.....                 | 4197, 4204 |
| show lldp command.....                       | 1842                                          | show route active-path command.....              | 2979       |
| show lldp local-info command.....            | 1847                                          | show route all command.....                      | 2984       |
| show lldp neighbors command.....             | 1849                                          | show route aspath-regex command.....             | 2986       |
| show lldp statistics command.....            | 1853                                          | show route best command.....                     | 2988       |
| show log                                     |                                               | show route brief command.....                    | 2991       |
| command.....                                 | 889, 1567, 2360, 3689, 6212, 6385, 6437, 6499 | show route command.....                          | 2974       |
| show log messages command .....              | 6391, 6500                                    | show route community command.....                | 2993       |
| show msdp command.....                       | 4545                                          | show route community-name command.....           | 2995       |
| show msdp source command.....                | 4547                                          | show route damping command.....                  | 2997, 3617 |
| show msdp source-active command.....         | 4549                                          | show route detail command.....                   | 2955, 3002 |
| show msdp statistics command.....            | 4552                                          | show route exact command.....                    | 3017       |
| show multicast flow-map command.....         | 4556                                          | show route export command.....                   | 3019       |
| show multicast interface command.....        | 4558                                          | show route extensive command.....                | 3022       |
| show multicast minfo command.....            | 4560                                          | show route flow validation command.....          | 3037       |
| show multicast next-hops command.....        | 4562                                          | show route forwarding-table command.....         | 3039       |
| show multicast pim-to-igmp-proxy             |                                               | show route forwarding-table family fibre-channel |            |
| command.....                                 | 4565                                          | command.....                                     | 5400       |
|                                              |                                               | show route inactive-path command.....            | 3052       |

- show route inactive-prefix command.....3055
- show route instance command.....1855, 3057
- show route label command.....3064
- show route label-switched-path command.....3066
- show route martians command.....3068
- show route next-hop command.....3070
- show route no-community command.....3076
- show route protocol command.....3079
- show route range command.....3091
- show route receive-protocol command.....3095
- show route resolution command.....3103
- show route snooping command.....3106
- show route source-gateway command.....3114
- show route summary command.....3120
- show route table command.....3124
- show route terse command.....2955, 3135
- show sflow collector command.....6442
- show sflow command.....6440
- show sflow interface command.....6443
- show snmp health-monitor command.....6445
- show snmp inform-statistics command.....6450
- show snmp mib
  - command.....1831, 6389, 6452, 6493
- show snmp rmon command.....6455
- show snmp rmon history command.....6459
- show snmp statistics
  - command.....1831, 1859, 6389, 6460, 6493
- show snmp v3 command.....6464
- show spanning-tree bridge command.....2235
- show spanning-tree interface command.....2240
- show spanning-tree mstp configuration
  - command.....2246
- show spanning-tree statistics command.....2248
- show subscribers command.....897
- show system alarms command.....914
- show system audit command.....915
- show system boot-messages command.....923
- show system buffers command.....930
- show system certificate command.....937
- show system commit command.....939, 1649
- show system configuration archival
  - command.....941, 1651
- show system configuration rescue
  - command.....942, 1652
- show system connections command.....944
- show system core-dumps command.....963
- show system directory-usage command.....976
- show system license command.....96, 980
  - explanation.....97
- show system license usage command.....97
  - explanation.....97
- show system processes command.....315, 983, 2627
- show system reboot command.....1010
- show system resource-cleanup processes
  - command.....1013
- show system rollback command.....1015, 1654
- show system services service-deployment
  - command.....1017
- show system software command.....1018
- show system statistics arp command.....2250
- show system statistics
  - command.....331, 372, 373, 939, 1013, 1026, 1067, 1148, 1150, 1631, 1647, 1649, 1656
- show system statistics igmp command.....4628
- show system storage command.....1061
- show system uptime command.....1067
- show system users command.....1072
- show system virtual-memory command.....1077
- show version command.....1135
- show vlans command.....2251
- show bfd session extensive command.....3668
- show isis interface brief command.....3636
- show isis interface detail command.....3636
  - explanation.....3636
- signaled LSPs
  - fate-sharing.....2882
- single-area network, OSPF.....3815
- single-connection statement.....291
  - usage guidelines.....1711
- single-rate statement.....4798
- size statement.....6378
  - archiving of all system log
    - files.....1381, 6366, 6369
  - archiving of individual system log file.....6368
  - system logging
    - usage guidelines.....6220
- SNMP
  - adding client lists and prefix lists.....6196
  - community string.....6194
  - configuration
    - versions 1 and 2.....1703, 6190
  - health monitor alarms, displaying.....6445
  - inform statistics, displaying.....6450
  - limiting interface access.....6197
  - MIB object values, displaying.....6452
  - MIB views.....6197
  - RMON alarms and events, displaying.....6455
  - show commands.....1831, 6389, 6493

|                                            |                 |
|--------------------------------------------|-----------------|
| statistics                                 |                 |
| clearing.....                              | 6403            |
| displaying.....                            | 1859, 6460      |
| system location.....                       | 1779, 6308      |
| tracing operations.....                    | 6393, 6495      |
| trap groups.....                           | 6195            |
| v3.....                                    | 6117            |
| version 3 configuration, displaying.....   | 6464            |
| snmp                                       |                 |
| rmon history.....                          | 6459            |
| snmp (tracing flag)                        |                 |
| class of service.....                      | 5909            |
| snmp history                               |                 |
| clear.....                                 | 6402            |
| SNMP inform statistics, displaying.....    | 6450            |
| SNMP informs.....                          | 6206            |
| snmp statement.....                        | 1810, 6339      |
| usage guidelines                           |                 |
| SNMPv1 and SNMPv2.....                     | 1703, 6190      |
| SNMP traps                                 |                 |
| spoofing.....                              | 6424            |
| snmp-community statement.....              | 6343            |
| SNMPv2                                     |                 |
| Passive Monitoring Traps MIB.....          | 6195            |
| SNMPv3.....                                | 6117            |
| informs, configuring.....                  | 6206            |
| minimum configuration.....                 | 6118            |
| snooping routes, displaying.....           | 3106            |
| sockets, displaying active IP.....         | 944             |
| software.....                              | 67, 1661        |
| version, displaying.....                   | 316, 2628       |
| <i>See also</i> JUNOS software             |                 |
| source filtering.....                      | 4231            |
| source gateway addresses, displaying.....  | 3114            |
| source statement                           |                 |
| commit scripts.....                        | 291, 6253       |
| IGMP.....                                  | 4442            |
| usage guidelines.....                      | 4305            |
| MSDP.....                                  | 4477            |
| op scripts.....                            | 6254            |
| PIM RPF selection.....                     | 4395, 4413      |
| SSM                                        |                 |
| usage guidelines.....                      | 4344            |
| source-active (tracing flag).....          | 4479            |
| source-active-request (tracing flag).....  | 4479            |
| source-active-response (tracing flag)..... | 4479            |
| source-address statement.....              | 6343            |
| NTP                                        |                 |
| usage guidelines.....                      | 181             |
| NTP, RADIUS, system logging, TACACS+.....  | 292             |
| RADIUS                                     |                 |
| usage guidelines.....                      | 1701            |
| RADIUS and TACACS+.....                    | 292             |
| system logging.....                        | 292             |
| usage guidelines                           |                 |
| usage guidelines, RADIUS.....              | 1699            |
| source-count statement                     |                 |
| IGMP.....                                  | 4443            |
| usage guidelines.....                      | 4305            |
| source-increment statement                 |                 |
| IGMP.....                                  | 4443            |
| usage guidelines.....                      | 4305            |
| source-ip statement.....                   | 6272            |
| source-port statement.....                 | 293             |
| usage guidelines.....                      | 151             |
| source-routing statement.....              | 2942            |
| source-specific multicast <i>See</i> SSM   |                 |
| Spanning Tree                              |                 |
| BPDU errors                                |                 |
| clearing.....                              | 2185            |
| speed statement.....                       | 2616            |
| fibrenchannel-options.....                 | 5259            |
| spf (tracing flag)                         |                 |
| IS-IS.....                                 | 3747            |
| OSPF.....                                  | 4023            |
| SPF algorithm                              |                 |
| overview.....                              | 3800            |
| SPF calculations, displaying.....          | 4077            |
| spf-options statement                      |                 |
| OSPF.....                                  | 4018            |
| split horizon technique.....               | 4106            |
| SPT.....                                   | 4225            |
| configuring threshold cutover policy.....  | 4287            |
| cutover control.....                       | 4229            |
| spt-threshold statement.....               | 4414            |
| usage guidelines.....                      | 4287            |
| SRC client information, displaying.....    | 1017            |
| SSB                                        |                 |
| firmware version, displaying.....          | 631             |
| ssh command.....                           | 1863            |
| usage guidelines.....                      | 76              |
| SSH key files.....                         | 158, 1687, 1702 |
| SSH service                                |                 |
| configuring.....                           | 1709, 6178      |
| limiting login attempts.....               | 1716            |

- 
- root login.....1710, 6179
  - SSH protocol version.....1710, 1711, 6179, 6180
  - ssh statement.....1814, 6234
    - usage guidelines.....1709, 6178
  - SSH, opening a connection.....1863
  - ssh-known-hosts statement.....4934
    - usage guidelines.....1707
  - ssh-rsa statement.....294
  - SSM.....4238, 4347
    - configuring.....4342
    - domains.....4344
    - mapping.....4344
  - SSM maps.....4350
    - example.....4351
  - SSM maps for different groups to different
    - sources.....4350
  - ssm-groups statement.....4484
    - usage guidelines.....4347
  - ssm-map statement
    - IGMP.....4444, 4485
      - usage guidelines.....4344
    - MLD
      - usage guidelines.....4344
    - SSM.....4485
      - usage guidelines.....4344
  - ssm-map-policy statement
    - IGMP interface.....4444, 4486
  - stale-routes-time statement.....2318, 3572
    - usage guidelines.....1376, 2282, 2286
  - standards documents
    - SNMP and MIBs.....6125
  - start shell command.....1148
  - start-time statement.....4935
    - system log file archiving.....6368
    - system logging
      - usage guidelines.....6220
  - startup messages, displaying.....923
  - startup period, VRRP.....2298
  - startup-alarm statement.....6344
  - startup-silent-period statement.....2335
    - usage guidelines.....2298
  - state (tracing flag)
    - routing protocols.....2950
  - state-machine-variables (tracing flag)
    - STP.....2138
  - stateless firewall filters
    - accepting Routing Engine traffic from trusted
      - sources
        - example: blocking TCP access.....3436
        - example: blocking Telnet and SSH
          - access.....3441
      - examples
        - blocking TCP access.....3436
        - blocking Telnet and SSH access.....3441
  - statement-name
    - statement.....2097, 2166, 2167, 2169, 4805, 4807, 4824, 4825, 4829, 4830, 4831
  - static routes.....2943
    - BFD.....2831, 2867
    - defining route selection.....2840
    - graceful restart.....2269
    - preferences.....3326
  - static statement.....2943
    - IGMP.....4445
      - usage guidelines.....4305
    - IGMP snooping.....4450, 4457
    - PIM.....4415
      - usage guidelines.....4266
  - static-host-mapping statement.....295
  - static-ip statement.....4827
  - statistics
    - interfaces, real-time.....2635
    - protocol-related, displaying.....1026
  - status command
    - usage guidelines.....82
  - storage space, freeing.....423
  - storing previous configurations.....1591
  - STP
    - bridge
      - displaying.....2235
    - interface
      - displaying.....2240
    - statistics
      - clearing.....2190
      - displaying.....2248
  - stp statement.....2133
  - Structure of Management Information MIB.....6132
  - structured-data statement.....296, 6379
    - usage guidelines.....6215
  - stub areas
    - configuring.....3822
    - description.....3821
    - overview.....3804
  - stub statement.....4020
    - usage guidelines.....3822



|                                               |                        |                                             |                  |
|-----------------------------------------------|------------------------|---------------------------------------------|------------------|
| sub-ASs, BGP.....                             | 3422                   | system logging                              |                  |
| subautonomous systems, BGP.....               | 3422                   | disabling.....                              | 6212             |
| subnet masks.....                             | 49                     | facilities                                  |                  |
| subscriber access                             |                        | alternate for remote machine.....           | 6224             |
| subscriber information, displaying.....       | 897                    | default for remote machine.....             | 6224             |
| subscriber-leave-timer statement.....         | 2945                   | for local machine.....                      | 6222             |
| subscribers                                   |                        | files, archiving.....                       | 6220             |
| displaying.....                               | 897                    | messages, displaying                        |                  |
| summaries statement.....                      | 4021                   | generated by Junos process.....             | 6219             |
| usage guidelines.....                         | 3822                   | structured-data format.....                 | 6216             |
| summary LSA.....                              | 3807                   | regular expression filtering.....           | 6227             |
| summary LSAs                                  |                        | regular expression operators.....           | 6228             |
| advertising LSP metric.....                   | 3908                   | single-chassis system.....                  | 6155             |
| support, technical See technical support      |                        | timestamp, modifying.....                   | 167, 6215        |
| switches                                      |                        | system login.....                           | 149              |
| user accounts.....                            | 1681                   | system overview                             |                  |
| switching                                     |                        | software.....                               | 67, 1661         |
| configuring.....                              | 2048                   | system process information, displaying..... | 316, 2628        |
| symmetric active mode, NTP                    |                        | system services                             |                  |
| configuring.....                              | 154                    | SSH.....                                    | 1709, 6178       |
| defined.....                                  | 52, 53, 153            | system statement.....                       | 299, 1815        |
| syntax conventions.....                       | cxxiii                 | system storage, displaying.....             | 317, 2629        |
| syntax of configuration files, verifying..... | 1150, 1656             | system time, displaying.....                | 316, 2629        |
| sysid statement.....                          | 295                    |                                             |                  |
| syslog.....                                   | 1227, 1372, 6156, 6162 | <b>T</b>                                    |                  |
| syslog statement                              |                        | TACACS+ accounting.....                     | 1714             |
| routing options.....                          | 2909                   | TACACS+ authentication                      |                  |
| system processes.....                         | 297, 1403, 6380, 6382  | configuring.....                            | 1711             |
| usage guidelines.....                         | 6208                   | overview.....                               | 1681             |
| syslog-subtag statement.....                  | 6345                   | tacplus-options statement.....              | 1821             |
| system authentication                         |                        | usage guidelines.....                       | 1713             |
| authentication order.....                     | 1676, 1694             | tacplus-server statement.....               | 305              |
| RADIUS                                        |                        | usage guidelines.....                       | 1711             |
| configuring.....                              | 1699                   | tag statement.....                          | 2946, 6345, 6346 |
| remote template accounts.....                 | 1701                   | tag-list statement.....                     | 6346             |
| TACACS+.....                                  | 1711                   | target-address statement.....               | 6347             |
| System Control Board See SCB                  |                        | target-parameters statement.....            | 6348             |
| system identification, displaying.....        | 316, 2628              | targets statement.....                      | 1822, 6349       |
| system identifier.....                        | 3626                   | usage guidelines.....                       | 6195             |
| system location, SNMP.....                    | 1779, 6308             | task (tracing flag).....                    | 2950             |
| system log.....                               | 1227, 6156             | TCC, graceful restart.....                  | 2269             |
| logging operations.....                       | 6159                   | tcp-mss statement.....                      | 3573             |
| system log file                               |                        | BGP                                         |                  |
| displaying.....                               | 6212, 6385, 6499       | usage guidelines.....                       | 3444             |
| system log messages                           |                        | te-metric statement                         |                  |
| routing protocol process.....                 | 2909                   | usage guidelines.....                       | 3914             |
| System Log MIB.....                           | 6132                   | technical support                           |                  |
|                                               |                        | contacting JTAC.....                        | cxxv             |



- telnet
  - service, limiting login attempts.....1716
- telnet command
  - usage guidelines.....76
- telnet statement.....6235
- template accounts.....1701
- term statement.....4782
- terminal type.....251, 6177
- terms
  - routing policy.....3922
- test configuration command.....1150, 1656
- test msdp command.....4632
- text message on craft interface
  - clearing.....328
  - displaying.....1469
- then statement.....2914, 4783, 4799
- three-color-policer statement.....4800
- threshold
  - BGP.....3574, 3576
  - PIM.....4416, 4417
- threshold statement
  - BFD (BGP)
    - usage guidelines.....3349
  - BGP.....3497
  - forwarding cache.....2947
  - IS-IS.....3697
  - MSDP.....4478
    - usage guidelines.....4332
  - RIP
    - usage guidelines.....4123
- time
  - configuration example.....185
  - security configuration example.....188
- time zone setting, routers.....172
- time, displaying.....1067
- time-format statement.....306, 6383
  - usage guidelines.....167, 6215
- time-zone statement.....308
  - usage guidelines.....172
- timeout statement.....307, 6349
  - authentication
    - usage guidelines, RADIUS.....1699
    - usage guidelines, TACACS+ .....1711
  - flow map.....2948
  - forwarding cache.....2949
- timer (tracing flag).....2950
- timers
  - OSPF.....3875
  - timers (tracing flag)
    - STP.....2138
- timestamp-and-timezone statement.....1821
- to statement.....2914
- top command
  - usage guidelines.....82
- topic1
  - subtopic184,188,184,215,329,202,217,254,857,486,486,494,5923
- topic2
  - subtopic184,188,184,215,329,202,217,254,857,486,486,494,5923
- topologies statement.....3746
- topology
  - sample BGP confederations.....3423
  - sample BGP MED use.....3209
  - sample BGP peer session.....3150
  - sample BGP route reflector (one cluster).....3406
  - sample BGP route reflectors (cluster of clusters).....3407
  - sample BGP route reflectors (multiple clusters).....3406
  - sample distance-vector routing.....4104
  - sample multiarea OSPF routing.....3814
  - sample OSPF network with stubs and NSSAs.....3821
  - sample poison reverse routing.....4107
  - sample split horizon routing.....4106
  - sample unidirectional routing.....4107
- topology statement
  - multitopology routing
    - OSPF.....4022
- topology-change-state-machine (tracing flag)
  - STP.....2138
- totally stubby areas
  - configuring.....3822
  - description.....3821
- trace files
  - logical systems
    - .....3477
- traceoptions
  - statement.2134, 2137, 5260, 5262, 5265, 5267, 5908, 6350
  - BFD
    - usage guidelines.....2835
  - BGP.....3577
    - description.....3476
  - commit scripts.....310
  - graceful restart
    - usage guidelines.....1378, 2284, 2291

|                                     |                  |                          |      |
|-------------------------------------|------------------|--------------------------|------|
| IGMP.....                           | 4446             | csn.....                 | 3747 |
| usage guidelines.....               | 4314             | damping.....             | 3577 |
| IGMP snooping.....                  | 4458             | dynamic                  |      |
| interfaces.....                     | 2618             | class of service.....    | 5908 |
| IS-IS.....                          | 3747             | error                    |      |
| LLDP.....                           | 1823             | IS-IS.....               | 3747 |
| MSDP.....                           | 4479             | OSPF.....                | 4023 |
| usage guidelines.....               | 4328             | RIP.....                 | 4193 |
| multicast snooping.....             | 2619             | events                   |      |
| nonstop active routing.....         | 2950             | STP.....                 | 2138 |
| op scripts.....                     | 310              | fabric                   |      |
| OSPF.....                           | 3954, 4023       | Fibre Channel.....       | 5262 |
| PIM.....                            | 4419             | fc2                      |      |
| usage guidelines.....               | 4248             | Fibre Channel.....       | 5262 |
| RIP.....                            | 4166, 4167, 4193 | fip                      |      |
| routing protocols.....              | 2950             | Fibre Channel.....       | 5262 |
| security.....                       | 4937             | flash.....               | 2950 |
| sFlow technology.....               | 6273             | flogi                    |      |
| SNMP                                |                  | Fibre Channel.....       | 5262 |
| usage guidelines.....               | 6393, 6495       | flooding.....            | 4023 |
| VRRP.....                           | 2336             | forwarding-database      |      |
| traceroute command.....             | 1151             | Fibre Channel.....       | 5262 |
| traceroute monitor command.....     | 1155             | general.....             | 2950 |
| Traceroute page                     |                  | graceful restart         |      |
| results for RIP.....                | 4115             | IS-IS.....               | 3747 |
| tracing.....                        | 312, 6236        | OSPF.....                | 4023 |
| destination-override.....           | 312, 6236        | graft                    |      |
| tracing flags                       |                  | PIM.....                 | 4419 |
| all.....                            | 2137, 2950       | hardware-database        |      |
| class of service.....               | 5908             | class of service.....    | 5908 |
| Fibre Channel.....                  | 5262             | hello                    |      |
| Fibre Channel fip.....              | 5260, 5265       | IS-IS.....               | 3747 |
| Fibre Channel proxy.....            | 5267             | PIM.....                 | 4419 |
| all-failures                        |                  | holddown.....            | 4193 |
| STP.....                            | 2137             | init                     |      |
| as-path.....                        | 3577             | class of service.....    | 5908 |
| assert.....                         | 4419             | interface                |      |
| asynch                              |                  | Fibre Channel.....       | 5262 |
| class of service.....               | 5908             | Fibre Channel proxy..... | 5267 |
| auth.....                           | 4193             | join.....                | 4419 |
| bootstrap.....                      | 4419             | keepalive                |      |
| bpdu.....                           | 2137             | BGP.....                 | 3577 |
| bridge-detection-state-machine..... | 2138             | MSDP.....                | 4479 |
| cache, PIM.....                     | 4419             | kernel.....              | 2950 |
| chassis-scheduler                   |                  | krt                      |      |
| class of service.....               | 5908             | Fibre Channel.....       | 5263 |
| config-internal.....                | 2950             | leave                    |      |
| cos-adjustment                      |                  | IGMP.....                | 4446 |
| class of service.....               | 5908             |                          |      |

|                                     |                  |  |
|-------------------------------------|------------------|--|
| lib                                 |                  |  |
| Fibre Channel.....                  | 5263             |  |
| lif                                 |                  |  |
| Fibre Channel.....                  | 5263             |  |
| lsp.....                            | 3747             |  |
| lsp-generation.....                 | 3747             |  |
| mt.....                             | 4419             |  |
| mtrace                              |                  |  |
| IGMP.....                           | 4314             |  |
| normal.....                         | 2950             |  |
| nsr-synchronization.....            | 4420             |  |
| packet-dump.....                    | 4023             |  |
| packets                             |                  |  |
| BGP.....                            | 3577             |  |
| IGMP.....                           | 4446             |  |
| IS-IS.....                          | 3747             |  |
| OSPF.....                           | 4023             |  |
| PIM.....                            | 4420             |  |
| RIP.....                            | 4193             |  |
| parse.....                          | 2950             |  |
| performance-monitor                 |                  |  |
| class of service.....               | 5909             |  |
| policy.....                         | 2950             |  |
| port-information-state-machine..... | 2138             |  |
| port-migration-state-machine.....   | 2138             |  |
| port-receive-state-machine          |                  |  |
| STP.....                            | 2138             |  |
| port-role-select-state-machine      |                  |  |
| STP.....                            | 2138             |  |
| port-role-transit-state-machine     |                  |  |
| STP.....                            | 2138             |  |
| port-state-transit-state-machine    |                  |  |
| STP.....                            | 2138             |  |
| port-transmit-state-machine         |                  |  |
| STP.....                            | 2138             |  |
| ppmd                                |                  |  |
| STP.....                            | 2138             |  |
| process                             |                  |  |
| class of service.....               | 5909             |  |
| prune                               |                  |  |
| PIM.....                            | 4420             |  |
| psn.....                            | 3747             |  |
| regex-parse.....                    | 2950             |  |
| register.....                       | 4420             |  |
| report                              |                  |  |
| IGMP.....                           | 4447             |  |
| restart                             |                  |  |
| class of service.....               | 5909             |  |
| route                               |                  |  |
| MSDP.....                           | 4479             |  |
| routing.....                        | 2950             |  |
| route-socket                        |                  |  |
| class of service.....               | 5909             |  |
| rp.....                             | 4420             |  |
| show                                |                  |  |
| class of service.....               | 5909             |  |
| snmp                                |                  |  |
| class of service.....               | 5909             |  |
| source-active.....                  | 4479             |  |
| source-active-request.....          | 4479             |  |
| source-active-response.....         | 4479             |  |
| spf                                 |                  |  |
| IS-IS.....                          | 3747             |  |
| OSPF.....                           | 4023             |  |
| state                               |                  |  |
| routing protocols.....              | 2950             |  |
| state-machine-variables             |                  |  |
| STP.....                            | 2138             |  |
| task.....                           | 2950             |  |
| timer.....                          | 2950             |  |
| timers                              |                  |  |
| STP.....                            | 2138             |  |
| topology-change-state-machine       |                  |  |
| STP.....                            | 2138             |  |
| trigger.....                        | 4193             |  |
| update                              |                  |  |
| RIP.....                            | 4193             |  |
| util                                |                  |  |
| class of service.....               | 5909             |  |
| vswitch                             |                  |  |
| Fibre Channel.....                  | 5263             |  |
| tracing IP multicast path           |                  |  |
| from receiver to source.....        | 4512             |  |
| from router to gateway.....         | 4520             |  |
| from server to router.....          | 4515             |  |
| tracing operations.....             | 6081             |  |
| BGP.....                            | 3476, 3577       |  |
| IGMP.....                           | 4314, 4446       |  |
| IS-IS.....                          | 3747             |  |
| MSDP.....                           | 4328, 4479       |  |
| OSPF.....                           | 3954, 4023       |  |
| PIM.....                            | 4419             |  |
| RIP.....                            | 4166, 4167, 4193 |  |
| routing protocols.....              | 2950             |  |
| SNMP.....                           | 6393, 6495       |  |
| tracing ospf                        |                  |  |
| configuring.....                    | 3956             |  |

|                                              |            |                                      |                        |
|----------------------------------------------|------------|--------------------------------------|------------------------|
| tracing routes                               |            | transient configuration changes      |                        |
| from the receiver to the source.....         | 4512       | compared to persistent changes.....  | 6091                   |
| from the source to the gateway router.....   | 4520       | overview.....                        | 6090                   |
| from the source to the receiver.....         | 4515       | transit areas                        |                        |
| monitoring.....                              | 4518       | overview.....                        | 3805                   |
| track statement.....                         | 2338       | transit-delay statement.....         | 4029                   |
| usage guidelines.....                        | 2301, 2302 | usage guidelines.....                | 3876                   |
| traffic                                      |            | translational cross-connect See TCC  |                        |
| controlling with incoming RIP metric.....    | 4143       | transmit-interval                    |                        |
| controlling with outgoing RIP metric.....    | 4144       | BGP.....                             | 3580                   |
| traffic control                              |            | PIM.....                             | 4418                   |
| OSPF.....                                    | 3861       | transmit-interval statement.....     | 4028                   |
| traffic engineering                          |            | BFD.....                             | 2867                   |
| advertising LSP metric in summary            |            | BGP.....                             | 3497                   |
| LSAs.....                                    | 3908       | IS-IS.....                           | 3697                   |
| database credibility value.....              | 3908       | transmit-rate statement.....         | 5911                   |
| enabling.....                                | 3910       | trap groups, SNMP.....               | 6195                   |
| ignoring RSVP LSP metrics.....               | 3908       | trap-group statement.....            | 1827, 6352             |
| igp shortcuts.....                           | 3908       | usage guidelines.....                | 6195                   |
| metric.....                                  | 3908, 3914 | trap-options statement.....          | 1828, 6353             |
| ospf protocol preference.....                | 3908       | traps                                |                        |
| support.....                                 | 3908       | spoofing.....                        | 6424                   |
| traffic engineering database                 |            | traps statement.....                 | 2621                   |
| OSPF support.....                            | 4026       | trigger (tracing flag).....          | 4193                   |
| traffic, real-time monitoring.....           | 6405       | Trivial Network Protocol family      |                        |
| traffic-control-profiles statement.....      | 5910       | interface addresses.....             | 2040                   |
| traffic-engineering statement                |            | troubleshooting                      |                        |
| IS-IS.....                                   | 3750       | general.....                         | 6473, 6485             |
| lsp-metric-info-summary                      |            | resources.....                       | 6471, 6483             |
| usage guidelines.....                        | 3910       | root password recovery.....          | 1159, 1717, 6478, 6524 |
| OSPF.....                                    | 4026       | trusted-key statement.....           | 313                    |
| OSPF passive TE mode                         |            | usage guidelines.....                | 152, 153               |
| usage guidelines.....                        | 3917       | TTY, displaying.....                 | 317, 2629              |
| shortcuts                                    |            | two-rate statement.....              | 4801                   |
| usage guidelines.....                        | 3910       | type statement.....                  | 6354                   |
| usage guidelines.....                        | 3910       | auxiliary port.....                  | 6177                   |
| transfer interval                            |            | console port.....                    | 6177                   |
| usage guidelines.....                        | 1611       |                                      |                        |
| transfer-interval statement                  |            | U                                    |                        |
| archiving of configuration.....              | 1618, 1825 | udp-port statement.....              | 6274                   |
| system log file archiving.....               | 6368       | uid statement.....                   | 313                    |
| system logging                               |            | usage guidelines.....                | 1681, 1692             |
| usage guidelines.....                        | 6220       | UIDs.....                            | 1682                   |
| transfer-on-commit statement.....            | 1619, 1826 | unit statement                       |                        |
| usage guidelines.....                        | 1612       | class of service.....                | 5914                   |
| transferring router configuration to archive |            | interfaces.....                      | 2622                   |
| site.....                                    | 1611       | UNIX-level shell, creating.....      | 1148                   |
|                                              |            | unknown-unicast-forwarding statement |                        |
|                                              |            | rate limiting.....                   | 2140                   |

up command  
     usage guidelines.....82  
 update (tracing flag)  
     RIP.....4193  
 update command  
     usage guidelines.....82  
 update messages  
     BGP.....3148  
 update-interval statement  
     RIP.....4196  
     usage guidelines.....4160, 4161  
 upgrading software  
     performing.....394, 408, 410, 1461, 2393  
 upstream-interface statement.....2953  
 uptime, displaying.....1067  
 URLs, specifying in commands.....73  
 user (system logging facility).....6223  
     option to facility-override statement.....6225  
 user access  
     login classes.....1683  
     user accounts.....1681, 1692  
 user accounts  
     configuring.....1681, 1692  
     security configuration example.....1744  
     shared user accounts.....1701  
 user authentication  
     methods for.....1681  
 user data, erasing.....433  
 user identifiers See UIDs  
 user permission flags.....1672  
 user statement  
     access.....314, 1829  
     usage guidelines.....1681, 1692  
     SNMPv3.....6355  
     system logging.....6384  
     usage guidelines.....6211  
 username  
     displaying.....317, 2629  
 users  
     CLI permissions, displaying.....876  
     displaying.....317, 2629  
     logged in, displaying.....1072  
     logging users out.....384  
     logs, displaying.....889, 1567, 2360, 6437  
     messages, displaying for.....372  
 usm statement.....6356  
 util (tracing flag)  
     class of service.....5909  
 Utility MIB.....6116

## V

v3 statement.....6358  
 vacm statement.....6360  
 validating software.....420  
 var/log/mib2d file.....6393, 6495  
 var/log/snmpd file.....6393, 6495  
 variable statement.....6361  
 vendor-id statement.....4832  
 verification  
     active licenses.....96  
     BFD for IS-IS.....3663  
     BGP session flap prevention.....3455  
     license usage.....97  
     licenses .....96  
     multicast topology for IS-IS.....3675  
     network interfaces.....3299, 3399  
     RIP host reachability.....4115  
     static route.....2839, 2844, 2852  
     tracing.....3482  
 verifying syntax of configuration file.....1150, 1656  
 version  
     BGP.....3581  
     firmware, displaying.....631  
     software, displaying.....316, 2628  
     general.....1135  
 version statement  
     BFD.....2867, 4422  
         usage guidelines.....2831  
     BFD (BGP)  
         usage guidelines.....3349  
     BGP.....3497  
     IGMP.....4448  
         usage guidelines.....4298  
     IS-IS.....3697  
     OSPF.....3971  
     PIM.....4423  
         usage  
             guidelines.....4246, 4255, 4266, 4290  
     RIP.....4175  
         usage guidelines.....4123  
     SNMP.....1830, 6362  
         usage guidelines.....6195  
 View Events page  
     field summary (filtering log messages).....2955  
 view statement  
     SNMP (associating with community).....6364  
         usage guidelines.....6194  
     SNMP (configuring MIB view).....6363  
         usage guidelines.....6197

|                                              |                        |  |
|----------------------------------------------|------------------------|--|
| views, MIB                                   |                        |  |
| SNMP.....                                    | 6197                   |  |
| virtual chassis configuration                |                        |  |
| procedures.....                              | 1308                   |  |
| virtual links                                |                        |  |
| overview.....                                | 3803                   |  |
| virtual memory, displaying.....              | 1077                   |  |
| virtual-address statement.....               | 2339                   |  |
| usage guidelines.....                        | 2296                   |  |
| VLAN MIB.....                                | 6133                   |  |
| vlan statement.....                          | 4833, 4834             |  |
| IGMP snooping.....                           | 4460                   |  |
| interfaces.....                              | 2141                   |  |
| MSTI.....                                    | 2142                   |  |
| port mirroring.....                          | 4919                   |  |
| rate limiting.....                           | 2143                   |  |
| vlan-id statement.....                       | 2144, 2154, 2156, 2157 |  |
| vlan-range statement.....                    | 2144                   |  |
| VLANs                                        |                        |  |
| configuring.....                             | 2048, 2145, 2158, 2170 |  |
| configuring VLAN range.....                  | 2144                   |  |
| vlangs statement.....                        | 2145, 2158, 2170       |  |
| vpls protocol family                         |                        |  |
| interface addresses.....                     | 2040                   |  |
| VPNs                                         |                        |  |
| BGP                                          |                        |  |
| preventing session flaps.....                | 3449                   |  |
| VPNs, graceful restart.....                  | 2269                   |  |
| VRRP                                         |                        |  |
| advertisement interval.....                  | 2299                   |  |
| authentication.....                          | 2297                   |  |
| basic configuration.....                     | 2296                   |  |
| group                                        |                        |  |
| inheritance.....                             | 2305                   |  |
| passive ARP learning.....                    | 2304                   |  |
| tracking logical interface status.....       | 2302                   |  |
| tracking route status.....                   | 2301                   |  |
| vrrp                                         |                        |  |
| failover-delay.....                          | 2326                   |  |
| vrrp-group statement.....                    | 2340                   |  |
| usage guidelines.....                        | 2296                   |  |
| vswitch (tracing flag)                       |                        |  |
| Fibre Channel.....                           | 5263                   |  |
| <b>W</b>                                     |                        |  |
| warning (system logging severity level)..... | 2909                   |  |
| wide-metrics-only statement.....             | 3753                   |  |
| wildcard-source statement                    |                        |  |
| PIM RPF selection.....                       | 4424                   |  |
| working directory                            |                        |  |
| current, displaying.....                     | 880                    |  |
| world-readable statement                     |                        |  |
| archiving of all system log files.....       | 6366                   |  |
| archiving of individual system log file..... | 6368                   |  |
| system logging                               |                        |  |
| usage guidelines.....                        | 6220                   |  |
| write-interval statement.....                | 4835                   |  |
| write-view statement.....                    | 6364                   |  |
| <b>X</b>                                     |                        |  |
| xe (Port) statement.....                     | 2625                   |  |
| xle (Port) statement.....                    | 2626                   |  |
| <b>Y</b>                                     |                        |  |
| yellow alarm condition.....                  | 147                    |  |
| yellow alarms See minor alarms               |                        |  |